

Orbit Reconstruction of Triton from Earth-Based Astrometric Observations

Dynamical Parameter Sensitivity, Pole Estimation,
and Weighting Strategy Assessment

MSc Thesis

Atanas Dzhurkov

Orbit Reconstruction of Triton from Earth-Based Astrometric Observations

Dynamical Parameter Sensitivity, Pole Estimation, and Weighting Strategy Assessment

by

Atanas Dzhurkov

Student Number: 5217512

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on 11 May 2026, at 13:00.

Supervisor: Dominic Dirx
Thesis committee: Stefano Speretta,
Marc Rovira Navarro,
Jonas Hener

Faculty of Aerospace Engineering
Delft University of Technology

Preface

This thesis was carried out at the Spaceflight Dynamics section of the Faculty of Aerospace Engineering at Delft University of Technology.

I would like to thank my supervisor, Dominic Dirx, for his guidance throughout this project. His consistently clear and structured approach to both the research itself and the inevitable problems along the way made a real difference. Every meeting left me with a sharper understanding of what to do next and, more importantly, why.

I am also grateful to my friends and family for their support and for the many conversations about orbits, moons, and uncertainty that they endured with genuine curiosity.

Working on this thesis has been, more than anything, an exercise in understanding how much I do not understand. It is tempting to trust numbers that come from real data, but the honest question is always: how well can I actually quantify what I claim to know? The most valuable lesson from this work has been to keep asking that question, to look carefully for where the truth is, and to be forthright about where the limits of that knowledge lie.

Atanas Dzhurkov
Delft, Monday 27th April, 2026

Abstract

This study investigates the orbit reconstruction of Triton, Neptune's largest moon, using exclusively Earth-based astrometric observations spanning 1963 to 2025. The estimation is performed within the Tudat software framework, employing a weighted least-squares procedure to estimate Triton's initial state and Neptune's pole orientation parameters. The JPL NEP097 ephemeris kernel serves as the benchmark for assessing solution accuracy.

A sensitivity analysis using simulated observations identifies Neptune's pole model parameters and gravitational zonal harmonics as the dynamical parameters to which Triton's orbit is most sensitive. An estimation including Triton's initial state, Neptune's pole position, and pole libration parameters achieves sub-kilometre agreement with the NEP097 kernel, demonstrating that the implemented dynamics are sufficient to reproduce the reference ephemeris within the estimation timespan.

When applied to real astrometric data, the inclusion of pole libration parameters as estimated quantities reduces the cross-track difference with NEP097 to below 100 km over the modern observational arc (1990 to 2025), while producing formal errors that closely follow the actual solution discrepancy. The pole libration declination correction of $+1.19^\circ$ is found to be statistically significant at 6.3σ .

Three base and two hybrid observation weighting strategies are derived, implemented, and compared. The conventional per-file weighting scheme, believed to be standard practice in the literature, is found to produce overconfident formal errors. A proposed scaled per-file weighting strategy, which applies per-timeframe deweighting while assigning weights at the file level, produces formal errors that are consistently calibrated with the actual solution accuracy across two independent initialisations.

The principal limitations of this work are the absence of Voyager 2 data and the inability to reliably estimate the system gravitational parameters. Recommendations for future work include the incorporation of additional observation types and a formal consider parameter analysis.

Contents

Preface	iii
Abstract	iv
1 Introduction	1
1.1 Background and Motivation	1
1.2 Problem Statement	1
1.3 Research Questions	2
1.4 Thesis Outline	2
2 Theoretical Background	3
2.1 Astronomical Data & Observation Models	3
2.1.1 Observation Vector	4
2.1.2 Observables	5
2.1.3 Space Observatories Data	7
2.1.4 Voyager 2 Data	9
2.1.5 Error Sources for Observations	11
2.1.6 Coordinate and Time Systems	12
2.2 Dynamics of Planetary Systems	15
2.2.1 Natural Satellite Dynamics	15
2.2.2 Non-inertial origins	16
2.2.3 Spherical Harmonic Gravity and Inner Moons	17
2.2.4 Neptune's Pole Precession	17
2.2.5 Tidal effects on Triton and Neptune	19
2.2.6 Numerical Integration Methods and Error Sources	21
2.3 Natural Satellite Orbit Determination	23
2.3.1 Weighted Least Squares	23
2.3.2 Covariance of the Estimated State	24
2.3.3 A priori constraints	25
2.3.4 Considered Parameters	25
2.3.5 Considerations & Coupled Parameters	26
2.4 Modern Research on the Neptunian System	27
2.4.1 Ephemerides and Purpose of Study	27
2.4.2 Summary of Modern Research	27
2.4.3 Conclusion and Potential Research Questions	30
3 Methodology	34
3.1 Dynamical Model	34
3.1.1 Neptune Gravitational Field Zonal Harmonics	34
3.1.2 Pole Precession Model	34
3.1.3 Numerical Integrator	35
3.2 Estimation Framework	35
3.2.1 Weighted Least Squares Formulation	35
3.2.2 Weighting Strategy	35
3.2.3 Observation Uncertainty from Weights	37
3.2.4 Observational Goodness-of-Fit Metrics	37
3.3 Astrometric Dataset and Initial Residual Analysis	39
3.3.1 O-C Residuals with respect to the NEP097 Kernel	39
3.4 Structure of Analyses	39
3.4.1 Simulated Observations Analysis	43

3.4.2	Pole Estimation with Astrometric Data Analysis	43
3.4.3	Weight Strategy Analysis	44
4	Parameter Estimation with Simulated Observations	45
4.1	Initial Propagation vs. NEP097	47
4.2	Initial State Estimation	48
4.3	Gravitational Parameter Estimations of Neptune and Triton	50
4.4	High Level Analysis of Best Performing Estimations	53
4.5	IAU Pole Parameter Estimations	55
4.5.1	RMS Analysis	55
4.5.2	Parameter Correlations and Conditioning	57
4.5.3	Residual Timeseries Analysis	57
4.5.4	Parameter Updates and Physical Interpretation	59
4.5.5	Pole Model Conclusion	62
4.6	Extended Body Gravity Estimations of Neptune	63
4.7	Conclusion	66
5	Pole Estimation with Astrometric Data	67
5.1	Single Value Fit Metrics	67
5.2	RSW Statistics with Respect to NEP097	71
5.3	Correlation Plots	72
5.4	RSW Timeseries Analysis	74
5.5	Parameter Updates	77
5.6	Pole Movement	79
5.7	Final Solution	80
6	Weighting Strategy Comparison Across Estimation Analyses	84
6.1	Comparison of the Two Analyses	85
6.2	Analysis WA-FIT-PL	88
6.2.1	Weights of Observations	88
6.2.2	RSW Statistics	91
6.2.3	RSW Timeseries Analysis	92
6.3	Conclusion	97
7	Comparison with Published Studies and Limitations	99
8	Conclusions and Recommendations	101
8.1	Answers to Research Questions	101
8.2	Key Contributions	101
8.3	Recommendations for Future Work	104
A	Initial Research Questions and Mapping	108
B	Integrator Step Size Validation	110
C	Additional Figures for Weight Strategy Comparison Across Estimation Analyses	112
C.1	Hybrid Weighting Scheme Uncertainties	112
C.1.1	Geometric Mean Hybrid	112
C.1.2	Arithmetic Mean Hybrid	113
C.2	Additional Uncertainties figures for WA-FIT-PL	113
C.2.1	Per File Uncertainty	113
C.2.2	Per Timeframe Analysis	114
C.2.3	RMS Update Per File	115
C.3	Analysis WA-IAU	117
C.3.1	RSW Statistics	117
C.3.2	RSW Timeseries Analysis	118
C.3.3	Per-File Uncertainty	121
C.3.4	Per-Timeframe Analysis	123
C.3.5	RMS Update Per File	124

List of Figures

2.1	Observation Vector of a Planet (figure from Emelyanov (2021))	4
2.2	Position of Topocenter, Planet and Satellite (figure from Emelyanov (2021))	5
2.3	Construction of Micrometric Observables	7
2.4	Residual distribution in the X - Y coordinate system and after transformation to the T - S coordinate system aligned with Gaia's scan direction. (figure from Emelyanov et al. (2023))	10
2.5	Schematic representation of the tidal deformation (solid red ellipsoid) of body i (left) at time t due to the tidal potential (dashed red ellipsoid) raised by point mass body j (right) at time t' . The rotation of body i and orbital motion of body j during the time lag $\Delta t = t - t'$ are also represented. (Figure from Fayolle (2025))	20
3.1		40
3.2	O-C Residuals NEP097 colored per file id, rejected observations are colored red and marked with x.	42
3.3	O-C Residuals with and without manual bias overlaid	42
4.1	RSW Positional Difference with NEP097 kernel for initial propagations with Neptune pole models IAU 2015 and Jacobson as described in Table 4.2.	47
4.2	RSW Difference of the initial propagations with different Neptune pole models IAU and Jacobson as described in Table 4.2.	48
4.3	RSW Difference with NEP097 of initial state estimation with simulated observations for both pole models (naming convention defined in Table 4.2)	49
4.4	Parameter Update of Initial State Estimation	50
4.5	Correlation heatmaps for the initial state estimations comparing the (a) IAU 2015 and (b) Jacobson 2009 Neptune pole models.	50
4.6	RSW difference with NEP097: Statistical comparison of estimations involving gravitational parameters.	52
4.7	RSW difference with NEP097: Timeseries comparison of all GM estimation variants.	52
4.8	Comparison of parameter updates for the Triton initial state and gravitational parameters across GM variants.	53
4.9	Correlation heatmaps for state estimations including (a) Triton's gravitational parameter and (b) Neptune's gravitational parameter.	53
4.10	RMS difference relative to NEP097 for estimations including initial state, pole parameters, and Neptune's extended body gravity (SH).	54
4.11	RMS difference relative to NEP097 decomposed into RSW components for estimations including initial state, pole parameters, and Neptune's extended body gravity (SH).	55
4.12	Total RMS difference relative to NEP097 for estimations including initial state and pole parameters.	56
4.13	RSW component decomposition relative to NEP097 for estimations including initial state and pole parameters.	56
4.14	Correlation heatmaps for separate pole position and libration estimations.	57
4.15	Correlation heatmaps for the combined IAU 2015 pole estimations, comparing the standard combined model (a) against the full model including rotation rates (b).	58
4.16	RSW Position difference with NEP097 for the baseline and top-performing pole estimations.	58
4.17	RSW Positional difference with NEP097 comparison between the <code>state + pos. + lib.</code> and <code>state + full variants</code> .	59
4.18	Estimated state updates for IAU pole cases.	60
4.19	RSW decomposition of estimated state updates.	60
4.20	Estimated pole position (α_0, δ_0) updates.	61
4.21	Estimated pole libration updates.	61
4.22	Neptune pole position evolution relative to initial IAU 2015 values.	62
4.23	Total RMS difference with NEP097: Comparison between the <code>state + SH. variant</code> , the <code>baseline state</code> estimation, and the <code>state + pos. + lib. model</code> .	63

4.24 RSW component RMS difference with NEP097: Comparison between state + SH., baseline state, and state + pos. + lib.	64
4.25 RSW difference timeseries with NEP097: Comparison of the state + SH. estimation against the selected state + pos. + lib. model.	65
4.26 Correlation heatmap for the state + SH. estimation variant.	65
5.1 RMS of the difference relative to NEP097 for pole estimations using astrometric data.	69
5.2 RMS of formal errors for the pole estimation variants.	69
5.3 Ratio of the RMS difference relative to NEP097 to the RMS of the formal errors.	69
5.4 Goodness-of-fit metrics derived from observation residuals for the pole estimation variants.	70
5.5 Observation residuals categorised by data source for the state+lib. (Fit.) estimation.	71
5.6 RSW statistics of the difference with respect to NEP097, formal errors, and their ratio for the pole estimation variants using astrometric data.	72
5.7 Correlations between estimated parameters for FitPole estimation for pole position and pole libration rate	73
5.8 Correlations between estimated parameters for FitPole estimation for pole libration rate	74
5.9 Correlations between estimated parameters for FitPole estimation for initial state only	74
5.10 RSW Difference with NEP097 for estimations state and state (Fit.)	75
5.11 RSW Difference with NEP097 for estimations state + lib. and state + lib. (Fit.)	76
5.12 RSW Formal Errors for estimations state + lib. and state + lib. (Fit.)	76
5.13 RSW-decomposed initial position and velocity update of Triton for state+lib. under both IAUPole and FitPole initialisations, and the benchmark state estimation.	77
5.14 Total initial position and velocity update of Triton for state+lib. under both IAUPole and FitPole initialisations, and the benchmark state estimation.	78
5.15 Neptune pole trajectory for state + lib. and state + lib. (Fit.) compared to the IAUPole and FitPole nominal trajectories.	80
5.16 RSW difference with NEP097 and propagated formal errors (1σ) for the final estimation state + lib. (Fit.).	82
6.1 Comparison of RMS derived true and formal errors for the two weighting strategy analyses	86
6.2 Comparison of the observational goodness of fit metrics for the weighting strategy analyses.	87
6.3 Uncertainties derived from weights of all observations for the per file weighting scheme.	89
6.4 Uncertainties derived from weights of all observations for the scaled per file weighting scheme.	89
6.5 Ratio of number of observations to number of timeframes (top) and total number of observations (bottom) per file.	89
6.6 Uncertainties derived from weights of all observations for the per timeframe weighting scheme.	90
6.7 Uncertainties derived from weights of all observations for the per timeframe free weighting scheme.	90
6.8 Uncertainties derived from weights of all observations for the scaled hybrid arith. weighting scheme.	91
6.9 Uncertainties derived from weights of all observations for the scaled hybrid geom. weighting scheme.	91
6.10 RMS difference with NEP097 and formal errors per RSW direction	92
6.11 Comparison of RSW position differences with respect to NEP097 (left) and formal errors (right) for per file and scaled per file weighting schemes. From top to bottom: radial (R), along-track (S), and cross-track (W) components. The full time range (1960–2025) is shown.	94
6.12 Comparison of RSW position differences with respect to NEP097 (left) and formal errors (right) for scaled per file and per timeframe weighting schemes. From top to bottom: radial (R), along-track (S), and cross-track (W) components. The full time range (1960–2025) is shown.	95
6.13 Comparison of RSW position differences with respect to NEP097 (left) and formal errors (right) for per timeframe and per timeframe free weighting schemes. From top to bottom: radial (R), along-track (S), and cross-track (W) components. The full time range (1960–2025) is shown.	96
6.14 RSW position difference with respect to NEP097 for scaled hybrid geom. and scaled hybrid arith. weighting schemes.	97
6.15 RSW formal errors for scaled hybrid geom. and scaled hybrid arith. weighting schemes.	97
B.1 Absolute position difference between propagations with consecutive step sizes as a function of time. The transition from truncation-error-dominated behaviour (upper curves) to numerical-noise-dominated behaviour (1 800 to 900 s pair) is clearly visible.	111

B.2	Accumulated position difference at the end of the simulation arc as a function of integrator step size. The flattening below 1 800 s indicates the onset of floating-point precision loss.	111
C.1	Per File weight uncertainty (defined in chapter 3) in RA (top) and Dec (bottom) for base weighting schemes: ID, ID ν 2, and TF	114
C.2	Per Timeframe Uncertainty for observation data file 874 nm0004 in RA (top) and Dec (bottom) for base weighting schemes: ID, ID ν 2, and TF. Grey bars indicate the number of observations per file (right axis)	115
C.3	Per-file RMS update (initial – final) in RA (top) and Dec (bottom) for base weighting schemes ID, ID ν 2, and TF. Grey bars indicate the number of observations per file (right axis).	116
C.4	Per-file RMS update (initial – final) in RA (top) and Dec (bottom) for hybrid weighting schemes hybrid G+ ν 2 and hybrid A+ ν 2, compared against ID ν 2 and TF. Layout as in Figure C.3.	117
C.5	RMS difference with NEP097, formal error RMS, and their ratio per RSW direction for all weighting schemes under the WA-IAU initialisation.	118
C.6	RSW position difference with respect to NEP097 for base weighting schemes ID, ID ν 2, and TF under the WA-IAU initialisation. Left panels show the full time range (1960–2025); right panels show a zoomed view in early 1963. From top to bottom: radial (R), along-track (S), and cross-track (W) components.	119
C.7	RSW formal errors for base weighting schemes ID, ID ν 2, and TF under the WA-IAU initialisation. Layout as in Figure C.6. From top to bottom: σ_R , σ_S , and σ_W	119
C.8	RSW position difference with respect to NEP097 for hybrid weighting schemes hybrid G+ ν 2, hybrid A+ ν 2, and TF, compared against ID ν 2, under the WA-IAU initialisation. Layout as in Figure C.6.	120
C.9	RSW formal errors for hybrid weighting schemes hybrid G+ ν 2, hybrid A+ ν 2, and TF, compared against ID ν 2, under the WA-IAU initialisation. Layout as in Figure C.6.	121
C.10	Per-file weight uncertainty in RA (top) and Dec (bottom) for base weighting schemes ID, ID ν 2, and TF under the WA-IAU initialisation.	122
C.11	Number of observations over number of timeframes ratio (top) and number of observations (bottom) per file under the WA-IAU initialisation.	123
C.12	Per-timeframe uncertainty for observation file 874 nm0004 in RA (top) and Dec (bottom) for base weighting schemes ID, ID ν 2, and TF under the WA-IAU initialisation. Grey bars indicate the number of observations per timeframe (right axis).	124
C.13	Per-file RMS update (initial – final) in RA (top) and Dec (bottom) for base weighting schemes ID, ID ν 2, and TF under the WA-IAU initialisation. Grey bars indicate the number of observations per file (right axis).	125
C.14	Per-file RMS update (initial – final) in RA (top) and Dec (bottom) for hybrid weighting schemes hybrid G+ ν 2 and hybrid A+ ν 2, compared against ID ν 2 and TF, under the WA-IAU initialisation. Layout as in Figure C.13.	126

List of Tables

2.1	Triton Observation (Part 1)	13
2.2	Triton Observations (Part 2)	14
2.3	Overview of Subtopics with Estimated Benefit and Difficulty	33
3.1	Neptune gravitational field zonal harmonic coefficients from Jacobson (2009).	34
3.2	Neptune pole model parameter values from Archinal et al. (2018) and Jacobson (2009).	35
3.3	Summary of goodness-of-fit metrics used in this work.	39
3.4	Summary of astrometric observation datasets used in this work. NSDC Listing gives the dataset identifier in the Natural Satellites Data Centre archive. MPC Code is the three-digit Minor Planet Center observatory code. N_{obs} is the total number of astrometric observations. Obs. Type indicates whether the original measurements are relative (Rel.) or absolute (Abs.) astrometry. RMS O–C RA and RMS O–C Dec are the root-mean-square observed minus computed residuals in right ascension and declination, respectively, evaluated against the NEP097 ephemeris [″].	41
3.5	Observations rejected by the outlier filter, grouped by dataset identifier.	41
4.1	Initial values of all estimated parameters used in the analysis. The Triton initial state is taken from the NEP097 ephemeris at 2006-10-01. Gravitational parameters (GM) and spherical-harmonic coefficients (\bar{C}_{20} , \bar{C}_{40}) are from the DE440 SPICE kernel and the Jacobson (2009), respectively; these values are identical for both rotation models. Pole parameters differ between the IAU 2015 (Archinal et al. (2018)) and Jacobson (2009) models. Degree-2 libration terms (α_2 , δ_2) are Jacobson (2009) specific and have no IAU 2015 counterpart.	45
4.2	Estimations naming conventions used in figures for Simulated Observations Analysis.	46
4.3	Estimated gravitational parameters for the GM simulation variants (IAU 2015 rotation model). Initial values are taken from the DE440 SPICE kernel. $\Delta = \text{Final} - \text{Initial}$; $\Delta[\%] = \Delta / \text{Initial} \times 100$.	53
4.4	Condition numbers of the correlation matrices for IAU 2015 variants.	57
4.5	Comparison of estimated updates versus formal uncertainties from Jacobson (2009).	59
4.6	Final estimated parameter values for best IAU 2015 variants.	62
4.7	Summary of spherical harmonics estimation for the <code>state+SH</code> variant. $\kappa(\mathbf{C})$ denotes the condition number of the correlation matrix; $\Delta\bar{C}_{20}$ and $\Delta\bar{C}_{40}$ represent the estimated updates to the normalized gravitational coefficients.	63
5.1	Weight-scheme naming conventions used in figures for Pole Estimation Analysis.	68
5.2	Initial parameter values for the IAU and FitPole rotation models. $\Delta = p_{\text{FitPole},0} - p_{\text{IAU},0}$ in the listed unit; $\Delta[\%] = \Delta / p_{\text{IAU},0} \times 100$.	68
5.3	Key parameter estimates for selected simulations. IAU_0 and FP_0 are the initial values for the IAU 2015 and NEP097-fitted pole models respectively. $\Delta_{\text{IAU}} = \text{Est.} - \text{IAU}_0$; $\Delta_{\text{FP}} = \text{Est.} - \text{FP}_0$.	79
5.4	RMS observed minus computed (O–C) residuals per observation dataset using post-estimation residuals (<code>state + lib. (Fit.)</code>). MPC Code is the three-digit Minor Planet Center observatory code. N_{obs} is the number of astrometric observations in the dataset. Obs. Type indicates relative (Rel.) or absolute (Abs.) astrometry. All RMS values in arcseconds [″].	81
5.5	Estimated parameters for simulation <code>state + lib. (Fit.)</code> . IAU_0 and FP_0 are the initial values for the IAU 2015 and NEP097-fitted pole models, respectively. $\pm 1\sigma$ is the formal error from the post-fit covariance matrix (square root of the diagonal element). $\Delta_{\text{IAU}} = \text{Final} - \text{IAU}_0$; $\Delta_{\text{FP}} = \text{Final} - \text{FP}_0$.	82
6.1	Naming convention for the two weighting strategy analyses.	85
6.2	Weight-scheme naming conventions used in figures for <code>WA-FIT-PL</code> , analogues for <code>WA-IUA</code> without <code>ref</code> , derivation of weighting schemes presented in chapter 3.	85
A.1	Mapping between the original research subtopics and the final research questions. Fully addressed questions are marked with \checkmark , partially addressed with \sim , and unaddressed with — .	109

1

Introduction

1.1. Background and Motivation

The study of the Solar System provides valuable insights into the formation of planets and the conditions that support life. Among the outer planets, Neptune is often overlooked in popular science, yet it continues to yield surprising discoveries. Neptune was the first planet to be predicted mathematically ([Galle and d'Arrest, 1846](#)), and only weeks later William Lassell discovered Triton, Neptune's largest moon ([Lassell, 1846](#)). Triton's unusual retrograde orbit indicates that it was likely captured from the Kuiper Belt rather than forming in situ. Its captured origin, suspected geological activity, and potential subsurface ocean make the Neptunian system a valuable analogue for studying planetary migration and satellite evolution.

Despite the scientific interest in the system, it has only been visited once by Voyager 2 ([Smith et al., 1989](#)). As future missions are in relatively early stages and travel times are prohibitively long, direct exploration remains limited. Nevertheless, active research has improved our understanding of tidal mechanics ([Lainey, 2016](#)), geological phenomena on icy satellites ([Sohl et al., 2002](#)), and the potential for life in subsurface oceans ([Chyba and Phillips, 2002](#)). Advancements in astrometric instrumentation and a renewed interest in the Neptunian system have also generated a significant amount of high-quality observational data.

In the absence of new flyby data, researchers have relied on dynamical modelling and orbit fitting techniques to study Triton's motion and to constrain the physical parameters of the system. In these methods, dynamical parameters are systematically adjusted so that simulated orbits best reproduce observational data through a least-squares fitting procedure ([Montenbruck and Gill, 2000](#)). The parameters of interest include, among others, the initial state of Triton, Neptune's pole orientation, gravitational field coefficients, and tidal parameters. Tidal parameters, in particular, provide constraints on Triton's internal structure and may indicate the presence of a subsurface ocean ([Tyler, 2008](#)).

The least-squares framework also provides the capability to identify and correct systematic errors in the observational data. This is especially valuable for older datasets, which can extend the temporal baseline of the solution significantly but also exhibit large biases and random errors. Successful efforts have been made to improve the accuracy of photographic plates through newer, more accurate star catalogues ([Robert et al., 2025](#)), which has greatly increased their impact on dynamical studies. Older micrometric observations, however, remain difficult to incorporate, as they are reported in apparent coordinates and require proper modelling of Earth rotation and atmospheric refraction ([Urban and Seidelmann, 2012](#)).

1.2. Problem Statement

While research on Triton's orbit is steadily growing, several gaps exist in the current literature. The most detailed dynamical estimation of the Neptunian system to date was performed by [Jacobson \(2009\)](#) using closed-source software, and ambiguity about the procedures makes replication of the results difficult. Since 2009, the volume of observational data for the Neptunian system has increased and dynamical modelling techniques have advanced. Yet most recent studies either focus on constraining the initial state of Triton to fit new data while neglecting the effects of other relevant dynamical parameters ([Tang et al., 2020](#)), or remain vague about their evaluation of data accuracy and the implementation of their estimation procedures ([Wang et al., 2025](#)).

The latest JPL solution for the orbit of Triton is provided in the NEP097 SPICE kernel. To the author's knowledge, no supporting publication exists that documents the solution generation process, its accuracy, the dataset used, or its limitations. The most recent JPL study on the topic is [Jacobson \(2009\)](#), from which the Neptune pole model and spherical harmonic coefficients do not appear to align with the solution presented in the NEP097 kernel, as discussed in [chapter 4](#).

This study aims to address these gaps by providing a transparent, open-source analysis of the orbit reconstruction of Triton, using the Tudat software framework developed at Delft University of Technology. Building on previous work by [Dahmani \(2024\)](#) and [Fayolle \(2025\)](#), the analysis incorporates Earth-based astrometric data spanning 1963 to 2025 and investigates the estimation of dynamical parameters beyond Triton's initial state. The NEP097 kernel serves as the benchmark against which all solutions are assessed. Particular attention is given to two aspects that are typically treated briefly or without justification in the existing literature: the sensitivity of the solution to Neptune's pole orientation parameters, and the design and assessment of observation weighting strategies.

1.3. Research Questions

Based on the objectives outlined above, this study addresses the following research questions:

- RQ1:** Which dynamical parameters of the Neptune system does Triton's orbit show the greatest sensitivity to, and can the dynamical model implemented in Tudat reproduce the NEP097 kernel?
- RQ2:** How does the inclusion of Neptune's pole model parameters as estimated quantities affect the accuracy of Triton's orbit reconstruction from Earth-based astrometric data?
- RQ3:** What weighting strategy for astrometric datasets produces formal errors that are statistically consistent with the actual solution accuracy?
- RQ4:** What are the principal limitations of Triton orbit reconstruction using exclusively Earth-based astrometric data, and which factors drive the remaining discrepancies with the NEP097 kernel?

The mapping between these research questions and the broader set of subtopics identified during the initial phase of this study is provided in [Appendix A](#).

1.4. Thesis Outline

The theoretical background of orbit estimation for solar system bodies is presented in [chapter 2](#), with a focus on the dynamical and observational peculiarities of Triton's orbit, as well as an overview of current research and identified gaps. The specific implementation choices, derivations of the weighting strategies, and the overall structure of the estimation analyses are documented in [chapter 3](#). A summary of the observational dataset used in this study, including the outlier rejection algorithm and a bias correction applied to a particular dataset, is presented in ??.

The results chapters follow the logical progression of the analyses. The Parameter Estimation with Simulated Observations ([chapter 4](#)) assesses the sensitivity of Triton's orbit to various dynamical parameters, and the estimation that most closely reproduces the benchmark NEP097 kernel is selected for further investigation with real astrometric data. As Neptune's pole model parameters are found to be the most influential, the Pole Estimation with Astrometric Data Analysis is presented in [chapter 5](#), where combinations of pole parameter estimations are compared based on their agreement with NEP097, the coherence of their formal errors, and the physical plausibility of the parameter updates. The Weighting Strategy Analysis ([chapter 6](#)), which is in circular dependence with the pole estimation chapter, uses the best-performing solution from [chapter 5](#) to compare the different weighting strategies derived in [chapter 3](#). A comparison with published studies and a discussion of the limitations of this work are presented in [chapter 7](#). Finally, the Conclusions and Recommendations ([chapter 8](#)) summarise the answers to the research questions, the key contributions of this study, and recommendations for future work.

2

Theoretical Background

The aim of this chapter is to provide a comprehensive overview of the principles, methods, and current practices in the orbit estimation of natural satellites. By consolidating theoretical foundations, observational techniques, and recent research developments, this review establishes the necessary background for the analysis and research questions addressed in later chapters.

The study is organized into four main sections. The first section, *Astronomical Data and Observation Models*, outlines the types of observational data used in orbit determination. Topics include the definition of the observation vector and common observables, special considerations for historical Earth-based measurements, data from space observatories and spacecraft imaging, common error sources, and differences in the employed coordinate systems.

The second section, *Dynamics of Planetary Systems*, reviews the dynamical models required for accurate orbit propagation. It covers the equations of motion for natural satellites, non-inertial reference frames, spherical harmonic representations of gravitational fields, pole precession effects, tidal interactions, and the implementation of numerical integration, with attention to error sources and stability.

The third section, *Natural Satellite Orbit Determination*, focuses on the estimation process itself. It introduces the weighted least-squares method, discusses the covariance of the estimated state, describes the use of a priori constraints and consider parameters, and ends with general considerations for practical implementation and possible difficulties with strongly coupled parameters.

The final section, *Modern Research on the Neptunian System*, reviews recent studies on Neptune's satellite system as a case study. It begins with a discussion of commonly used ephemerides and their purpose, followed by a synthesis of recent research results, and concludes by identifying gaps in the literature and formulating potential research questions.

Together, these sections provide a structured foundation for understanding the methods and challenges in natural satellite orbit estimation, while highlighting specific opportunities for further investigation in the Neptunian system.

2.1. Astronomical Data & Observation Models

Systematic observation of celestial objects has been practised since antiquity. Early efforts focused on recording the rising and setting of the Sun and Moon, along with the apparent motions of stars and planets. Over time, advances in technology and in the understanding of Earth's motion and celestial mechanics have driven a continual demand for increasingly precise measurements.

Most of the astronomical observation data used in this study is available through the Natural Satellite DataBase maintained by the ([Institute of Celestial Mechanics and Computation of Ephemerides \(IMCCE\), 2025](https://nsdb.imcce.fr/obspos/bnepomae.htm))¹. A chronological summary of the Earth-based observations, along with Gaia astrometry, is presented in [Table 2.1](#) and [Table 2.2](#). This listing excludes observations obtained by the Hubble Space Telescope (HST) and Voyager 2, which are provided separately by [Showalter et al. \(2019\)](#) and [Jacobson \(1991\)](#), respectively.

¹<https://nsdb.imcce.fr/obspos/bnepomae.htm>

The purpose of this section is to outline the construction of the Observation Model. It begins with the formulation of the observation vector, followed by a discussion of the measured quantities (observables). Special considerations for specific data types, such as micrometric observations, space-based observatories, and spacecraft imaging (Voyager 2), are summarized with references for further detail. The most common sources of observational error are then reviewed, with particular emphasis on the limitations of historical Earth-based data. Finally, a summary of the spatial and temporal coordinate systems used in these datasets is provided. Each dataset is distributed in the coordinate system most convenient for its producers; however, these reference frames are not always optimal for use in orbit determination (e.g., the time system employed by Gaia or the Earth rotation reference frame used for older observations).

2.1.1. Observation Vector

To construct the observation vector some key positions need to be defined. The barycenter of the Solar system, the topocenter of the Earth, and the observed planet or natural satellite are denoted by $\mathbf{B}(t)$, $\mathbf{T}(t)$ and $\mathbf{P}(t)$. The calculations of these vectors are performed by specific models. It is assumed that the errors introduced by the models of the barycenter and the Earth are negligibly small, while the model of the observed planet or natural satellite is part of the fitting process which will be explained in later sections.

Considering the most simple of cases, the observation of a planet in the Solar System on the surface of the Earth. The time of observation is t_0 . Then the planet observation vector is denoted \mathbf{P}_T , as the observation is made on the topocenter of the Earth. The moment at which the photons departed the observed planet is denoted by t_1 . Then the observation vector can be constructed as in [Figure 2.1](#). The observation vector equation is then

$$\mathbf{P}_T = \mathbf{P}(t_1) - \mathbf{T}(t_0) \quad (2.1)$$

As the speed of light is not infinite an auxiliary equation, the light time, is needed and formulated as

$$t_0 - t_1 = \frac{|\mathbf{P}_T|}{c}, \quad (2.2)$$

where c is the speed of light. The two equations can be solved by iteration when initially setting the times equal as $t_1 = t_0$.

This construction of the observation vector is with respect to the barycentre of the Solar System. The book of [Emelyanov \(2021\)](#) presents also a formulation with respect to the centre of the Sun. There consideration of the motion of the Sun is necessary. The mathematics will not be shown here as they are quite involved and would seldom be required. The reader is advised to read the book for further information.




Figure 2.1: Observation Vector of a Planet (figure from [Emelyanov \(2021\)](#))

In the situation when both a planet and a satellite of the planet are observed the construction of the observation vector is more involved. Let's denote the satellite observation vector by $\mathbf{S}_T^{(1)}$ and the planetocentric vector of the satellite, that is the relative vector of the satellite to the planet by $\mathbf{S}_P^{(1)}$. Then \mathbf{S}_T can be constructed as

$$\mathbf{S}_T^{(1)} = \mathbf{S}_P^{(1)}(t_1) + \mathbf{P}(t_1) - \mathbf{T}(t_0) \quad (2.3)$$

$$t_0 - t_1 = \frac{|\mathbf{S}_T^{(1)}|}{c} \quad (2.4)$$

where t_0 is the moment the photons hit the detector and t_1 is the moment the photons depart from the observed satellite S . As the distance between the topocentre to the planet is different from the distance from the topocentre to the satellite, then the moment the photons depart from the planet is different, denoted by t_2 . The observation vector of the planet becomes

$$\mathbf{P}_T = \mathbf{P}(t_2) - \mathbf{T}(t_0) \quad (2.5)$$

$$t_0 - t_2 = \frac{|\mathbf{P}_T|}{c} \quad (2.6)$$

The construction of the vectors can be seen in [Figure 2.2](#). To find the times t_1 and t_2 the same procedure is used as previously shown.




Figure 2.2: Position of Topocenter, Planet and Satellite (figure from [Emelyanov \(2021\)](#))

[Emelyanov \(2021\)](#) also discusses the simultaneous observation of two satellites of a planet, mutual events and star occultations. At this stage this is outside the scope of this study the reader is referred to the book instead.

2.1.2. Observables

Mathematical models of the motion of planets and their satellites make it possible to determine their rectangular position at any point in time. When dealing with optical observations from Earth or space, the observables are typically angles either relative to the planet the object orbits or to the celestial sphere. Therefore to connect the observation vectors discussed previously to the observables a conversion from rectangular position to the specific measured angles needs to be performed. It is necessary to have precise knowledge of the position at which the observation is made. When dealing with Earth based observations, the topocentric position of the measuring station is necessary, while for space observatories and spacecraft data an accurate dynamical model of the spacecraft is necessary.

The angular coordinates are measured in the geoequatorial coordinate system associated with the equator and ecliptic of a certain epoch. In modern work, the equatorial coordinate system refers to the equator and ecliptic of the J2000 epoch, which is not based on the rotation of the Earth, but rather to the International Celestial Reference Frame (ICRF). This was not the case for publications of the past decade, where mean equator and equinox coordinates were used and the precession and nutation of the Earth's axis were simulated.

In this section it is assumed that the axes of all considered rectangular coordinate systems are parallel to the axes of some non-rotating geoequatorial system (ICRF), if not stated otherwise. Therefore the angular coordinates determined from the observation vector are called astrometric coordinates.

As the observables are angles it is useful to use spherical coordinates. The latitude in the geoequatorial system is assigned declination and the longitude - right ascension. For any rectangular observation vector x, y, z the right ascension α and the declination δ are determined by the equation

$$\tan \alpha = \frac{y}{x}, \quad \tan \delta = \frac{z}{\sqrt{x^2 + y^2}}, \quad (2.7)$$

which are also called absolute coordinates and can be measured quantities.

The celestial coordinate difference between two bodies is called relative coordinates. For any two bodies the difference is

$$\Delta\alpha = \alpha_2 - \alpha_1, \quad \Delta\delta = \delta_2 - \delta_1 \quad (2.8)$$

Due to the converges of the declination angle close to the poles, the sensitivity to error in right ascension changes with the change in declination. To account for this the equations

$$X_d = (\alpha_2 - \alpha_1) \cos \delta_1, \quad Y_d = \delta_2 - \delta_1 \quad (2.9)$$

are more commonly used as measured quantities, as X_d and Y_d are decoupled and in same scale. Therefore errors can be assessed simultaneously, but also separately if necessary.

The computation of absolute coordinates through the observation vector is straightforward and already defined by [Equation 2.7](#). When dealing with relative coordinates the subtraction of the topocentric angular coordinates of bodies consists in subtracting tow close numbers, which leads to some loss of accuracy. Therefore the planetocentric coordinates (that is the planet that the observed satellites orbit) should be used instead.

The first satellite has observation vector with components X, Y, Z as

$$\{X, Y, Z\} = \mathbf{S}_T^{(1)} \quad (2.10)$$

Then the difference in rectangular coordinates between two satellites is

$$\{\Delta x, \Delta y, \Delta z\} = \mathbf{S}_T^{(2)} - \mathbf{S}_T^{(1)} = \mathbf{S}_P^{(2)}(t_2) - \mathbf{S}_P^{(1)}(t_1) + \mathbf{P}(t_2) - \mathbf{P}(t_1) \quad (2.11)$$

The exact derivations of the relative coordinates $\Delta\alpha$ and $\Delta\delta$ from the observation vectors above are described in detail in [Emelyanov \(1999\)](#). Approximate solutions are also presented if necessary.

[Emelyanov \(2021\)](#) mentions the tangential coordinates X_t, Y_t . However, the description is limited and no supporting figure is provided. Since no observations of the Neptunian system in tangential coordinates are known to the present study, they are not included in the analysis.

Another type of observable is the meridian transit circle, for which at least one data set exists for Triton. [Jacobson \(2009\)](#) briefly mentions the inclusion of such data in their fitting process, but provides limited methodological detail. Fully incorporating these observations into the present study would require developing the reduction procedure from first principles, with significant time investment. Given the limited expected impact on the final results, these observations are not included here. Readers interested in their application are referred to [Jacobson \(2009\)](#) and related sources.

As discussed previously, mutual events and stellar occultations differ significantly from astrometric angle measurements in both the construction of the observation vector and the nature of the observables. While their accuracy can exceed that of any other Earth-based data, their limited number and the substantial effort required for implementation have led to their exclusion from the present analysis. Instead, the focus of this study is on the inclusion of older Earth-based observations, which, despite their lower precision, extend the temporal coverage of the dataset and improve long-term orbital constraints. Alongside these, space observatory measurements (e.g., Hubble, Gaia) and spacecraft imaging data (e.g., Voyager 2) are prioritized, as they offer high-accuracy information with distinct modelling considerations.

Micrometric Observations

In historical astrometric practice, particularly during the 19th and early 20th centuries, observational data were often published in the form of apparent coordinates. These coordinates are expressed relative to the true equator and equinox of date, including corrections for precession, nutation, and aberration, but typically excluding atmospheric refraction.

Micrometric observations were a prevalent method before the adoption of photographic and CCD techniques. In such observations, positional measurements were made using an instrument-mounted micrometer to record separations s and position angles p between celestial objects (for example, between a planet and its satellites or between two stars). To convert these micrometer readings into celestial coordinates, astronomers determined scale and orientation constants by observing known star pairs are used. This process is functionally analogous to deriving plate constants in photographic astrometry, albeit typically less precise due to manual measurement limitations. Therefore, the position angle and separation were usually not measured at the same time. The equations of the angles are

$$p = \tan^{-1} \left(\frac{\Delta\alpha \cdot \cos \delta_0}{\delta - \delta_0} \right), \quad s = \sqrt{(\Delta\alpha \cdot \cos \delta_0)^2 + (\delta - \delta_0)^2}, \quad \text{Wanget al. (2023)} \quad (2.12)$$

where δ_0 and δ are the declination of Neptune and Triton, respectively. While α_0 and α are the right ascension of these two bodies, $\Delta\alpha = \alpha - \alpha_0$ is the difference in right ascension between them. The value of p is in degrees, while for s , the value is measured in arcseconds.

Figure 2.3 presents the construction of the observables as described by Urban and Seidelmann (2012). The angular distance vector s is constructed between the Planet body and the satellite, while p is the angle between s and the direction of increasing declination.




Figure 2.3: Construction of Micrometric Observables

2.1.3. Space Observatories Data

Hubble Space Telescope

Showalter et al. (2019) obtained 40 relative astrometric measurements of Triton using the Hubble Space Telescope. The primary focus of the study was the inner moons of Neptune, and as a result, most observations were scheduled so that Triton remained outside the field of view. This may have been done to avoid potential issues related to Triton's brightness relative to the faint inner satellites.

The following passages outline the observation program and data processing steps as described in the paper. While some technical details may be more relevant to the inner moons, they are included here for completeness, with particular attention to points relevant to Triton.

The reference epoch used in the analysis is midnight 2009 January 1 UTC, selected as it falls near the midpoint of all observations. In barycentric dynamical time (TDB), this corresponds to 284,040,066.184 seconds after the J2000 epoch. The relevant Hubble observation programs are GO 10398, GO 11656, and GO 14217.

Neptune's tracking accuracy by HST is assumed to be sufficient for each HST orbit. Images taken with the same filter during a single orbit share a common navigation solution. However, changing filters alters the optical path,

introducing shifts of up to 0.5 pixels; each filter was therefore treated independently. For this program, the visual-band filters CLEAR, F606W, and F350LP were used. The CLEAR filter on ACS/HRC exhibited a plate-scale error, which the authors corrected by applying a 0.9987 scale factor derived empirically.

The observing sequences were designed so that Neptune was placed near the center of the images, using the ultraviolet filter F330W. This centering was part of the image acquisition strategy, rather than a determination of HST's pointing, which is controlled by its internal guidance system. For astrometric reduction, Neptune itself was not used as a reference object due to its large apparent size, frequent saturation, and unstable centroid. Background stars were generally absent from the field of view. As a result, initial relative alignment of the images was performed manually by identifying the brightest moons by eye, achieving a precision of 1–2 pixels. These alignments were later refined through automated processing.

For each detected body, a model point-spread function (PSF) was fitted to a small region around it, with additional processing applied to mitigate Neptune's glare. The final centroid positions (x, y) were obtained via nonlinear least-squares fitting. Uncertainties were initially estimated from a linear model around the best-fit solution, producing a covariance matrix. In some cases, the formal errors were unrealistically small; to address this, the authors imposed a minimum uncertainty of $\sigma = 0.1$ pixels.

The paper provides a CSV file containing the processed measurements, including times in TDB, pixel coordinates, relative astrometric positions, and their uncertainties. Additional columns include Neptune's pixel and astrometric position, the reference image used for navigation, and a column of navigational measurements. The purpose of the latter is not explained, but it may be relevant for assessing HST pointing accuracy.

In the context of natural satellite orbit determination, HST astrometry is structurally similar to ground-based CCD observations but typically provides higher precision, making it a valuable constraint in orbit fitting.

([Mikulski Archive for Space Telescopes \(MAST\)](#), 2025)² serves as the primary repository for space telescope data. However, it does not currently support direct search queries for Solar System objects, which makes the retrieval of potentially relevant observations a labour-intensive process.

Gaia Data Release 3

This section outlines the construction of Gaia's observables and the associated covariance matrix, as these differ significantly from other astrometric data sources.

Gaia's measurement precision varies strongly along and across its scanning direction, leading to a highly anisotropic error distribution in the focal plane. This asymmetry also results in a characteristic pattern of residuals that must be taken into account in the processing of Gaia astrometry.

All Gaia observations are time-stamped in the Barycentric Coordinate Time (TCB) scale. Since many orbit determination frameworks operate in Barycentric Dynamical Time (TDB) or Coordinated Universal Time (UTC), careful conversion between time scales is required. [Tanga et al. \(2023\)](#) provides the conversion from TCB to TDB as

$$\text{TDB} = \text{TCB} - L_B(\text{JD}_{\text{TCB}} - 2\,443\,144.500\,3725) \times 86400 \text{ s} - 6.55 \times 10^{-5} \text{ s}, \quad (2.13)$$

where the time is expressed in seconds, and $L_B = 1.550\,519\,768 \times 10^{-8}$ is a defining constant in the astronomical system of units.

Gaia observations are grouped into batches, each corresponding to a single transit of the satellite across the Gaia focal plane. A transit typically consists of no more than nine position measurements, spaced at approximately 40 s intervals.

The data provided on ([Institute of Celestial Mechanics and Computation of Ephemerides \(IMCCE\)](#), 2025)³ however, the first recorded batch for Triton contains 15 measurements — a number that exceeds the stated maximum in the literature. The reason for this discrepancy is unclear from the documentation and may be related to the grouping procedure or data formatting.

As Gaia operates with two telescopes, consecutive transits can be separated by as little as 106 minutes or as much as several days, depending on the scanning geometry.

²<https://archive.stsci.edu/>

³<https://nsdb.imcce.fr/obspos/bnepomae.htm>

The astrometric observables provided by Gaia are Gaia-centric right ascension (α) and declination (δ), expressed in the Barycentric Celestial Reference System (BCRS) with the origin at the Gaia spacecraft. This is achieved by referencing Gaia's attitude to the Gaia-CRF3 axes, which are aligned with the ICRF3 [Gaia Collaboration et al. \(2022\)](#).

These positions are corrected for annual aberration but do not include corrections for relativistic light deflection caused by the gravitational field of the Solar System. A procedure for applying such corrections is provided by [Emelyanov et al. \(2023\)](#). The position of Gaia itself at the time of each observation is also provided in the ICRF, allowing transformation of the observables to other reference frames or correction for additional effects.

The observation equations are constructed by first calculating the residuals between the measured and computed positions as

$$X = \Delta\alpha \cos \delta, \quad Y = \Delta\delta \quad (2.14)$$

where δ is the computed declination. These residuals represent the differences between observed and model-predicted values, although the exact dynamical model used for the predictions is not specified in [Emelyanov et al. \(2023\)](#).

The axes X and Y shown in Figure 2.4 correspond to residual offsets in right ascension and declination, respectively. In principle, these should align with the ICRF-based coordinate frame, as the differences are expressed in equatorial coordinates. However, since the residuals are model-dependent, there is some ambiguity in the precise orientation of X and Y with respect to the Gaia focal-plane frame.

The position angle P between the scan direction S and the Y -axis is provided in the Gaia data. This allows the scan/across-scan components of the residuals to be obtained by applying a rotation by angle θ . The interpretation here assumes that the X - Y axes are sufficiently close to the ICRF orientation for this transformation to be meaningful, though the details of the underlying model alignment remain unstated.

For Gaia data to be effectively used in a fitting procedure, the observations must be assigned appropriate weights. Because Gaia's measurement accuracy differs along and across the scan direction, the weighting for these two directions must be handled separately. The covariance matrix of the observation errors is initially provided in the X - Y frame, where X and Y are highly correlated [Tanga et al. \(2023\)](#).

To simplify the weighting, the covariance matrix is rotated into the T - S frame (aligned with Gaia's along-scan and across-scan directions), where it becomes diagonal, eliminating covariance between the two directions. This transformation is outlined in [Emelyanov et al. \(2023\)](#) and the mathematics are described in more detail in [Deakin \(2005\)](#).

The covariance matrix includes contributions from both random and systematic errors, denoted by subscripts r and s respectively. The combined maximum (a) and minimum (b) error components are defined as:

$$a = \sqrt{a_r^2 + a_s^2}, \quad b = \sqrt{b_r^2 + b_s^2}. \quad (2.15)$$

[Emelyanov et al. \(2023\)](#) applies these a and b values to determine the weighting coefficients used in the least-squares fit. While the paper does not explicitly detail how a and b are computed prior to fitting, they are treated as known quantities for the weighting process.

[Tanga et al. \(2023\)](#) provides an extensive discussion of potential Gaia error sources and the structure of the covariance matrix, although no explicit numerical values are given.

Observation data from Gaia is available via the [Institute of Celestial Mechanics and Computation of Ephemerides \(IMCCE\) \(2025\)](#) and is summarized in [Table 2.2](#).

2.1.4. Voyager 2 Data

Voyager 2 optical images of Neptune's largest satellites, Triton and Nereid, were processed by [Jacobson \(1991\)](#) to derive astrometric positions. The relevant images are those originally acquired for spacecraft navigation. These were taken with the Voyager narrow-angle camera, which recorded one or more Neptunian satellites against a star background. This combination provided the necessary information to determine both the satellites' astrometric positions and the pointing direction of the camera.

The Voyager narrow-angle camera is a vidicon imaging system, a predecessor of the modern CCD camera. The instrument itself is described in detail by [Smith et al. \(1977\)](#).




Figure 2.4: Residual distribution in the X – Y coordinate system and after transformation to the T – S coordinate system aligned with Gaia’s scan direction. (figure from [Emelyanov et al. \(2023\)](#))

[Jacobson \(1991\)](#) outlines the full procedure for reducing the raw image data to astrometric positions. Here, only a brief summary of the process will be presented, with the focus on the key steps relevant to this study.

In total, 413 images of Triton were obtained between 13 July 1988 and 24 September 1989, covering distances from approximately 2.5 to 590 million kilometres. For Nereid, 83 images were processed over the same general time period.

Although the reduction from raw images to astrometric coordinates is complex, the resulting dataset is sufficiently accurate for use in modern orbit determination. Re-reducing the raw images would require substantial effort and detailed knowledge of the original instrument calibration, for only marginal improvement given the available time. For this reason, the current study will make direct use of the astrometric positions provided by [Jacobson \(1991\)](#).

A brief summary of the reduction approach is included here for completeness.

The surface of the Voyager narrow-angle camera sensor is approximately 1 cm^2 , divided into 800×800 pixels. The pixel and line coordinates (p, l) correspond to the horizontal and vertical axes of the detector, and are related to the focal plane coordinates (x, y) by

$$\begin{bmatrix} p \\ l \end{bmatrix} = \begin{bmatrix} K_x & K_{xy} & K_{xxy} \\ K_{yx} & K_y & K_{yyx} \end{bmatrix} \begin{bmatrix} x \\ y \\ xy \end{bmatrix} + \begin{bmatrix} p_0 \\ l_0 \end{bmatrix} \quad (2.16)$$

The matrix \mathbf{K} contains the scale (pixels per mm), rotation, and potential non-orthogonality of pixel and line axes. The constant term shifts the coordinate origin to the image centre. Solutions for x and y are given explicitly in [Jacobson \(1991\)](#).

The focal plane coordinates are related to the unit pointing vector $\hat{\mathbf{P}}$ in the camera body frame through a gnomonic projection:

$$\begin{bmatrix} x \\ y \end{bmatrix} = \frac{f}{P_3} \begin{bmatrix} P_1 \\ P_2 \end{bmatrix} \quad (2.17)$$

where f is the focal length of the camera. The pointing vector in inertial space, $\hat{\mathbf{T}}'$, is then obtained by

$$\hat{\mathbf{T}}' = TITV^T \hat{\mathbf{P}} \quad (2.18)$$

where $TITV$ is the rotation matrix from inertial space to the camera body frame. Due to stellar aberration, $\hat{\mathbf{T}}'$ corresponds to the apparent, not true, position of the object.

The apparent inertial coordinates of the object are then

$$\alpha' = \arctan \frac{\hat{T}'_2}{\hat{T}'_1}, \quad \delta' = \arcsin \hat{T}'_3. \quad (2.19)$$

A further transformation from the inertial frame to the Sun–Canopus coordinate system is performed to estimate the uncertainties in these apparent angles.

The inertial observation vector of a satellite is given by

$$\mathbf{T} = \mathbf{S}(t - \tau) + \mathbf{B}(t - \tau) - \mathbf{R}(t) - \mathbf{B}(t) \quad (2.20)$$

where t is the observation time, τ is the light-time from satellite to spacecraft, \mathbf{S} is the satellite position relative to the planetary system barycentre, \mathbf{R} is the spacecraft position relative to the system barycentre, and \mathbf{B} is the planetary system barycentre position relative to the Solar System barycentre. This formulation accounts for stellar aberration due to spacecraft motion.

For convenience, Equation 2.20 can be expressed in the planet-barycentric frame by defining the spacecraft position relative to the planet barycentre as \mathbf{R}' . The expression then reduces to

$$\mathbf{T} = \mathbf{S}(t - \tau) - \mathbf{R}'(t, \tau), \quad (2.21)$$

analogous to Equation 2.1 with $t = t_0$ and $t - \tau = t_1$, except that all vectors are planet-barycentric and the motion of the barycentre must be included in \mathbf{R}' .

Tables 1 and 2 in Jacobson (1991) provide the final astrographic positions of Triton and Nereid. Each entry lists: image ID, observation time in Julian Ephemeris Date (timescale not specified), Voyager-centric right ascension and declination in the FK4/B1950 frame (stellar aberration corrected), and uncertainties.

Ancillary tables (3 and 4) include: light-time τ from satellite to spacecraft, spacecraft position relative to Neptune barycentre in FK4/B1950 coordinates, and the right ascension/declination of the asymptotic approach direction of Voyager's trajectory. The asymptote direction may be useful in estimating spacecraft positional uncertainties, as the trajectory is approximated by a hyperbola.

Jacobson (1991) compared fits using the astrographic reductions to fits using the raw optical navigation measurements. Including the asymptote direction as an additional estimated parameter altered the final orbits by at most ~ 2.5 km for Triton and ~ 40 km for Nereid, both negligible for the purposes of this study.

The FK4/B1950 frame used in the reductions is a potential source of error. Jacobson (1991) improved the reductions by transforming reference star positions to the ICRF using the UCAC2 catalogue. Further improvements could be obtained using the Gaia star catalogue, but this would require a full re-reduction of the raw Voyager images, which is beyond the scope of the present study.

2.1.5. Error Sources for Observations

All observational datasets contain errors. In the context of parameter estimation, these errors can manifest as Gaussian random errors, systematic biases in the measured quantities, or biases in the recorded time of observation. A realistic assessment of observational accuracy is essential, as it directly informs the weight matrix in the fitting procedure (discussed in detail in the fitting section).

The exact physical origin of an error is not always critical for fitting itself. However, understanding potential sources is useful for deciding whether to model them explicitly (e.g., refraction corrections) or to account for them through statistical bias and covariance adjustments.

All Earth-based observations are affected by atmospheric distortion, primarily due to refraction. While modern refraction models correct most of this effect, older or site-specific corrections can be inaccurate. Jacobson (1991), for example, notes mismatches for Bordeaux Observatory observations.

Absolute astrometric accuracy also depends on the reference star catalogue used. Prior to Gaia, catalogue errors were a significant source of systematic bias. Reprocessing older observations with Gaia DR2 or DR3 can substantially improve their accuracy, this process will be discussed in more detail in the section on older observations.

For photographic plates and CCD data, measured image coordinates must be transformed into celestial coordinates, usually by solving for plate constants. These parameters describe the rotation, scale, and possible optical distortion in the image.

Ideally, plate constants are determined separately for each image using multiple reference stars from a high-accuracy catalogue (Gaia DR3 being ideal).

If too few reference stars are available, or no catalogue was used, alternative calibration methods can be applied, but at the cost of reduced precision, especially if atmospheric or instrumental conditions varied between exposures.

The limitations described above become more significant for older Earth-based datasets, where catalogue quality, calibration methods, and environmental control were all less advanced. The following section addresses these older datasets, which require more extensive preprocessing before use in modern orbit determination.

Considerations for Older Earth Based Observations

Older datasets, whether from early photographic plates or micrometric measurements, require special treatment.

Older photographic plate data can often be improved significantly through digitization and recalibration against modern catalogues. For example, [Robert et al. \(2025\)](#) scanned and re-reduced large numbers of historic plates against Gaia DR3, reducing their systematic errors to near the intrinsic limit of the photographic process. This improvement is enabled by precision scanning techniques and Gaia's exceptional astrometric accuracy.

Micrometric measurements present greater challenges. These were not tied to star catalogues but instead relied on accurate determination of True North at the observation epoch. While modern Earth rotation models can reconstruct this orientation more precisely, the historical rotation model used is rarely documented, and the time standard applied is often unclear. Refraction corrections were also generally absent and must be applied retrospectively.

Some micrometric datasets, for instance [Hall \(1877\)](#), were reduced to the mean heliocentric distance of Neptune. While convenient for comparing results across epochs, this approach is inconsistent with modern modelling and requires re-reduction. Additionally, the manual nature of micrometer readings introduces the possibility of errors in scale and rotation constants, which must be accounted for in any modern analysis.

Different researchers address these issues in different ways. [Wang et al. \(2023\)](#), for instance, avoided raw reductions by using the preprocessed OCANS-8-2020 catalogue [Yuan et al. \(2021\)](#), which includes time-space unification and catalogue bias correction relative to Gaia DR2. Others, such as [Emelyanov and Samorodov \(2015\)](#) and [Jacobson \(1991\)](#), applied their own corrections to older observations, but their methods are described only briefly.

2.1.6. Coordinate and Time Systems

When working with astrometric observations, two key aspects must be carefully defined: the time scale and the spatial reference frame. Both influence how observations are expressed and how they are transformed into a consistent model for orbit determination.

Time Scales

Two main factors determine the choice and definition of a time scale: (1) the physical realisation of time through a reference clock (e.g., atomic clocks), and (2) relativistic effects, which cause clocks in different gravitational potentials or moving at different velocities to tick at different rates.

The most physically rigorous time scale is the **Barycentric Coordinate Time** (TCB). TCB corresponds to the time coordinate of the **Barycentric Celestial Reference System** (BCRS) and models a clock co-moving with the Solar System barycentre but outside its gravitational well. This is the natural time variable for high-precision barycentric dynamical models, and is also the time scale used by the *Gaia* mission.

In Solar System dynamics, the more common choice is **Barycentric Dynamical Time** (TDB). TDB is linearly related to TCB in such a way that it stays close to **Terrestrial Time** (TT), which itself is offset by a fixed amount from **International Atomic Time** (TAI). TAI is the weighted average of many atomic clocks on the Earth's surface, and thus reflects the influence of Earth's gravitational potential.

A third important time scale is **Coordinated Universal Time** (UTC). UTC is adjusted to match the irregular rotation of the Earth through the use of leap seconds. This makes UTC suitable for civil timekeeping, but less convenient for dynamical modelling. Most modern observations are timestamped in UTC, while older datasets may use historical or obsolete time scales that require careful conversion (see Appendix D of [Emelyanov \(2021\)](#) for a detailed summary).

For orbit determination, UTC is typically converted to TDB (or sometimes to TCB, depending on the chosen dynamical model). This ensures consistency and avoids discontinuities caused by leap seconds.

Spatial Reference Frames

Modern dynamical and observational data are expressed in the **International Celestial Reference Frame** (ICRF). The ICRF is a quasi-inertial, barycentric frame realised by precise positions of extragalactic radio sources, and is aligned to the *Gaia* Celestial Reference Frame (Gaia-CRF) to within microarcseconds.

Older observations, however, may be expressed in an **apparent reference frame**, in which positions are tied to the true equator and equinox of date. This frame reflects the effects of precession, nutation, annual aberration, and light deflection appropriate to the epoch of observation. Converting such data to the ICRF requires applying a model for Earth's precession and nutation, aberration corrections, and often atmospheric refraction (for ground-based observations).

In practical terms:

- **Modern datasets** (e.g., Gaia, Hubble, most CCD astrometry) are already in ICRF or an aligned frame.
- **Older datasets** (e.g., micrometric or early photographic observations) may require transformation from apparent coordinates to ICRF using contemporary precession-nutation models. Care must be taken to understand which model, if any, was implicitly assumed at the time of observation.

Table 2.1: Triton Observation (Part 1)

Site	Date	Source	Type	N1	N2
503 - Cambridge	1847-1847	Lassell (1849)	(<i>p,s</i>)	5	5
992 - Liverpool	1848-1848	Lassell (1849)	(<i>p,s</i>)	4	4
992 - Liverpool	1849-1849	Lassell (1849)	(<i>p,s</i>)	5	5
992 - Liverpool	1849-1849	Lassell (1849-1857)	(<i>p,s</i>)	140	140
992 - Liverpool	1863-1863	Lassell (1864)	(<i>p,s</i>)	20	20
802 - Harvard	1866-1868	Winlock, Pickering (1888)	(<i>p,s</i>)	18	18
787 - USNO, Wash.	1873-1873	USNO (1873)	(<i>p,s</i>)	33	33
787 - USNO, Wash.	1874-1874	Davis (1874)	(<i>p,s</i>)	61	61
787 - USNO	1875-1877	Hall (1876, 1877)	(<i>p,s</i>)	123	123
U.S.N.O.	1877-1877	Holden (1881)	(<i>X,Y</i>)	29	29
007 - Paris	1883-1884	Henry, Boinot, Sy (1886)	(<i>p,s</i>)	31	31
007 - Paris	1884-1884	Henry (1884,1884)	(<i>p,s</i>)	21	21
020 - Nice	1886-1887	Perrotin (1887)	(<i>p,s</i>)	27	27
662 - Lick Obs.	1892-1895	Barnard (1893,1894,1895)	(<i>p,s</i>)	320	320
787 - USNO, Wash.	1894-1906	USNO (1911)	(<i>p,s</i>)	388	388
690 - Lowell Obs.	1896-1897	Drew (1897)	(<i>p,s</i>)	85	85
754 - Yerkes Obs.	1897-1899	Barnard (1898,1899)	(<i>p,s</i>)	330	330
690 - Lowell Obs.	1898-1898	Drew (1899)	(<i>p,s</i>)	36	36
000 - Greenwich	1899-1899	Greenwich (1899)	(<i>p,s</i>)	12	12
000 - Greenwich	1899-1899	Greenwich (1900)	(<i>p,s</i>)	28	28
000-Greenwich	1899-1907	Christie (1901-1909)	(<i>p,s</i>)	60	60
084 - Pulkovo	1899-1920	Balanovskii (1923)	(<i>p,s</i>)	101	101
754 - Yerkes Obs.	1899-1901	Barnard (1901)	(<i>p,s</i>)	361	361
786 - USNO, Wash.	1899-1900	See (1900)	(<i>p,s</i>)	129	129
754 - Yerkes Obs.	1901-1903	Barnard (1903)	(<i>p,s</i>)	223	223
662 - Lick Obs.	1902-1903	Perrine (1903)	(<i>p,s</i>)	51	51
786 - USNO, Wash.	1902-1903	Dinwiddie (1903)	(<i>p,s</i>)	40	40
522 - Strasbourg	1903-1905	Wirtz (1905)	(<i>p,s</i>)	13	13
754 - Yerkes Obs.	1903-1922	Barnard (1906-1927)	(<i>p,s</i>)	989	989
786 - USNO, Wash.	1904-1908	Hammond (1905,1906,1908)	(<i>p,s</i>)	100	100
662 - Lick Obs.	1905-1905	Greenwich (1899)	(<i>p,s</i>)	6	6
754 - Yerkes obs.	1905-1906	Barnard (1906)	(<i>p,s</i>)	32	32
786 - USNO	1908-1910	Hall (1911)	(<i>p,s</i>)	64	64

Table 2.2: Triton Observations (Part 2)

Site	Date	Source	Type	N1	N2
000 - Greenwich	1909-1910	Greenwich (1913)	(p,s)	59	59
786 - USNO	1911-1911	Hall, Burton (1913)	(p,s)	72	72
786 - USNO, Wash.	1911-1912	Burton (1913)	(p,s)	85	85
786 - USNO	1912-1919	Hall, Burton (1910)	(p,s)	282	282
786 - USNO	1919-1923	Hall (1920)	(p,s)	84	84
786 - USNO	1920-1920	Hall (1922)	(p,s)	46	46
084 - Pulkovo	1923-1923	Neuimin, Pokrovskii (1926)	(p,s)	118	118
662 - Lick Obs	1928-1928	Crawford (1928)	(p,s)	15	15
083, 188	1963-1990	Yizhakevych et al. (2016)	(α, δ)	10	10
999 - Bordeaux	1970-1970	Soulie, G. (1975)	(α, δ)	3	3
U.S.N.O.	1975-1977	Walker et al. (1978)	(X, Y)	28	28
U.S.N.O.	1979-1983	Harrington, Walker (1984)	(X, Y)	114	114
U.S.N.O.	1984-1986	Walker, Harrington (1988)	(X, Y)	56	56
119 - Abastumani	1986-1993	Chanturia, Kisseleva (2006)	(α, δ)	54	54
999 - Bordeaux	1995-2007	Arlot et al. (2008)	(α, δ)	95	95
337 - Sheshan	1996-2006	Qiao R.C. et al.(2007)	(α, δ)	943	943
337 - Sheshan	1996-2006	Yan et al. Icarus(2020)(*1)	(α, δ)	1006	1006
Flagstaff	1998-1999	Stone R.C.(2001)	(α, δ)	124	124
Table Mountain	1999-1999	Communicated by Owen	(α, δ)	3	3
Table Mountain	1999-1999	Communicated by Owen	(X, Y)	3	3
Flagstaff	2000-2012	Stone (2000, 2001)	(α, δ)	874	874
Table Mountain	2001-2001	Communicated by Owen (2001)	(α, δ)	3	3
Table Mountain	2001-2001	Communicated by Owen (2001)	(X, Y)	3	3
673-Table Mount.	2002-2002	Table Mountain Obs.(2016)	(α, δ)	201	201
337 - Sheshan	2005-2009	Zhang et al.(2021)	(α, δ)	2299	2299
327 - Peking Obs.	2006-2006	Yu et al. (2019)	(α, δ)	34	34
Peking Sheshan	2007-2009	Qiao et al. (2014)	(α, δ)	1095	1095
337 - Sheshan	2010-2014	Zhang et al.(2022)	(α, δ)	604	604
286 - Yunnan Obs.	2013-2019	Yan et al.(2021)	(α, δ)	1135	1135
Flagstaff	2013-2014	Stone (2000, 2001)	(α, δ)	116	116
258 - Gaia	2014-2017	Not published	(α, δ)	153	153
258 - Gaia	2014-2019	Gaia Archive	(α, δ)	468	468
286 - Yunnan Obs.	2014-2016	Wang et al. (2017)	(α, δ)	755	755
286 - Yunnan Obs.	2020-2024	Yan et al. (2025) in prep.	(α, δ)	1918	1918

2.2. Dynamics of Planetary Systems

Planetary systems are shaped by the mutual gravitational interactions of their constituent bodies, with additional perturbations from non-spherical gravity fields, tidal effects, and solar radiation pressure. The relative importance of these effects depends on the masses, distances, and orbital configurations involved. For natural satellites, assessing the magnitude of each perturbation is critical in selecting an appropriate force model for orbit determination and dynamical analysis.

This section presents the mathematical formulation of a dynamical model for a natural satellite orbiting its host planet. The discussion begins with the central gravitational attraction of the host planet, followed by the transformation to non-inertial reference frames. The extended gravity field is then described using spherical harmonics. For the Neptunian system specifically, the motion of Neptune's pole is modelled, followed by tidal effects acting on both Neptune and Triton. The section concludes with an overview of numerical integration methods and potential sources of numerical error.

2.2.1. Natural Satellite Dynamics

Triton, Neptune's largest moon, was discovered in 1846 and follows a retrograde orbit with an inclination of approximately 157° relative to Neptune's equator. Its orbit is nearly circular, with a semi-major axis of about 350 000 km (~ 14.3 Neptune radii).

The dominant force on Triton is Neptune's central gravitational attraction. Perturbations from Neptune's gravity field (spherical harmonics), tides, other Neptunian moons, the Sun, and other planets have been considered in prior studies. One goal of this work is to quantify the relative magnitudes of these perturbations, rank them, and recommend which should be included given current observational accuracy.

While this analysis focuses on Triton, similar assessments are necessary for other Neptunian moons. Inner moons are more strongly influenced by Neptune's non-spherical gravity and mutual interactions, with possible resonant effects. Triton's large mass can also induce measurable perturbations on their orbits. Conversely, outer moons are more strongly affected by solar perturbations, which can produce significant variations in eccentricity, inclination, and semi-major axis over time.

Following the notation in Tudat [Dirkx et al. \(2022\)](#) the gravitational acceleration of body B acting on point mass body A, expressed in an inertial reference frame can be written as

$$\mathbf{a}_{BA} = \nabla U_B(\mathbf{r}_{BA}) \quad (2.22)$$

where U_B represents the gravitational potential of body B and ∇ the gradient w.r.t inertial position. The equation can be decomposed by splitting the contribution of point masses and extended bodies as

$$\mathbf{a}_{BA} = \mathbf{a}_{\bar{B}\bar{A}} + \mathbf{a}_{\hat{B}\bar{A}} + \mathbf{a}_{\bar{B}\hat{A}} + \mathbf{a}_{\hat{B}\hat{A}} \quad (2.23)$$

1. $\mathbf{a}_{\bar{B}\bar{A}}$: Exerted by point mass B on point mass A .
2. $\mathbf{a}_{\hat{B}\bar{A}}$: Exerted by extended body B on point mass A .
3. $\mathbf{a}_{\bar{B}\hat{A}}$: Exerted by point mass B on extended body A .
4. $\mathbf{a}_{\hat{B}\hat{A}}$: Exerted by extended body B on extended body A , which can be safely neglected as shown in [Fayolle \(2025\)](#).

The point mass gravity for 1 can be simply expressed as

$$\mathbf{a}_{\bar{B}\bar{A}} = -\frac{\mu_B}{\|\mathbf{r}_{BA}\|^2} \hat{\mathbf{r}}_{BA}, \quad (2.24)$$

where \mathbf{r}_{BA} is the vector from B to A and $\hat{\mathbf{r}}_{BA}$ is the normalized vector from B to A.

The effects of the extended body B on A under 2 are usually described using the spherical harmonic potential. The terms $U_{B,lm}$ are a function of the body-fixed position of body A w.r.t body B. Therefore the potential gradient is more easily calculated in the frame fixed to body B (denoted as $\nabla^{(B)}$). To convert the acceleration to inertial reference

frame:

$$\mathbf{a}_{\hat{B}\hat{A}} = \nabla \left(\sum_{l=1}^{l_{\max}} \sum_{m=0}^l U_{B,lm}(\mathbf{r}_{BA}) \right) \quad (2.25)$$

$$= \mathbf{R}^{(I/B)} \left(\sum_{l=1}^{l_{\max}} \sum_{m=0}^l \left(\nabla^{(B)} U_{B,lm}(\mathbf{r}_{BA}) \right) \right) \quad (2.26)$$

where Equation 2.26 gives the extended body acceleration effect from body B on point mass of body A , without the inclusion of point mass of body B effects ($l = 0$).

When the body undergoing acceleration A is treated as having an extended body. We are interested in the force exerted by the point mass B on extended body A under point 3. A useful trick is to apply Newton's third law which can be represented as

$$\mathbf{a}_{\hat{B}\hat{A}} = -\frac{\mu_B}{\mu_A} \mathbf{a}_{\hat{A}\hat{B}} \quad (2.27)$$

$$= -\frac{\mu_B}{\mu_A} \nabla U_{\hat{A}}(\mathbf{r}_{AB}) \quad (2.28)$$

2.2.2. Non-inertial origins

So far only the accelerations under inertial reference frames have been considered. In practice it is useful to define some propagation w.r.t to the central body. For example, the dynamics of Triton around Neptune.

It is important to distinguish two different cases:

1. The computation of $(\mathbf{a}_{BA})_B$: acceleration due to body B, acting on body A, with the frame origin at body B's centre of mass.
2. The computation of $(\mathbf{a}_{BA})_C$: acceleration due to body B, acting on body A, with the frame origin at some other body C. Useful for third body perturbations.

For the first case:

$$(\mathbf{a}_{BA})_B = \mathbf{a}_{BA} - \mathbf{a}_{AB} \quad (2.29)$$

where the second term in the r.h.s of the equation takes into account the accelerations of body A on body B.

When only point mass interactions are considered the equations becomes:

$$(\mathbf{a}_{BA})_B = -\frac{\mu_A + \mu_B}{\|\mathbf{r}_{BA}\|^2} \hat{\mathbf{r}}_{BA} \quad (2.30)$$

Similarly for case 2, the third body perturbation is:

$$(\mathbf{a}_{BA})_C = \mathbf{a}_{BA} - \mathbf{a}_{BC} \quad (2.31)$$

where for point masses only it reduces to:

$$(\mathbf{a}_{BA})_C = \mu_B \left(-\frac{\hat{\mathbf{r}}_{BA}}{\|\mathbf{r}_{BA}\|^2} + \frac{\mathbf{r}_{\hat{B}C}}{\|\mathbf{r}_{BC}\|^2} \right) \quad (2.32)$$

The same can be applied to extended bodies as:

$$(\mathbf{a}_{BA})_B = \nabla U_B(\mathbf{r}_{BA}) - \nabla U_A(\mathbf{r}_{AB}) \quad (2.33)$$

NOTE: This equation is not in the reader but I think its an important distinction

$$(\mathbf{a}_{BA})_C = \nabla U_B(\mathbf{r}_{BA}) - \nabla U_B(\mathbf{r}_{BC}) \quad (2.34)$$

The equations can be formulated fully for the dynamics of a natural satellite in a planetary system w.r.t a central body with subscript 0, and considering the dynamics of N extended bodies with subscripts $i = 1 \dots N$. Then the equations formulate above are:

$$(\mathbf{a}_i)_0 = (\mathbf{a}_0, i)_0 + \sum_{j=1}^{N(i \neq j)} (\mathbf{a}_j, i)_0 \quad (2.35)$$

For completeness, the accelerations due to tidal dissipation and relativistic effects can be included as

$$(\mathbf{a}_i)_0 = (\mathbf{a}_0, i)_0 + \sum_{j=1}^{N(i \neq j)} (\mathbf{a}_j, i)_0 + \mathbf{a}_T + \mathbf{a}_R \quad (2.36)$$

The later sections describe in more detail the formulation of the spherical harmonics of extended body gravity, Neptune's pole precession which will be integrated separately, but is implicitly related to the computation of the forces on Triton, Tidal effects and the list of numerical integrators available in Tudat.

2.2.3. Spherical Harmonic Gravity and Inner Moons

As described previously the acceleration due to the gravity of an extended body k on body i can be described by the gradient in body-fixed reference frame as

$$\mathbf{a}_{k,i} = \nabla^{(k)} U_k(\mathbf{r}_{k,i}) \quad (2.37)$$

The gravity field of an extended body k is typically modelled by a spherical harmonic expansion of its gravitational potential at the position of body i described by the parameters r, ϕ, λ . The full formulation is then:

$$U_k(r, \phi, \lambda) = \frac{Gm_k}{r} \sum_{l=1}^{l_{\max}} \left(\frac{R_k}{r} \right)^l \sum_{m=0}^l [\check{C}_{lm}^k \cos(m\lambda) + \check{S}_{lm}^k \sin(m\lambda)] \check{P}_{lm}(\sin \phi). \quad (2.38)$$

Where m_k is the total mass of body k and R_k represents its equatorial radius. The radial distance, r , latitude, ϕ , and longitude, λ , are the spherical coordinates of the point at which the potential is evaluated, in the reference frame fixed to body k . \check{C}_{lm}^k and \check{S}_{lm}^k are body k 's fully normalised cosine and sine spherical harmonics coefficients of degree n and order m , respectively, and \check{P}_{lm} designates the normalised associated Legendre polynomial. The fully normalised spherical harmonics coefficients can be obtained from the unnormalised coefficients as described in [Fayolle \(2025\)](#) or [Dirkx et al. \(2022\)](#). These coefficients are directly related to the internal density distribution of body k .

The potential given by [Equation 2.38](#) is expressed in the reference frame fixed to body k and can be directly used in [Equation 2.26](#).

[Emelyanov \(2021\)](#) proposed modelling Neptune's inner moons as mass disks around the equator, effectively increasing the magnitude of the planet's J_2 and J_4 gravity terms. In the present study, this approximation is not considered necessary due to the relatively low computational cost of the simulations. However, it remains a useful option should the simulations become computationally intensive.

2.2.4. Neptune's Pole Precession

The pole of Neptune precesses due to moments exerted on the planet by Triton and other perturbing bodies.

The most simple formulation of rotational dynamics for a non-deformable solid body is given by the Euler equation:

$$\frac{d(\mathbf{I}\boldsymbol{\omega})}{dt} + \boldsymbol{\omega} \times (\mathbf{I}\boldsymbol{\omega}) = \sum_i \boldsymbol{\Gamma}_i, \quad (2.39)$$

where \mathbf{I} is the inertia tensor of a rigid body, $\boldsymbol{\omega}$ is the rotation vector, and the right-hand side of the equation represents the sum of all torques acting on the body. In reality, Neptune cannot be accurately treated as a solid body; internal fluid layers and differential rotation can produce coupling between the inertia tensor and the applied torques. In such cases, [Equation 2.39](#) no longer applies directly, and more sophisticated treatments—such as models based on the Laplace tidal equations—are required.

Nevertheless, rotation models developed by [Jacobson \(2009\)](#) and [Tang et al. \(2020\)](#) adopt the solid-body approximation for Neptune. These models describe the planet's precession under the assumption that the total angular momentum remains constant in both magnitude and direction. The pole motion is incorporated into the equations of motion, and the precession rate and direction can be introduced as estimable parameters, as demonstrated by [Jacobson \(2009\)](#).

Formulation by Jacobson

The pole precession can be modelled as the pole direction vector, \hat{h} , precession about an axis, \hat{k} , (the direction vector of the system angular momentum), used by [Jacobson \(1991\)](#) as:

$$\hat{h} = \cos(\dot{\Omega}T)\hat{h}_0 + (\hat{h}_0 \cdot \hat{k})[1 - \cos(\dot{\Omega}T)]\hat{k} + \sin(\dot{\Omega}T) \times (\hat{h}_0 \times \hat{k}) \quad (2.40)$$

with \hat{h}_0 being the pole at starting epoch, $\dot{\Omega}$ the precession rate, and T the time from epoch. The unit vector normal to Triton's orbit, $\hat{\omega}$, also precesses about the system angular momentum with the same rate as the Neptune pole. Its direction is specified by with \hat{n}_0 being the pole at epoch, $\dot{\Omega}$ the precession rate, and T the time from epoch. The unit vector normal to Triton's orbit, \hat{w} , also precesses about the system angular momentum with the same rate as the Neptune pole. Its direction was specified by

$$\hat{w} = \frac{\mathbf{r}_1 \times \dot{\mathbf{r}}_1}{|\mathbf{r}_1 \times \dot{\mathbf{r}}_1|} = \hat{k} \cos I + \left(\hat{p} \sin(\Omega_0 + \dot{\Omega}T) - \hat{q} \cos(\Omega_0 + \dot{\Omega}T) \right) \sin I, \quad (2.41)$$

where $\mathbf{r}_1, \dot{\mathbf{r}}_1$ are Triton's barycentric position and velocity, I is the angle between Triton's orbit normal and the system angular momentum, and Ω_0 is the epoch value of ascending node of Triton's orbit on the plane normal to \hat{k} . The coordinate axes were

$$\hat{p} = (e_3 \times \hat{k}) \sec \delta_r, \quad \hat{q} = \hat{k} \times \hat{p}, \quad \text{Jacobson (1991)} \quad (2.42)$$

where e_3 is the unit vector $(0, 0, 1)$ and δ_r is the declination of \hat{k} .

The components of the Neptune and Triton angular momentum normal to the system angular momentum were related by

$$\left[1 - \left(1 - \frac{J_2}{\gamma} \right) \left(\frac{\dot{\Omega}}{\dot{W}} \right) \cos \epsilon \right] \cos \epsilon = \frac{\mu_1 \mu_2^2}{(\mu_s - \mu_1)^3} \frac{\sin I}{\gamma R^2 \dot{W}} \frac{|\mathbf{r}_1 \times \dot{\mathbf{r}}_1|}{|\mathbf{r}_1|}, \quad (2.43)$$

where ϵ is the angle between \hat{n} and \hat{k} , J_2 is the second zonal harmonic of Neptune's gravitational potential, γ is Neptune's axial moment of inertia factor, \dot{W} is Neptune's spin rate, μ_s is the GM of the Neptunian system, μ_1 is the GM of Triton, and R is the equatorial radius of Neptune. GM is the product of the Newtonian constant of gravitation G and the mass M .

To constrain the solution in a fitting procedure, the following expressions are required to hold in a least-squares sense over the data processing time span:

$$\hat{h} \cdot \hat{k} = \cos \epsilon \quad (2.44)$$

$$\hat{w} \cdot \hat{k} = \cos I \quad (2.45)$$

$$\hat{w} \cdot \hat{p} = \sin I \sin(\Omega_0 + \dot{\Omega}T) \quad (2.46)$$

Equation (2.44) enforces the angle ϵ between the total angular momentum direction \hat{h} and Neptune's spin pole \hat{k} . Equation (2.45) constrains the angle I between Triton's orbit normal \hat{w} and \hat{k} , while Equation (2.46) constrains the projection of \hat{w} onto a reference axis \hat{p} , using the precession rate $\dot{\Omega}$ and the epochal ascending node Ω_0 .

Rather than representing the pole with a vectorial model, [Jacobson \(2009\)](#) adopts a formulation in terms of right ascension and declination angles (α, δ) , expanded in series to second order in ϵ :

$$\tan(\alpha - \alpha_r) = \frac{\tan \epsilon \sec \delta_r \sin \Omega}{1 + \tan \epsilon \tan \delta_r \cos \Omega}, \quad \sin \delta = \sin \delta_r \cos \epsilon - \cos \delta_r \sin \epsilon \cos \Omega \quad (2.47)$$

where (α_r, δ_r) denote the right ascension and declination of the system angular momentum, and $\Omega = \Omega_0 + \dot{\Omega}T$.

Using a series expansion to $\mathcal{O}(\epsilon^2)$:

$$\alpha = \alpha_r + \epsilon \sec \delta_r \sin \Omega - \frac{1}{2} \epsilon^2 \sec \delta_r \tan \delta_r \sin 2\Omega \quad (2.48)$$

$$\delta = \delta_r - \epsilon \cos \Omega - \frac{1}{4} \epsilon^2 \tan \delta_r (1 - \cos 2\Omega) \quad (2.49)$$

These expansions are advantageous for compatibility with JPL planetary navigation software and IAU/IAG orientation standards.

In a fitting procedure, the estimated parameters include α_r , δ_r , ϵ , I , Ω_0 , and $\dot{\Omega}$. As with the model proposed by [Jacobson \(1991\)](#), [Equation 2.43](#) and [Equation 2.45](#) are required to constrain the solution, while [Equation 2.44](#) is replaced with the numerically superior formulation:

$$\hat{\boldsymbol{w}} \cdot (\hat{\boldsymbol{p}} \cos \Omega + \hat{\boldsymbol{q}} \sin \Omega) = 0 \quad (2.50)$$

Formulation by Tang

[Tang et al. \(2020\)](#) proposes a dynamical model of Neptune's pole precession by integrating the simplified Euler equations of motion for a rigid axially symmetric body under external torque. Unlike [Jacobson \(2009\)](#), who assumes that the angular momentum vector of the Neptune-Triton system remains fixed in inertial space (and models Neptune's pole precessing around it), Tang explicitly computes the time evolution of Neptune's pole using the torque induced by Triton.

Neptune's orientation is defined using two Euler angles, φ and θ , with respect to an intermediate reference frame $O - X'Y'Z'$ from the Neptunian reference system by applying the rotation obtained via the rotation $R_1(\theta)R_3(\varphi)$. The Z' axis of this intermediate reference system is aligned with Neptune's pole. The equations governing the evolution of these angles are derived from Euler's equations, neglecting higher-order terms as

$$A \frac{d\dot{\theta}}{dt} + C\omega\dot{\varphi} \sin \theta = T_{X'}, \quad (2.51)$$

$$A \frac{d}{dt}(\dot{\varphi} \sin \theta) - C\omega\dot{\theta} = T_{Y'}, \quad (2.52)$$

where ω is Neptune's spin rate, A , C are Neptune's equatorial and polar moments of inertia, φ and θ are Euler angles relating the inertial and body-fixed frames and $T_{X'}$, $T_{Y'}$ are the torque components in the $X'Y'Z'$ frame.

The angles are related to Neptune's pole right ascension, α_N , and declination, δ_N , via:

$$\varphi = \frac{\pi}{2} + \alpha_N, \quad (2.53)$$

$$\theta = \frac{\pi}{2} - \delta_N, \quad (2.54)$$

The torque T is computed as the gravitational torque from Triton acting on Neptune's equatorial bulge. This torque arises due to the misalignment of Triton's inclined orbit with Neptune's equator and depends on Neptune's J_2 term, equatorial radius R , and the mass of Triton m_T . Although the detailed expression for the torque is not explicitly given in [Tang et al. \(2020\)](#).

The authors compare the model to [Jacobson \(2009\)](#) and find differences no larger than 2.5" deemed sufficiently small. A key topic in this study is to examine both models and derive a full mathematical model of the pole motion due to Triton's dynamics.

2.2.5. Tidal effects on Triton and Neptune

The gravitational pull of a body j (here treated as a point mass) generates a so-called tidal bulge on body i , responsible for its deformation. The tide-inducing potential at any point on body i 's surface caused by j is given by




Figure 2.5: Schematic representation of the tidal deformation (solid red ellipsoid) of body i (left) at time t due to the tidal potential (dashed red ellipsoid) raised by point mass body j (right) at time t' . The rotation of body i and orbital motion of body j during the time lag $\Delta t = t - t'$ are also represented. (Figure from [Fayolle \(2025\)](#))

$$V_i^{(j)}(\mathbf{r}^*) = \frac{Gm_j}{R_i} \sum_{n=2}^{\infty} \left(\frac{R_i}{r_{ij}}\right)^{n+1} P_n(\hat{\mathbf{r}}^* \cdot \hat{\mathbf{r}}_{ij}), \quad (2.55)$$

where \mathbf{r}^* is the position of the point at which the potential is evaluated, and r_{ij} is the position of the tidal-inducing body j with respect to body i .

The deformation of body i 's gravity field caused by the tidal bulge described by [Equation 2.55](#) in turn raises the following gravitational potential:

$$U_i^{(j)}(\mathbf{r}) = \frac{Gm_j}{R_i} \sum_{n=2}^{\infty} k_n^i \left(\frac{R_i}{r}\right)^{n+1} \left(\frac{R_i}{r_{ij}}\right)^{n+1} P_n(\hat{\mathbf{r}} \cdot \hat{\mathbf{r}}_{ij}), \quad (2.56)$$

with r defining the position at which the potential is calculated with respect to body i . The gravity field variations caused by tides, and the resulting tidal bulge potential that they raise are driven by the tidal Love numbers k_{nm}^i , which describe the body's visco-elastic response to tidal forcing (n and m respectively denoting the degree and order of the Love number). These are frequency-dependent complex numbers whose imaginary part defines the phase lag of the tidal response.

The second degree Love number k_2 has the strongest effect on the tides potential and often [Equation 2.56](#) is reduced to its degree two component as:

$$U_i^{(j)}(\mathbf{r}) = \frac{Gm_j k_2^i}{2R_i} \left(\frac{R_i}{r}\right)^3 \left(\frac{R_i}{r_{ij}}\right)^3 (3(\hat{\mathbf{r}} \cdot \hat{\mathbf{r}}_{ij})^2 - 1). \quad (2.57)$$

The phase lag between the perturbing potential and the body's response is usually represented as a quality factor Q , which is directly related to k_2 's imaginary part $\text{Im}(k_2)$ as:

$$Q = \frac{|k_2|}{\text{Im}(k_2)} \quad (2.58)$$

Q is frequency-dependent, and depends on the body's rheology. A full description of this frequency dependency is outside the scope of this study. Therefore, certain assumptions have to be made to continue the analysis. The tidal forcing for both planet- and satellite-raised tides is dominated by a single leading frequency, driven by the rotation period of the tide-raising body and by the moon's orbital period.

The most widely-used approximations are the constant phase lag (CPL) and constant time lag (CTL) models. Both rely on important and much disputed assumptions [Fayolle \(2025\)](#). The CPL model, for example assumes constant

angular phase lag and therefore does not hold for eccentric orbits. Both the CPL and CTL models dismiss the frequency-dependence of the response of the body's interior to the tidal forcing. Which makes them unsuitable for long-term evolution studies.

For shorter timescales, such as typical ephemerides studies, the CTL model is suitable [Fayolle \(2025\)](#). The time lag of dissipation occurring in body i due to tides raised by body j is related to the tidal quality factor $Q_i^{(j)}$ by the following

$$\Delta t_i^{(j)} = \frac{T_i^{(j)} \arcsin 1/Q_i^{(j)}}{2\pi}, \quad (2.59)$$

where $Q_i^{(j)}$ denotes the tidal quality factor of body i at the forcing frequency of the tides raised by j , while $T_i^{(j)}$ refers to the period of this tidal forcing. For tides raised by the satellite i on the planet 0, this period depends on the satellite's orbital period and on the planet's rotational period defined as $T_i^{\text{orb}} = \frac{2\pi}{n_i}$ and $T_0^{\text{rot}} = \frac{2\pi}{\omega_0}$, respectively, with n_i the mean motion of satellite i and ω_0 the planet's rotational rate:

$$T_0^{(i)} = \frac{T_i^{\text{orb}} T_0^{\text{rot}}}{2|T_0^{\text{rot}} - T_i^{\text{orb}}|} = \frac{2\pi}{2|\omega_0 - n_i|}. \quad (2.60)$$

For the tides raised once per orbit by the planet on a fully synchronous satellite, on the other hand, the tidal forcing period is

$$T_i^{(0)} = \frac{2\pi}{n_i}. \quad (2.61)$$

The tides raised by satellite i on central planet 0 generate a force acting on satellite i due to the planet's tidal bulge. The force in body-fixed frame is given by

$$\mathbf{F}_i^{(0)} = -m_i \nabla U_0^{(i)}(\mathbf{r}_i) \quad (2.62)$$

Substituting [Equation 2.56](#) in the expression above, after extensive simplifications, a truncation at $O((\Delta t_i^{(0)})^2)$ and decomposing the velocity vector in radial and tangential components (assuming that the satellite's rotation pole is aligned with the normal to the orbit plane), the formulation for tides raised on the central planet becomes:

$$\mathbf{F}_i^{(0)} = -\frac{3Gm_i^2 k_2^0 R_0^5}{r_i^8} \left[\mathbf{r}_i + \Delta t_0^{(i)} \left(2 \frac{\mathbf{r}_i \cdot \dot{\mathbf{r}}_i}{r_i^2} \mathbf{r}_i + (\dot{\mathbf{r}}_i + \mathbf{r}_i \times \boldsymbol{\omega}_0) \right) \right]. \quad (2.63)$$

For the tides raised on the moon, the fact that energy dissipated by librational tides (tangential component) over one orbit is equal to 4/3 that dissipated by radial tides [Murray and Dermott \(1999\)](#) can be used. Then the tidal force on the satellite becomes:

$$\mathbf{F}_i^{(i)} = -\frac{7Gm_0^2 k_2^i R_i^5}{r_i^7} \left(1 + 3\Delta t_i^{(0)} \frac{\hat{\mathbf{r}}_i \cdot \dot{\mathbf{r}}_i}{r_i} \right) \hat{\mathbf{r}}_i. \quad (2.64)$$

The tides raised on the satellite are negligible compared to the tides raised on the central body and therefore will not be part of the fitting procedure. Therefore only [Equation 2.63](#) will be used in further analysis. The inclusion of tidal forces in the fitting procedure is explained in more detail in [Lainey \(2016\)](#) and [Wang et al. \(2025\)](#).

2.2.6. Numerical Integration Methods and Error Sources

Tudat [Dirkx et al. \(2022\)](#) offers a wide range of numerical integration schemes, making it important to select the method that is most efficient and accurate for a given problem. This section provides a brief overview of the two principal types of numerical error—*round-off* and *truncation errors*—and the classification of integration schemes available. For a more detailed treatment, the reader is referred to standard numerical analysis references.

Error Sources in Integration

The two dominant sources of error in numerical integration are **numerical (rounding) errors** and **truncation errors**.

- **Numerical errors** stem from the finite precision of floating-point arithmetic. For example, double-precision variables typically store around 16 decimal digits, meaning that information can be lost during repeated mathematical operations such as summation or multiplication. These errors accumulate over time. Although modern CPUs are powerful enough to rely on large variable sizes, some accumulation is inevitable, especially for more intricate integration methods. Numerical errors can be reduced—though not eliminated—by employing variables with higher precision (e.g., *quadruple precision*, which typically provides 34 digits), or by applying algorithmic techniques to minimize cancellation or overflow. The reader is referred to [Oliveira and Stewart \(2006\)](#).
- **Truncation errors** arise from the approximations inherent in the numerical integration method. For instance, Euler's method assumes the derivative of the state remains constant over a time step, which introduces a local truncation error proportional to the square of the step size. Higher-order methods reduce this error, but do not eliminate it.

Integration Methods

Integration methods generally fall into two categories: **multi-stage** and **multi-step** methods.

- **Multi-stage methods** (e.g., Runge-Kutta methods) evaluate the derivative multiple times per step. Examples include:
 - *Fixed step-size* methods, such as Euler or classical 4th-order Runge-Kutta (RK4).
 - *Adaptive step-size* methods (e.g., Runge-Kutta-Fehlberg or Dormand-Prince), which estimate local truncation error by comparing the results of a lower-order and a higher-order method on the same interval.

These adaptive methods require the user to set *absolute* and *relative* tolerances to guide error control:

- The **absolute tolerance** defines the minimum acceptable error per integration step.
- The **relative tolerance** defines acceptable error relative to the current state's magnitude.

Loose tolerances will lead to faster but less accurate integrations, while strict tolerances increase computational cost and risk reaching machine precision, where numerical noise can dominate and cause integrator failure. Since numerical errors are stochastic, overly tight tolerances can cause the solver to reduce the step size indefinitely without satisfying the error bounds.

- **Multi-step methods** (e.g., Adams-Bashforth, Adams-Moulton) differ in that they reuse previous steps to estimate the next state, rather than computing multiple function evaluations within one step. These methods can be efficient for long-duration, smooth problems where the function's derivative varies slowly. However, they require good initial values and are less suitable for high stiffness problems.

Practical Considerations

In practice, the most suitable integrator depends on the specific problem being solved. Problems involving short integration windows but high accuracy (e.g., spacecraft trajectories near flybys) may benefit from adaptive Runge-Kutta methods, while long-term simulations of celestial dynamics may be more efficient with multi-step schemes.

2.3. Natural Satellite Orbit Determination

Weighted least-squares estimation is the preferred method for natural satellite orbit determination when a real-time solution is not required. It offers a robust and well-established approach for deriving accurate orbital parameters by fitting a dynamical model to observational data.

The core idea of least-squares estimation in this context is to determine the initial state of the satellite at a selected epoch, along with relevant model parameters, such that the squared differences between the observed and modelled measurements are minimized. Because observational data typically varies in type and reliability, a weighting factor is applied to each residual to reflect its relative confidence before computing the squared differences. This results in the minimization of the weighted residuals rather than the unweighted ones.

In this section, the mathematical foundation of the weighted least-squares method is briefly outlined, following the notation introduced by [Montenbruck and Gill \(2000\)](#). For a more detailed mathematical treatment, the reader is referred to their work.

2.3.1. Weighted Least Squares

The first step is to define a time-dependent, m -dimensional vector vector

$$\mathbf{x}(t) = \begin{pmatrix} r(t) \\ v(t) \\ p \\ q \end{pmatrix}$$

which consists of the satellite state vector r, v as well as the free parameters p and q that affect the force and measurement model. This way the parameter estimation and the trajectory could be treated simultaneously. The time evolution can always be described by an ordinary differential equations of the form

$$\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x}) \quad (2.65)$$

with initial value

$$x_0 = x(t_0) \quad (2.66)$$

The n -dimensional vector of measurements taken at time t_1, \dots, t_n is defined as

$$\mathbf{z} = \begin{pmatrix} z_1 \\ \cdot \\ \cdot \\ \cdot \\ z_n \end{pmatrix}$$

The observations are described by

$$z_i(t_i) = g_i(t_i, x(t_i)) + \epsilon_i = h_i(t_i, x_0) + \epsilon_i \quad (2.67)$$

or briefly

$$z = \mathbf{h}(\mathbf{x}_0) + \epsilon \quad (2.68)$$

Where g_i denotes the model value of the i th observation as a function of time and the instantaneous state. This value can also be computed through h_i , as function of time but only dependent on the initial state x_0 . Where ϵ denotes the difference between the actual measurements and modelled observations due to measurement errors, usually treated with a random function with mean zero.

The least-squares orbit determination problem can be defined as finding the state x_0^{lsq} which minimizes the loss function

$$J(\mathbf{x}_0) = \rho^T \rho = (\mathbf{z} - \mathbf{h}(\mathbf{x}_0))^T (\mathbf{z} - \mathbf{h}(\mathbf{x}_0)) \quad (2.69)$$

Here to avoid under-constraining the problem the number of observations n should be at least equal to the number of unknowns m .

In order to solve this equation practically the problem needs to be linearized as \mathbf{h} is highly non-linear. Usually an approximate value \mathbf{x}_0^{appr} of the initial state vector is known. This value is used to linearize all quantities around as \mathbf{x}_0^{ref} , then the residual vector approximation is

$$\boldsymbol{\rho} = \mathbf{z} - \mathbf{h}(\mathbf{x}_0) \quad (2.70)$$

$$\boldsymbol{\rho} \approx \mathbf{z} - \mathbf{h}(\mathbf{x}_0^{ref}) - \frac{\delta \mathbf{h}}{\delta \mathbf{x}_0} (\mathbf{x}_0 - \mathbf{x}_0^{ref}) = \Delta \mathbf{z} - \mathbf{H} \Delta \mathbf{x}_0 \quad (2.71)$$

Where the Jacobian is defined as

$$\mathbf{H} = \left. \frac{\partial \mathbf{h}(\mathbf{x}_0)}{\partial \mathbf{x}_0} \right|_{\mathbf{x}_0 = \mathbf{x}_0^{ref}} \quad (2.72)$$

The problem can now reduced to a linear least-squares cost function minimisation by finding $\Delta \mathbf{x}_0^{lsq}$ such that

$$J(\Delta \mathbf{x}_0) = (\Delta \mathbf{z} - \mathbf{H} \Delta \mathbf{x}_0)^T (\Delta \mathbf{z} - \mathbf{H} \Delta \mathbf{x}_0) \quad (2.73)$$

becomes a minimum. If the columns of \mathbf{H} are linearly independent, then the minimum is uniquely determined when the partial derivatives of J with respect to $\Delta \mathbf{x}_0$ vanish.

After some cumbersome mathematical operations the problem can be simplified to solving the m-dimensional normal equations

$$(\mathbf{H}^T \mathbf{H}) \Delta \mathbf{x}_0^{lsq} = \mathbf{H}^T \Delta \mathbf{z} \quad (2.74)$$

Due to the non-linearity of \mathbf{h} the linear loss function will differ slightly from the exact solution. It can be further improved by substituting in the reference value \mathbf{x}_0^{ref} . The procedure can be iterated until a prescribed tolerance is met. The updated of the Jacobian \mathbf{H} in each step is necessary for optimum convergence.

So far the derivation of an unweighted least square method have been explained. To properly account for the accuracy of each individual measurement the solution to the weighted least-squares problem may be written as

$$\Delta \mathbf{x}_0^{lsq} = (\mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} (\mathbf{H}^T \mathbf{W} \Delta \mathbf{z}) \quad (2.75)$$

where the weighting matrix is

$$\mathbf{W} = \text{diag}(\sigma_1^{-2}, \dots, \sigma_n^{-2}) \quad (2.76)$$

the mean measurement error σ_i corresponds to each measurement z_i and should consider the total expected error due to random noise and systematic errors. In this case the weighting matrix is diagonal, meaning the measurement errors are uncorrelated, but it can be generalized even further to include correlated errors in which case \mathbf{W} becomes non-diagonal. This makes the computation more expensive and can lead to numerical instability. In practice, if the observations are correlated (e.g. Gaia X-Y plane) a transformation is performed before the estimation procedure to make them uncorrelated.

2.3.2. Covariance of the Estimated State

When treating observations with no errors and a perfect physical model the determined \mathbf{x}_0^{lsq} is the initial state for which all residuals vanish. In practice, observations exhibit some source of uncertainties and a practical question is: What is the influence of the observational accuracy on the least-squares solution?

To answer this question, the observation vector with measurement errors has to be defined as

$$\mathbf{z} = \mathbf{h}(\mathbf{x}_0) + \boldsymbol{\epsilon} \quad (2.77)$$

After linearization the solution to the least square problem is

$$\mathbf{x}_0^{lsq} = \mathbf{x}_0 + (\mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} (\mathbf{H}^T \mathbf{W} \boldsymbol{\epsilon}), \quad (2.78)$$

which shows the difference between the computed \mathbf{x}_0^{lsq} and the actual state \mathbf{x}_0 differ due to the measurement noise $\boldsymbol{\epsilon}$.

Assuming a mean measurement noise of zero and uncorrelated measurements then the covariance of the computed initial state can be derived as

$$\text{Cov}(\mathbf{x}_0^{\text{lsq}}, \mathbf{x}_0^{\text{lsq}}) = \mathbb{E} \left[\left(\mathbf{x}_0^{\text{lsq}} - \mathbf{x}_0 \right) \left(\mathbf{x}_0^{\text{lsq}} - \mathbf{x}_0 \right)^T \right] \quad (2.79)$$

that describes the mean squared deviation of $\mathbf{x}_0^{\text{lsq}}$ from the actual state. Inserting the expression for $\mathbf{x}_0^{\text{lsq}}$ yields

$$\text{Cov}(\mathbf{x}_0^{\text{lsq}}, \mathbf{x}_0^{\text{lsq}}) = (\mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} (\mathbf{H}^T \mathbf{W}) \mathbb{E}(\epsilon \epsilon^T) (\mathbf{W} \mathbf{H}) (\mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} \quad (2.80)$$

which may further be simplified, provided that the weighting matrix has been chosen in accordance with the measurement standard deviation. Then

$$\text{Cov}(\mathbf{x}_0^{\text{lsq}}, \mathbf{x}_0^{\text{lsq}}) = (\mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} \quad (2.81)$$

which is just the inverse of the normal equations matrix. The diagonal elements of the covariance matrix yield the standard deviation of the components of $\mathbf{x}_0^{\text{lsq}}$, while the off-diagonal terms are a measure of the correlation between errors of individual components.

2.3.3. A priori constraints

Aside from the initial state $\mathbf{x}_0^{\text{ref}}$ which is required to start the least-squares orbit determination, usually some information about the accuracy of the initial state is also available. This is contained in the a priori covariance $\mathbf{P}_0^{\text{ref}}$, which can be incorporated in the loss function [Equation 2.69](#) as

$$J = (\mathbf{x}_0 - \mathbf{x}_0^{\text{ref}}) \mathbf{\Lambda} (\mathbf{x}_0 - \mathbf{x}_0^{\text{ref}}) + \rho^T \rho \quad (2.82)$$

where $\mathbf{\Lambda} = (\mathbf{P}_0^{\text{ref}})^{-1}$ which is known as the information matrix, used to penalize any deviation from $\mathbf{x}_0^{\text{ref}}$ by an appropriate contribution to the loss function.

The formulation of the least-square solution then, through rigorous mathematics, can be reduced to

$$\Delta \mathbf{x}_0^{\text{lsq}} = (\mathbf{\Lambda} + \mathbf{H}^T \mathbf{W} \mathbf{H})^{-1} (\mathbf{\Lambda} \Delta \mathbf{x}_0^{\text{ref}} + \mathbf{H}^T \mathbf{W} \Delta \mathbf{z}) \quad (2.83)$$

The covariance of the estimation \mathbf{P}_0 is related to the a priori covariance and the measurement information matrix by

$$(\mathbf{P}_0)^{-1} = (\mathbf{P}_0^{\text{ref}})^{-1} + (\mathbf{H}^T \mathbf{W} \mathbf{H}) \quad (2.84)$$

2.3.4. Considered Parameters

While the covariance matrix gives a measure of the orbit determination accuracy, it is often optimistic in the presence of systematic force and measurement model errors. This has a significant impact if a large number of redundant measurements is processed, since the computed covariance is inversely proportional to the number of measurements within a given data arc. The inclusion of systematic errors does not depend on the number of data points and eventually limits the attainable orbit determination accuracy. It is therefore worthwhile to consider the effect of unmodelled, systematic errors in the computation of the covariance matrix. This can be done through the consider covariance matrix to attain realistic estimates of the achievable orbit determination accuracy.

First, the observation vector with systematic errors is defined as

$$\mathbf{z} = \mathbf{h}(\mathbf{x}_0, \mathbf{c}) + \epsilon \quad (2.85)$$

where the consider parameters \mathbf{c} is a vector comprising the force and measurement model parameters that are uncertain but are not adjusted as part of the least-squares estimation.

Then the least-squares solution can be written as

$$\mathbf{x}_0^{\text{lsq}} = \mathbf{x}_0 + (\mathbf{H}_x^T \mathbf{W} \mathbf{H}_x)^{-1} \mathbf{H}_x^T \mathbf{W} (\mathbf{H}_c \mathbf{c} + \epsilon) \quad (2.86)$$

where H_x and H_c denote the partial derivatives of the modelled measurements h with respect to x_0 and c , respectively. This solution differs from the true values of the estimation parameters by an offset that depends on both the consider parameters c and the measurement noise ϵ .

The noise only covariance is given by

$$\mathbf{P} = (\mathbf{H}_x^T \mathbf{W} \mathbf{H}_x)^{-1} \quad (2.87)$$

Then the consider covariance matrix \mathbf{P}^c , which is also designated as formal or computed covariance. It is given by

$$\mathbf{P}^c = \mathbf{P} + (\mathbf{P} \mathbf{H}_x^T \mathbf{W}) (\mathbf{H}_c \mathbf{C} \mathbf{H}_c^T) (\mathbf{P} \mathbf{H}_x^T \mathbf{W})^T \quad (2.88)$$

In the above formulation, both the consider parameters c and the measurement noise ϵ are modelled as random quantities, but their roles differ fundamentally. Measurement noise affects individual observations within a single orbit determination. When a large number of measurements are collected, this noise is averaged out, leading to reduced uncertainty in the estimated parameters.

Consider parameters, on the other hand, are treated as constant during a single orbit determination but are subject to known uncertainties. These uncertainties propagate into the estimated parameters and contribute directly to the consider covariance. Unlike measurement noise, the additional term introduced in Equation 2.88 does not diminish with higher data rates. It remains essentially constant for a given data arc.

The consider covariance framework is therefore particularly effective for evaluating the impact of systematic errors in orbit determination. Common examples of consider parameters include measurement biases, station location errors, and uncertainties in physical models (higher order extended body gravity) which can not be reliably estimated with the available data quality.

2.3.5. Considerations & Coupled Parameters

The method presented here introduces the fundamental mathematical framework of the weighted least-squares approach. In practice, when applied to natural satellite orbit determination, numerous challenges can arise at different stages of the process.

For instance, both subsection 2.3.1 and the final expression derived in subsection 2.3.2 assume that the weight matrix is constructed directly from the measurement standard deviations. In reality, however, observational data often originates from multiple sources, and scaling factors are sometimes applied to different datasets. Determining the appropriate values for the weighting matrix can be a labor-intensive and nuanced task, yet it is rarely described in detail in the literature. One of the objectives of this study is to provide a clear and rigorous explanation of how data weighting is performed in practice.

Another important aspect is the coupling between estimated parameters. For example, the initial velocity along the orbital track and the radial position of a satellite relative to its central planet are strongly correlated, as both are tied to the system's orbital energy. In the case of the Neptunian system, the pole orientation of Neptune and its extended-body gravity field are similarly coupled in their influence on Triton's orbit. Such parameter coupling is inherent to the system's dynamical properties and the nature of the available data. Consequently, it is challenging—if not impossible—to fully decouple the parameters unless additional types of observations are available (e.g., direct measurements of Neptune's gravity field). Since this study focuses on optical tracking data, and direct measurements of the Neptunian system are limited or unavailable, this is not a practical solution to the coupling problem.

Coupling can also lead to multiple plausible solutions. For instance, different combinations of pole orientation and extended-body gravity parameters may produce nearly identical orbital fits for Triton. In such cases, the least-squares method may fail to converge, or it may produce non-physical values for some parameters. One potential remedy is the application of a priori constraints to stabilize the estimation of strongly coupled parameters.

In summary, while the weighted least-squares method provides a robust mathematical foundation, its practical application to natural satellite orbit determination requires careful attention to parameter weighting, correlation effects, and solution stability. Addressing these challenges is essential for obtaining physically meaningful and reliable results.

2.4. Modern Research on the Neptunian System

Modern research on the Neptunian system aims to refine orbital solutions, characterize dynamical parameters, and improve the physical understanding of Neptune and its satellites. These studies are motivated by both fundamental scientific interest—such as constraining formation scenarios—and practical applications, including the planning of future observation campaigns.

This section reviews the current state of research on the Neptunian system, focusing on the methods, datasets, and dynamical models employed in recent work. The discussion highlights the ephemerides and dynamical parameters in use, evaluates the treatment of observational data, and assesses the strengths and limitations of current approaches. This overview provides the necessary context for identifying gaps in the literature and formulating the research questions addressed in the present study.

2.4.1. Ephemerides and Purpose of Study

Modern research in natural satellite dynamics focuses on improving our understanding of the motion of planets and their satellites. Such advancements not only support the planning of future space missions but also contribute to a deeper comprehension of Solar System formation, as well as the internal structure and evolution of planets and satellites—including Earth and the Moon.

The primary aim of such studies is to produce accurate ephemerides for the target body, derived from current models of its governing dynamics and the available observations. These results can improve long-term predictions of the satellite's position and enable the estimation of key dynamical parameters for both the satellite and its parent body—such as Neptune's extended gravity field, tidal dissipation factors, and pole precession rates.

Contemporary research builds upon a long history of rigorous astronomical and dynamical studies. Foundational contributions include the accurate determination of planetary and solar positions, the construction of inertial reference frames, the development of precise star catalogues, and the detailed characterization of the gravity fields of major bodies.

In this section, it is important to briefly summarize the differences among commonly used planetary ephemerides. Different research groups have developed their own versions, often with subtle variations. Since ephemerides are essential for modelling third-body perturbations, understanding the methods by which they are computed is critical. Comparing ephemerides of similar accuracy can also help identify potential observational biases that might otherwise remain unaccounted for.

2.4.2. Summary of Modern Research

A selected set of research papers was examined in detail. While the overall structure of these studies is similar—typically addressing observational data, dynamical modelling, estimation procedures, and the assessment of accuracy—there are notable differences among them. These differences include the choice of ephemerides, the treatment of older datasets, and variations in the quality and quantity of available observations. Moreover, the satellites under investigation differ significantly in their dynamical characteristics and observational histories, further influencing the approaches taken in each study.

Among the studies considered, [Jacobson \(2009\)](#) provides the most comprehensive determination of the dynamics and orbits of Triton, Nereid, and Proteus, and its results are widely used as a benchmark in subsequent research. Although more extensive and precise datasets are now available, no recent study has matched the depth of Jacobson's analysis. More recent works—such as [Tang et al. \(2020\)](#), [Wang et al. \(2023\)](#), and [Wang et al. \(2025\)](#)—focus primarily on Triton, incorporating newly acquired Earth-based observations reduced with the Gaia star catalogue, as well as Hubble Space Telescope (HST) data. In addition, [Wang et al. \(2025\)](#) extends the analysis to estimate Neptune's tidal dissipation factor using a broad dataset. In contrast, [Brozović and Jacobson \(2022\)](#) investigates the irregular satellites of Uranus and Neptune (excluding Triton), emphasizing their long-term orbital evolution and identifying optimal epochs for future observations. Meanwhile, [Emelyanov et al. \(2023\)](#) evaluates the contribution of Gaia data to the study of faint irregular satellites of Jupiter, Saturn, and Uranus. Although Gaia astrometry for Triton exists, it has not yet been employed in a dedicated dynamical study to the author's knowledge.

Observational Data

Across all studies, the primary datasets comprise Earth-based astrometric observations (spanning from the late 19th century to the present), Voyager 2 imaging data, Gaia astrometry, and Hubble Space Telescope (HST) observations.

[Jacobson \(2009\)](#) incorporates the most comprehensive observational dataset available at the time of publication. This includes extensive historical ground-based observations, stellar occultations, HST observations of Proteus, Voyager 2 imaging, and radiometric tracking. Importantly, Jacobson reprocessed the Voyager 2 raw images, converting them from the older FK4 reference frame to the International Celestial Reference Frame (ICRF) using the UCAC2 star catalogue. This transformation improved the accuracy of the already high-precision Voyager data for Triton and Nereid.

Subsequent works—such as [Tang et al. \(2020\)](#), [Wang et al. \(2023\)](#), and [Wang et al. \(2025\)](#)—draw on a similar dataset, primarily using the extensive astrometric catalogue compiled by [Yuan et al. \(2021\)](#), which was reduced with the Gaia DR2 star catalogue. This significantly improved the precision of Earth-based astrometry. These studies exclude stellar occultations and Voyager radiometric tracking, but they incorporate 1,306 additional Earth-based observations collected between 2009 and 2016. Furthermore, [Wang et al. \(2023\)](#) and [Wang et al. \(2025\)](#) include high-accuracy HST observations of Triton that were not available to Jacobson.

In contrast, [Brozović and Jacobson \(2022\)](#) works with a dataset similar to that of [Jacobson \(2009\)](#) but adopts a different weighting strategy—down-weighting older Earth-based observations rather than applying explicit biases. Brozović also applies biases to the reduced astrometric data of Voyager images of Nereid (as published by [Jacobson \(1991\)](#)) instead of reprocessing the raw images, as done in [Jacobson \(2009\)](#).

Finally, [Emelyanov et al. \(2023\)](#) is the only study in the reviewed set to rely directly on high-precision Gaia data. To maintain consistency in observational weights, Emelyanov limits the inclusion of older ground-based observations to those obtained after 1996.

Dynamical Models and Ephemerides

Across all studies, the dynamical models include Neptune's point mass and extended gravity field (terms J_2 and J_4), as well as third-body perturbations from the Sun and other planets. However, the planetary ephemerides used for these perturbations vary between studies, and in some cases are explicitly compared to assess their impact on orbit fitting accuracy (e.g., [Emelyanov et al., 2023](#)).

The choice of planetary ephemerides differs slightly from one study to another. [Jacobson \(2009\)](#) employs DE421, while [Brozović and Jacobson \(2022\)](#), [Wang et al. \(2023\)](#), and [Wang et al. \(2025\)](#) use DE440. [Tang et al. \(2020\)](#) adopts INPOP19a, and [Emelyanov et al. \(2023\)](#) uses DE431.

The treatment of Neptune's rotational dynamics also varies. [Jacobson \(2009\)](#) develops a complete model for Neptune's pole precession. By contrast, [Tang et al. \(2020\)](#), [Wang et al. \(2023\)](#), and [Wang et al. \(2025\)](#) use simplified Euler equations combined with fitted parameters from [Jacobson \(2009\)](#). [Brozović and Jacobson \(2022\)](#) omits Neptune's pole precession entirely, since the study focuses on distant irregular satellites for which the effect is negligible.

There are also differences in the treatment of satellite masses. In [Jacobson \(2009\)](#), Proteus and Nereid are modelled as massless test particles. [Brozović and Jacobson \(2022\)](#) analyses Nereid separately from Triton. The more recent studies by [Tang et al. \(2020\)](#), [Wang et al. \(2023\)](#), [Wang et al. \(2025\)](#), and [Emelyanov et al. \(2023\)](#) include the gravitational effects of Neptune's inner satellites indirectly, by augmenting the J_2 and J_4 terms of Neptune's gravity field.

Fitting Procedure and Accuracy Estimation

The complexity of the fitting procedures varies significantly among the studies. [Jacobson \(2009\)](#) applies the most comprehensive approach, estimating a broad set of parameters that includes the system gravitational parameter (GM), Neptune's J_2 and J_4 gravity terms, pole precession parameters, the total angular momentum of the system, numerous observational biases, and the initial state vectors of Triton, Nereid, and Proteus.

In contrast, [Tang et al. \(2020\)](#), [Brozović and Jacobson \(2022\)](#), and [Emelyanov et al. \(2023\)](#) estimate only the initial state vector of the satellite under study. [Wang et al. \(2023\)](#) expands the parameter set to include GM , J_2 , J_4 , and Triton's initial state vector. [Wang et al. \(2025\)](#) further extends the estimation to Neptune's tidal quality factor Q , while assuming a constant Love number k_2 .

All studies employ numerical integration schemes appropriate for their dynamical models, and all report that no significant numerical error accumulates over their integration time spans.

The methods used to assess the accuracy of the fitted solutions also differ. [Jacobson \(2009\)](#) evaluates formal uncertainties by down-weighting or removing specific batches of observations and examining the resulting changes in the solutions. These tests indicate that the adopted dynamical model and data weighting are appropriate. Jacobson

also includes consider parameters in the estimation, which increase the reported uncertainties but remain consistent with expected values.

Brozović and Jacobson (2022) and Emelyanov et al. (2023) propagate uncertainties through covariance analysis—Brozović uses covariance matrix mapping, while Emelyanov applies covariance Monte Carlo simulations. In contrast, Tang et al. (2020), Wang et al. (2023), and Wang et al. (2025) perform Monte Carlo analyses on a more limited set of dynamical parameters to assess ephemeris uncertainty.

Discussion and Notable Conclusions

Comparisons between modern solutions and earlier work highlight both continuity and refinement in the understanding of the Neptunian system. Jacobson (2009) established a benchmark orbital solution for Triton, Nereid, and Proteus, identifying changes in mean motion across all three satellites—though the origin of these changes remains unresolved. Differences in Nereid’s out-of-plane motion were attributed to inclination variations, while discrepancies in Triton’s and Proteus’s precession rates were linked to Neptune’s extended gravity field.

Subsequent studies build on this foundation with expanded datasets and updated dynamical models. Tang et al. (2020) determined that Triton’s orbit relative to Neptune is largely unaffected by the choice of planetary ephemeris (e.g., DE430), while its absolute position in space exhibits greater sensitivity. The inclusion of Voyager 2 data continues to strongly constrain the solution at the initial epoch, as expected from its high precision. Similarly, Wang et al. (2023) confirmed that Triton’s orbit is most sensitive to Neptune’s J_2 term, with negligible effects from J_4 or GM . However, they note that J_2 and Neptune’s pole precession parameters are strongly coupled, making their simultaneous estimation problematic. Extending this work, Wang et al. (2025) incorporated Neptune’s tidal dissipation factor Q , finding a lower value than previous assumptions and reporting a reduction in RMS residuals. This reduction is most pronounced for 19th-century observations, which are known to be affected by systematic biases in Earth orientation models. The study does not explicitly address the potential influence of these biases on the reported improvement.

The long-term stability of Neptune’s satellites has been addressed by Brozović and Jacobson (2022), who examined the irregular satellites of both Neptune and Uranus over 30,000 years. Their precessing ellipse models capture the secular evolution of orbital precession rates, though accuracy can degrade for satellites near or beyond the Hill sphere. These models nevertheless provide a valuable framework for predicting optimal epochs for future observations.

Finally, Emelyanov et al. (2023) demonstrated improvements in predicted ephemeris accuracy through the incorporation of Gaia astrometry into dynamical models. Their analysis shows that Gaia-based reductions substantially enhance the precision of the predicted orbits and further validate the suitability of the adopted weighting and fitting strategies.

Taken together, these studies demonstrate that modern ephemerides for Neptune’s satellites are increasingly constrained by improved astrometric precision, particularly from Voyager 2, HST, and Gaia. While Triton’s orbit remains the most precisely characterized, the coupling of certain dynamical parameters (notably J_2 and pole precession) continues to limit parameter estimation accuracy. Long-term stability analyses of irregular satellites highlight dynamical complexities beyond the Hill sphere.

The breadth of topics covered in this review reflects the wide range of open questions in the field. The present study does not attempt to address all of them; instead, it focuses on three specific aspects that are either underexplored or insufficiently documented in the existing literature: the sensitivity of Triton’s orbit to dynamical parameters (particularly Neptune’s pole orientation), the design and assessment of observation weighting strategies, and the limitations of orbit reconstruction from Earth-based astrometric data alone. The full set of research subtopics identified during the initial phase of this study, together with a mapping to the research questions actually addressed, is provided in Appendix A.

Gaps in Modern Research

The preceding sections have reviewed the methods of observation, dynamical modelling, and parameter estimation for natural satellites, followed by an assessment of modern practices, their strengths, and their limitations. Understanding the state of current research is essential for defining focused research questions for the present study. This work aims to address specific gaps in the existing literature, with the intent of contributing both methodological and scientific value. A comparison of results with similar studies, along with the identification and interpretation of discrepancies, forms a core part of this process.

The most comprehensive investigation of the Neptunian system remains Jacobson (2009), which incorporated nearly all data available at the time and simultaneously estimated a wide set of dynamical and observational parameters.

As discussed earlier, this work serves as a benchmark against which subsequent studies are often evaluated.

Since then, significant advances have been made, particularly in the accuracy of star catalogues—formerly a dominant source of error. Additional observational data are now available from both ground-based campaigns and space missions such as Gaia and the Hubble Space Telescope. Furthermore, the digitization of older photographic plates [Robert et al. \(2025\)](#) has substantially improved the precision of historical observations, often reaching the intrinsic accuracy limits of the original measurement techniques.

The continued growth in both the quantity and quality of observations opens opportunities to constrain additional dynamical parameters. For instance, [Wang et al. \(2025\)](#) investigates Neptune’s tidal dissipation factor Q through its effect on Triton’s orbit. Constraints on the masses of Neptune’s inner moons represent another potential research direction, as their gravitational influence can introduce measurable perturbations on Triton’s motion. To the author’s knowledge, no dedicated study has yet addressed this specific problem.

Another notable gap concerns the treatment of observational data. In many modern studies, procedures for handling observational errors—particularly for pre-CCD era measurements—are described only briefly. Transparent documentation of these steps, especially bias corrections and weighting strategies for older data, would improve reproducibility and facilitate cross-study comparisons.

In summary, while modern research has made substantial progress in the modelling of the Neptunian system, the availability of improved observational datasets and unresolved dynamical questions provide clear opportunities for further investigation. The present study is designed with these gaps in mind, aiming to address selected open questions and to contribute results that can be meaningfully compared with established benchmarks.

2.4.3. Conclusion and Potential Research Questions

Observational Data

The primary objective of this study is to incorporate all available astrometric data relevant to Triton’s orbit. This includes both historical and recent Earth-based observations, as well as measurements from space-based observatories such as Gaia and the Hubble Space Telescope (HST), in addition to the close-encounter imaging data obtained by Voyager 2.

Due to time constraints, stellar occultation and mutual event observations will not be included in the present analysis, as their reduction requires a significantly different treatment from that of other astrometric data. The Earth-based and Gaia datasets used in this study are summarized in [Table 2.1](#) and [Table 2.2](#). HST observations of Neptune’s moons are provided by [Showalter et al. \(2019\)](#), while [Jacobson \(1991\)](#) supplies reduced astrometric measurements derived from Voyager 2 imaging data.

Estimated Parameters

As discussed previously, numerous parameters could be considered in the estimation process. Foremost among these is the initial state vector of Triton, which is included as a fitted parameter in all prior studies. The reference epoch for the simulation will be set at the time of Voyager 2’s closest approach to Triton, as this represents the epoch with the most precise observational constraints.

The next most significant parameters are Neptune’s pole precession and the higher-degree spherical harmonic coefficients J_2 and J_4 . Triton’s orbit is strongly influenced by Neptune’s extended gravity field, while Neptune’s pole precession is itself affected by Triton’s dynamics. This coupling makes the simultaneous estimation of these parameters essential for accurately reproducing Triton’s motion.

The combined gravitational parameter GM of the Neptune–Triton system is another key driver of Triton’s orbit. Although observational data have shown limited sensitivity to this parameter ([Wang et al., 2023](#)), reproducing earlier estimates would be valuable to confirm or challenge this conclusion.

Perturbations from the Sun and other planets also play a significant role in Triton’s dynamics. The positions of these bodies are taken from planetary ephemerides, but because multiple ephemerides exist—derived from different models and not always open source—assessing their influence on the solution will be necessary.

Tidal dissipation in Neptune due to Triton is an additional important effect. Following the approach of [Wang et al. \(2025\)](#), a lower bound on Neptune’s tidal dissipation factor Q could be estimated. However, the ratio $\frac{Q}{k_2}$, combining Q with Neptune’s Love number k_2 , may be a more physically meaningful quantity to evaluate, as there is no direct measurement of both parameters.

Finally, the gravitational influence of Neptune’s inner moons remains an open question. Their small size and close orbits mean their masses and positions are uncertain. While their net effect on Triton’s orbit may be small, exploring their contribution remains a potential gap in current research. Even if results are inconclusive, a formal assessment would help quantify their importance.

Dynamical and Observation Models

To estimate the parameters described above, both dynamical and observational models must be developed. The open-source Tudat software will be used for the numerical integration of Triton’s dynamics, with potential modifications to accommodate different planetary ephemerides and alternate formulations of Neptune’s pole precession dynamics.

The observational model will be constructed following the approaches used in previous works [Dahmani \(2024\)](#) and [Fayolle \(2025\)](#). Special consideration will be required for heterogeneous datasets such as older photographic plates, micrometric observations, and Gaia astrometry. Photographic plate data may require corrections for scale factor and reference star catalogue. Micrometric observations are reported in the Earth rotation reference frame and therefore necessitate the application of appropriate Earth rotation and atmospheric correction models.

Older observations also present challenges related to time standards and conventions, which are not always well documented. Where possible, conversions will be performed as described in [Emelyanov \(2021\)](#). If such conversions cannot be applied, a time bias correction or a down-weighting of the affected observations will be incorporated into the model.

Analysis

The analysis will begin by evaluating the impact of individual observational datasets on their effectiveness in constraining the selected dynamical parameters. Such an evaluation is also useful for diagnosing potential deficiencies in the observational model, such as incorrect weighting strategies or unaccounted-for biases.

Several previous studies have employed techniques such as down-weighting or removing specific data batches to verify their proper integration into the estimation process. In this work, a similar approach will be adopted, supplemented by the formulation of an “efficiency parameter” that quantifies the contribution of each dataset to the overall solution. For instance, the high-precision imaging data from Voyager 2 provides exceptionally strong constraints, but only over a brief observation window of a few months. In contrast, older micrometric data—though significantly less precise and subject to systematic biases—extends the temporal coverage to epochs well outside the most densely observed periods of recent decades. Gaia observations also warrant careful assessment: as Triton is brighter than the typical targets of Gaia, [Tanga et al. \(2023\)](#) reports a measurable degradation in astrometric accuracy for bright objects. An analysis of the different datasets, similar to that performed in [Fayolle \(2025\)](#), will be conducted.

A sensitivity analysis will also be performed on the estimated parameters to assess the reliability of the formal errors reported by the estimation procedure. Parameters for which the available observations offer limited sensitivity will be treated as consider parameters: these will not be estimated directly in the fit, but their uncertainties will be propagated into the final covariance. Their impact on the orbital solution will be analysed separately. A key part of this study will involve identifying which dynamical parameters should be treated as consider parameters and quantifying their effects.

Finally, an analysis of Triton’s long-term orbital evolution will be undertaken. Such simulations can provide insight into the origin of Triton and the dynamical history of Neptune’s satellite system. They can also be used to predict optimal future observation windows—such as intervals when Triton’s motion is nearly perpendicular to the line of sight, maximizing sensitivity to along-track and out-of-plane components. Predictions of mutual events or stellar occultations will also be considered, as these can provide additional high-precision constraints on the system.

Research Questions Formulation

A comprehensive overview of the potential research subtopics, grouped into the categories of Data, Estimated Parameters, Modelling, and Analysis, is presented in [Table 2.3](#). For each subtopic, the table also provides an estimated benefit (in terms of potential scientific return) and an approximate difficulty of implementation.

One advantage of the proposed research structure is its modularity: subtopics can be added or omitted as time and resources allow, since most can be addressed independently. The table is intended as a planning tool to identify which topics offer the greatest scientific return for a given level of effort. Implementing all subtopics would represent the most comprehensive analysis of the Triton system to date.

To frame these subtopics as targeted goals, they can be formulated as specific research questions:

1. Data Evaluation and Processing

- (a) What are the benefits and limitations of each data type (e.g., micrometric, CCD, Gaia, Hubble, Voyager) in constraining Triton's orbit?
- (b) How should Earth-based and space-based observations be weighted, and what biases should be applied?
- (c) How sensitive is the orbital fit to the inclusion/exclusion of older (19th century) micrometric data?
- (d) What are the observational biases in historic datasets, and can they be identified and corrected systematically?
- (e) What are Gaia's brightness-related limitations when observing objects like Triton?
- (f) How should inconsistent time standards in older datasets be corrected or mitigated?

2. Estimation of Physical and Dynamical Parameters

- (a) How well can the initial state of Triton at the Voyager epoch be constrained using current observations?
- (b) What is the sensitivity of Triton's orbit to Neptune's pole precession and gravitational harmonics (J_2 , J_4)?
- (c) To what degree is the Neptune–Triton gravitational parameter (GM) constrained by current observations?
- (d) How strongly is Triton's orbit affected by tidal dissipation, and what lower bound can be placed on Neptune's Q or Q/k_2 ?
- (e) How do uncertainties in inner satellite masses and orbits impact Triton's orbit over long timescales?
- (f) What observational biases can be estimated with current data?

3. Modelling Framework and Technical Implementation

- (a) How does the choice of planetary ephemeris (e.g., DE421 vs. DE440) influence Triton's fitted orbit?
- (b) How can Tudat be adapted to simulate pole precession and tidal interactions?
- (c) What corrections must be applied to micrometric and photographic observations to align them with modern dynamical models?
- (d) How does Earth's rotation model and atmospheric refraction affect observational reduction for old data?
- (e) What level of model complexity is justified by the available data quality?

4. Analysis, Uncertainty, and Scientific Implications

- (a) Which data batches contribute most significantly to reducing uncertainties in fitted parameters?
- (b) Which parameters are weakly constrained and should instead be treated as considered parameters?
- (c) How robust are the uncertainties of the estimated parameters to different weighting strategies or data removal techniques?
- (d) Can linearized covariance propagation accurately represent orbit uncertainty growth over time?
- (e) What is the forecasted uncertainty in Triton's orbit in the next 10, 20, or 30 years based on current models?
- (f) Is the developed model accurate enough to predict mutual events or star occultations?
- (g) What does the long-term evolution of Triton's orbit suggest about its origin or Neptune's dynamical history?

While [Table 2.3](#) summarizes the potential subtopics, together with their expected benefits and implementation difficulty. These rankings are qualitative and intended to help prioritize efforts during the study.

In summary, the structured list of research questions and prioritized subtopics provides a clear roadmap for the study. By selecting topics with the highest ratio of potential benefit to implementation difficulty, the research can maximize its impact while ensuring feasibility within the available time frame.

Table 2.3: Overview of Subtopics with Estimated Benefit and Difficulty

Category	Subtopic	Benefit (1–5)	Difficulty (1–5)
Data	Earth-Based CCD	3	1
	Photographic Plates	4	3
	Micrometric	5	5
	Voyager Astrometric Data	5	2
	Gaia Data	4	5
	Hubble Data	4	2
Estimated Parameters	Initial State	5	1
	Pole Precession	4	3
	Extended Body (J_2, J_4)	4	2
	Tidal Dissipation	3	4
	Inner Satellites	2	5
	Observational Biases	5	5
Modelling	Point Mass, Extended Gravity, Third Bodies	5	1
	Pole Precession Dynamics	5	3
	Observation Vector Geometry	5	3
	Earth Rotation Model	4	5
	Atmospheric Refraction	4	4
Analysis	Data Usefulness Metric	5	5
	Parameter Uncertainty Estimation	5	4
	Considered Parameter Effects	3	3
	Bias Estimation Confidence	5	4
	Dynamical Evolution Study	3–5	3–5

3

Methodology

This chapter documents the specific implementation choices and deviations from standard practice that form the basis of the analyses presented in this work. Where a method is implemented as described in [chapter 2](#), only the key parameters and settings are reported without re-derivation. The chapter is structured as follows: the dynamical model is presented first, followed by the estimation framework, the observational data, and finally the structure of the three analyses performed in this study. Key terminology used throughout this work, including the distinction between propagation, estimation, and analysis, is defined in [section 3.4](#).

3.1. Dynamical Model

The dynamical model describes the motion of Triton with respect to Neptune. The gravitational attraction of all planets and the Sun is modelled as point mass perturbations. Neptune’s gravitational field additionally includes zonal harmonic coefficients up to degree four, which have been found to have a significant impact on Triton’s orbit. The orientation of Neptune’s pole is modelled using the IAU standard pole precession and nutation formulation, as the pole orientation directly affects the direction of the zonal harmonic perturbations. Finally, the numerical integrator used to propagate the equations of motion is described and validated. The following subsections address each of these components in turn.

3.1.1. Neptune Gravitational Field Zonal Harmonics

The zonal harmonic coefficients of Neptune’s gravitational potential, whose role in perturbing Triton’s orbit is discussed in [chapter 2](#), are taken from [Jacobson \(2009\)](#) and listed in [Table 3.1](#). Both J_2 and J_4 are also considered as estimatable parameters in the Simulated Observations Analysis ([subsection 3.4.1](#), results in [chapter 4](#)).

Table 3.1: Neptune gravitational field zonal harmonic coefficients from [Jacobson \(2009\)](#).

Parameter	Value
$J_2 (\times 10^{-6})$	3408.42853071795 ± 4.5
$J_4 (\times 10^{-6})$	-33.398917590066 ± 2.9

3.1.2. Pole Precession Model

The pole orientation of Neptune is modelled following the IAU standard formulation ([Archinal et al., 2018](#)). While the formulations of [Jacobson \(2009\)](#) and [Wang et al. \(2025\)](#), derived in [chapter 2](#), are based on dynamical propagation of the torques exerted by Triton on Neptune, the IAU standard formulation adopts a much simpler parametric form and is already implemented in Tudat. The right ascension and declination of the pole are given by:

$$\alpha = \alpha_0 + \dot{\alpha}_0 \cdot T + \alpha_1 \sin(N) \tag{3.1}$$

$$\delta = \delta_0 + \dot{\delta}_0 \cdot T + \delta_1 \cos(N) \tag{3.2}$$

where T is the time from the reference epoch in Julian centuries and N is the longitude of Neptune’s ascending node, advancing at 52.316 per century. The model parameters include the pole position (α_0, δ_0) , pole precession rate $(\dot{\alpha}_0, \dot{\delta}_0)$, set to zero in both models), and pole libration amplitudes (α_1, δ_1) . Two sets of parameter values are used in this work, from Archinal et al. (2018) and Jacobson (2009), and are summarised in Table 3.2.

Table 3.2: Neptune pole model parameter values from Archinal et al. (2018) and Jacobson (2009).

Parameter	IAU 2015	Jacobson (2009)	Unit
α_0	299.36	299.46	deg
δ_0	43.46	43.40	deg
$\dot{\alpha}_0$	0.0	0.0	deg/century
$\dot{\delta}_0$	0.0	0.0	deg/century
α_1	0.70	0.635	deg
δ_1	−0.51	−0.462	deg

3.1.3. Numerical Integrator

An overview of numerical integrators, their limitations, and their applications is presented in chapter 2. In this work, the equations of motion are integrated using a fixed-step Runge–Kutta–Fehlberg 7(8) integrator.

A fixed step size is appropriate here given Triton’s near-circular orbit, which does not exhibit the rapid variation in orbital velocity that would otherwise require adaptive step-size control. The chosen step size of 60 minutes is comfortably above the threshold below which numerical precision loss begins to dominate the integration error.

This threshold is higher than would be expected for a Neptune-centred frame. The integration is performed in the Solar System Barycentre (SSB) frame, as this is required for the construction of the Earth-based observation geometry. In the SSB frame, the absolute position coordinates of Triton are on the order of the Neptune–Sun distance (~ 30 AU), several orders of magnitude larger than the Neptune-centred coordinates. This leads to floating-point precision loss in the relative position of Triton with respect to Neptune, raising the effective numerical noise floor and consequently the minimum viable step size. Step sizes below approximately 30 minutes were found to exhibit this numerical noise dominance. Validation of the integrator settings is presented in Appendix B.

3.2. Estimation Framework

3.2.1. Weighted Least Squares Formulation

The weighted least-squares formulation and the derivation of the normal equations are presented in chapter 2. The implementation here follows the standard formulation directly, using the Tudat estimation framework. The estimated parameter vector, a priori constraints, and consider parameters are discussed in chapter 2 and are not re-derived here; specific parameter choices for each analysis are documented in the relevant results chapters (chapter 4 and chapter 5).

3.2.2. Weighting Strategy

Throughout this section, an *observation file* refers to a discrete dataset submitted by a single observatory for a specific observational campaign, corresponding to a single entry in Table 3.4. Multiple files may originate from the same observatory if the data are presented in a different observational campaign.

Three base weighting schemes are considered, each differing in the granularity at which observation quality is assessed. In the following, $r_{i,f}^\alpha$ and $r_{i,f}^\delta$ denote the right-ascension and declination residuals for the i -th observation in file f . The notation is presented for the right-ascension component; the declination component follows identically.

Per-File Weighting

The per-file weighting scheme assigns a single weight to all observations within a given observation file. For a file f containing N_f observations, the root-mean-square error is:

$$\epsilon_f^\alpha = \sqrt{\frac{1}{N_f} \sum_{i=1}^{N_f} (r_{i,f}^\alpha)^2} \quad (3.3)$$

The weight is then defined as the inverse square of the RMSE:

$$w_f^\alpha = \frac{1}{(\epsilon_f^\alpha)^2} \quad (3.4)$$

This scheme assumes that all observations within a file are drawn from a Gaussian distribution centred at zero, with comparable accuracy across observations and no correlated errors. It provides a straightforward approach for weighting files according to their relative accuracies. However, it performs poorly when observation quality varies significantly within a file — a common occurrence when files span extended time periods during which atmospheric conditions fluctuate, observatory instruments degrade, or data-processing algorithms evolve. Additionally, it does not account for temporal clustering: when multiple observations are collected within a short timespan, their errors are likely correlated, which can artificially reduce the computed RMSE and inflate the file weight.

Per-Timeframe Weighting

The per-timeframe weighting scheme addresses the limitations of per-file weighting by dividing each file into temporal segments and computing weights independently for each. For a file f containing T_f timeframes, where timeframe t has n_t observations, the RMSE for each timeframe is:

$$\epsilon_t^\alpha = \sqrt{\frac{1}{n_t} \sum_{i=1}^{n_t} (r_{t,i}^\alpha)^2} \quad (3.5)$$

To account for correlated errors between observations within a single night, the per-timeframe RMSE is scaled by the square root of the number of observations, which is nearly equivalent to treating the observations per night as a single composite observation:

$$\tilde{\epsilon}_t^\alpha = \epsilon_t^\alpha \cdot \sqrt{n_t} \quad (3.6)$$

The weight for each observation in timeframe t is then:

$$w_t^\alpha = \frac{1}{(\tilde{\epsilon}_t^\alpha)^2} \quad (3.7)$$

This approach allows the estimator to automatically down-weight portions of a file that exhibit poorer quality while maintaining higher weights for higher-quality timeframes. However, when observations are sparsely distributed and most timeframes contain only one or a few observations, the assigned weights become dominated by individual residuals rather than robust statistical measures. To mitigate this, a lower bound is imposed on the minimum allowable RMSE per timeframe ($\epsilon_t \geq v_{\min} = 10$ mas).

Scaled Per-File Weighting

The scaled per-file weighting scheme combines the strengths of both previous approaches by computing timeframe-level statistics while assigning a single weight per file. For each timeframe t in file f , the scaled RMSE is computed as in the per-timeframe scheme. The file-level RMSE is then obtained by treating each timeframe as a single composite observation:

$$\epsilon_f^\alpha = \sqrt{\frac{1}{T_f} \sum_{t=1}^{T_f} (\tilde{\epsilon}_t^\alpha)^2} \quad (3.8)$$

with the final file weight defined as $w_f^\alpha = 1/(\epsilon_f^\alpha)^2$. This scheme provides a middle ground between per-file and per-timeframe weighting: by computing and scaling the RMSE at the timeframe level, it captures temporal variations in observation quality and correlated errors, while assigning weights at the file level prevents individual observations from artificially skewing the result. This is the weighting scheme adopted for all estimations in the pole estimation analysis of [chapter 5](#).

Hybrid Weighting Schemes

The hybrid weighting schemes combine the file-level weight w_f from the scaled per-file scheme (ID v2) with the timeframe-specific weight w_t from the per-timeframe scheme (TF), producing a weight for each timeframe t in file f that reflects both the overall file quality and the local observation quality.

Two combination rules are considered. The geometric mean hybrid assigns:

$$w_{t,f}^{(g)} = \sqrt{w_f \cdot w_t} \quad (3.9)$$

This formulation ensures that the resulting weight is always bounded between the two constituent weights and gives equal multiplicative influence to each. If either the file-level or timeframe-level weight is very small (indicating poor quality), the hybrid weight is pulled down accordingly, while a high weight in one component cannot fully compensate for a low weight in the other.

The arithmetic mean hybrid assigns:

$$w_{t,f}^{(a)} = \frac{w_f + w_t}{2} \quad (3.10)$$

Unlike the geometric mean, the arithmetic mean is dominated by the larger of the two weights. In practice, the per-timeframe weights tend to be several orders of magnitude larger than the per-file weights, causing the arithmetic mean to be skewed toward the per-timeframe component. This makes the arithmetic mean hybrid less balanced than the geometric mean variant.

Both hybrid schemes inherit the cap on the minimum allowable per-timeframe RMSE ($v_{\min} = 10$ mas) from the per-timeframe scheme. The derivation of the corresponding observation uncertainties for the hybrid schemes is provided in [Appendix C](#).

3.2.3. Observation Uncertainty from Weights

To facilitate interpretation and comparison of the weighting schemes, an observation uncertainty parameter v is defined from the computed weights. For any weight w with units of $[\text{rad}^{-2}]$, the uncertainty is:

$$v = \frac{1}{\sqrt{w}} \quad (3.11)$$

This parameter has units of $[\text{rad}]$ (or equivalently $[\text{mas}]$ when converted), making it directly interpretable as the expected scatter or precision of the observations. For the scaled per-file scheme, this reduces to:

$$v_f = \sqrt{\frac{N_f}{T_f}} \sqrt{\frac{1}{N_f} \sum_{i=1}^{N_f} (r_i)^2} \quad (3.12)$$

which is the per-file RMSE scaled by $\sqrt{N_f/T_f}$, reflecting the effective number of independent observations after accounting for temporal clustering.

3.2.4. Observational Goodness-of-Fit Metrics

Three scalar metrics are used to assess the quality of the least-squares orbit estimation based on the residuals: the cost function J , the weighted root-mean-square (WRMS) residual, and the unweighted root-mean-square (RMS) residual. Each is computed from the observation residuals at both the initial (pre-estimation) and final (post-estimation) iteration, allowing a direct assessment of the improvement achieved by the estimator.

Observation Residuals

Let N denote the total number of astrometric observations. Each observation i yields a right-ascension residual $\Delta\alpha_i$ and a declination residual $\Delta\delta_i$, obtained as the difference between the observed and modelled angular positions:

$$\Delta\alpha_i = \alpha_i^{\text{obs}} - \alpha_i^{\text{calc}}, \quad \Delta\delta_i = \delta_i^{\text{obs}} - \delta_i^{\text{calc}} \quad (3.13)$$

expressed in radians. The statistical weight assigned to each residual is:

$$w_{\alpha,i} = \frac{1}{\sigma_{\alpha,i}^2}, \quad w_{\delta,i} = \frac{1}{\sigma_{\delta,i}^2} \quad (3.14)$$

where $\sigma_{\alpha,i}$ and $\sigma_{\delta,i}$ are the one-sigma uncertainties of observation i in right ascension and declination, respectively.

Cost Function

The least-squares cost function is the weighted sum of squared residuals over all observations and both angular components:

$$J = \sum_{i=1}^N [w_{\alpha,i} \Delta\alpha_i^2 + w_{\delta,i} \Delta\delta_i^2] \quad (3.15)$$

Because J is dimensionless (the weights carry units of rad^{-2}), it provides a scale-free measure of fit quality. Larger values of J indicate a poorer fit; a perfect fit would yield $J = 0$.

Weighted Root-Mean-Square Residual

The weighted root-mean-square residual ρ_w is a weighted quadratic mean of the residuals, normalised by the total weight:

$$\rho_w = \sqrt{\frac{\sum_{i=1}^N w_{\alpha,i} \Delta\alpha_i^2 + \sum_{i=1}^N w_{\delta,i} \Delta\delta_i^2}{\sum_{i=1}^N w_{\alpha,i} + \sum_{i=1}^N w_{\delta,i}}} \quad (3.16)$$

Denoting the total weight $W \equiv \sum_{i=1}^N w_{\alpha,i} + \sum_{i=1}^N w_{\delta,i}$, the numerator of (3.16) is identical to J , so:

$$\rho_w^2 = \frac{J}{W} \implies J = \rho_w^2 W \quad (3.17)$$

For a well-calibrated solution the residuals scatter at the level of the assigned uncertainties, so $\rho_w \approx \sigma_{\text{typical}}$. The analogous dimensionless goodness-of-fit criterion is the reduced chi-squared:

$$\frac{J}{\nu} = \frac{\sum_{i=1}^N \left[\frac{\Delta\alpha_i^2}{\sigma_{\alpha,i}^2} + \frac{\Delta\delta_i^2}{\sigma_{\delta,i}^2} \right]}{\nu} \approx 1 \quad (3.18)$$

where $\nu = 2N - k$ is the number of degrees of freedom and k is the number of estimated parameters. Converting ρ_w to milliarcseconds ($1 \text{ rad} \approx 206,264,806 \text{ mas}$) yields a value that can be compared directly with the astrometric precision of the observations.

Unweighted Root-Mean-Square Residual

The unweighted root-mean-square residual ρ treats every observation equally regardless of its assigned uncertainty, providing a weight-independent measure of residual scatter:

$$\rho = \sqrt{\frac{1}{2N} \sum_{i=1}^N (\Delta\alpha_i^2 + \Delta\delta_i^2)} \quad (3.19)$$

The factor of $2N$ accounts for the two independent angular components across N observations. After conversion to milliarcseconds, ρ can be compared directly with ρ_w to assess whether the weighting scheme is well-matched to the actual scatter.

Relation Between the Metrics

The three metrics carry complementary information. When all weights are equal ($w_{\alpha,i} = w_{\delta,i} = w$ for all i), ρ_w reduces to ρ . More generally, the ratio ρ_w/ρ indicates whether the weighting scheme up-weights or down-weights the observations with the largest residuals. [Table 3.3](#) summarises the three metrics.

Table 3.3: Summary of goodness-of-fit metrics used in this work.

Metric	Formula	Normalisation	Unit
Cost function J	$\sum_i (w_{\alpha,i} \Delta\alpha_i^2 + w_{\delta,i} \Delta\delta_i^2)$	None	dimensionless
Weighted RMS ρ_w	$\sqrt{J/W}$	Total weight W	rad (mas)
Unweighted RMS ρ	$\sqrt{(\sum_i \Delta\alpha_i^2 + \sum_i \Delta\delta_i^2)/2N}$	N_{obs}	rad (mas)

3.3. Astrometric Dataset and Initial Residual Analysis

This section presents a brief summary of the dataset utilized in this study. A comprehensive overview of all observations is provided in [Table 3.4](#), which includes the observatory name, the NSDC identifier, the Minor Planet Center (MPC) observatory code, and the total number of observations (N_{obs}). The data spans from 1963 to 2025 and a temporal visualisation of the O-C residuals with NEP097 are presented in [Figure 3.2](#).

The table further specifies the observation types, restricting the scope to absolute and relative astrometric observations, and lists the Root Mean Square (RMS) residuals of the observations relative to the NEP097 kernel. All observational data were retrieved from the NSDC website ([Institute of Celestial Mechanics and Computation of Ephemerides \(IMCCE\), 2025](#))¹.

The temporal distribution of these observations and their quantity per file are illustrated in [Figure 3.1](#). Furthermore, the residuals calculated with respect to the NEP097 kernel are visualized in [Figure 3.2](#).

3.3.1. O-C Residuals with respect to the NEP097 Kernel

The residuals relative to the NEP097 kernel are shown in [Figure 3.2](#). Because certain observations exhibit significantly larger residuals than others within the same file, they are rejected based on a fixed maximum threshold, indicated by red crosses. The total count of rejected observations per file ID is listed in [Table 3.5](#). The resulting distribution without these rejections is "flat," meaning there is no visible trend in either Right Ascension or Declination. While [Jacobson \(2009\)](#) suggests that some files may exhibit per-night biases, such an analysis was not conducted in this research.

The file 689 nm077 (shown in brown in the figure) exhibited a clear bias of approximately 0.2" in Declination. Although it was initially intended to treat this bias as an estimated parameter, time constraints necessitated a manual adjustment by subtracting 0.2" from the observations. This bias is not reported from any of the consulted papers using the same dataset ([Jacobson \(2009\)](#), [Wang et al. \(2023\)](#), [Dolgakov and Pavlov \(2025\)](#)). The result of this correction is shown in [Figure 3.3](#), which brings the residuals closer to the zero-line.

3.4. Structure of Analyses

Before describing the analyses themselves, it is useful to establish precise definitions for the terminology used throughout this report.

A *propagation* refers to the numerical integration of the equations of motion governing a dynamical system, producing a trajectory for Triton given an initial state and a set of model parameters. An *estimation* is an iterative procedure consisting of a bundle of propagations, in which the result of each propagation is used to update the estimated parameters, their formal uncertainties, and the correlations between them through a weighted least-squares procedure. Two types of observations are used as input to the estimation: *simulated observations*, derived from the three-dimensional positions of Triton provided by the NEP097 ephemeris, and *real astrometric observations*, comprising both absolute and relative angular measurements collected over the historical observational record. An *analysis* refers to a collection of estimations performed under a consistent observational dataset, in which one or more of the dynamical model parameters, initial states, estimated parameter sets, or weighting strategies are varied, and the results of which are collectively evaluated and compared.

¹<https://nsdb.imcce.fr/obspos/bnepomae.htm>

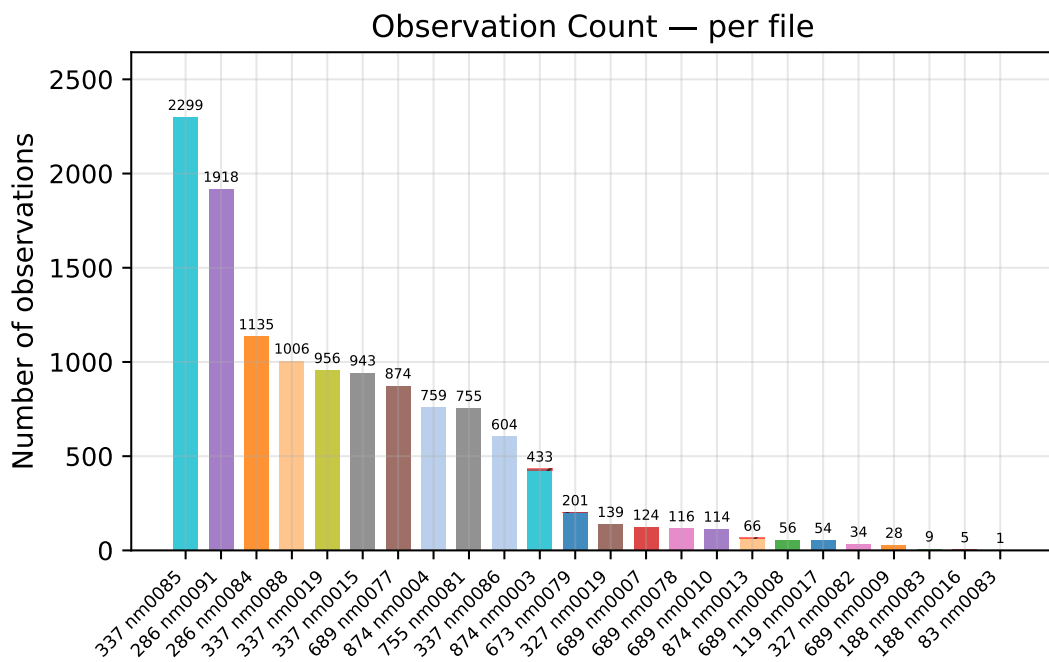
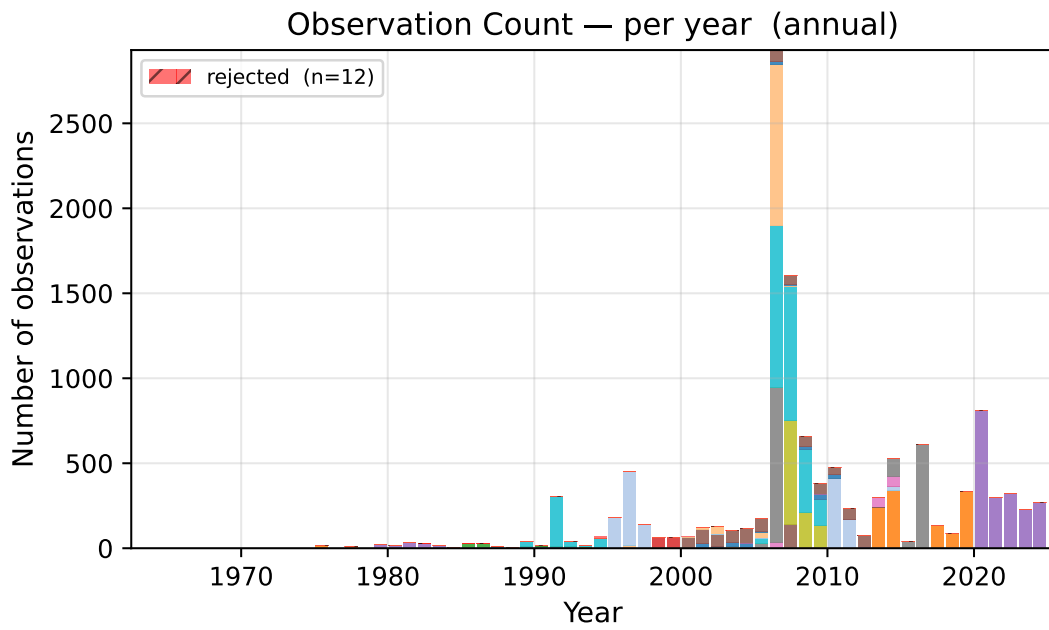


Figure 3.1

Table 3.4: Summary of astrometric observation datasets used in this work. NSDC Listing gives the dataset identifier in the Natural Satellites Data Centre archive. MPC Code is the three-digit Minor Planet Center observatory code. N_{obs} is the total number of astrometric observations. Obs. Type indicates whether the original measurements are relative (Rel.) or absolute (Abs.) astrometry. RMS O–C RA and RMS O–C Dec are the root-mean-square observed minus computed residuals in right ascension and declination, respectively, evaluated against the NEP097 ephemeris [’].

Observatory	NSDC Listing	MPC Code	N_{obs}	Obs. Type	RMS O–C vs NEP097 [’] RA	Dec
Holosiivskiyi district-Kyiv	nm0083	83	1	Abs.	0.951	0.013
Abastuman	nm0017	119	54	Abs.	0.448	0.414
Majdanak	nm0016	188	5	Abs.	0.137	0.328
Majdanak	nm0083	188	9	Abs.	0.600	0.551
Yunnan Observatory	nm0084	286	1135	Abs.	0.032	0.029
Yunnan Observatory	nm0091	286	1918	Abs.	0.029	0.034
Peking Observatory, Xinglong Station	nm0019	327	139	Abs.	0.112	0.112
Peking Observatory, Xinglong Station	nm0082	327	34	Abs.	0.063	0.061
Sheshan, formerly Zo-Se	nm0015	337	943	Abs.	0.098	0.044
Sheshan, formerly Zo-Se	nm0019	337	956	Abs.	0.063	0.059
Sheshan, formerly Zo-Se	nm0085	337	2299	Abs.	0.056	0.051
Sheshan, formerly Zo-Se	nm0086	337	604	Abs.	0.062	0.049
Sheshan, formerly Zo-Se	nm0088	337	1006	Abs.	0.046	0.040
Table Mountain Observatory, Wrightwood	nm0079	673	201	Abs.	0.176	0.078
U.S. Naval Observatory, Flagstaff	nm0007	689	124	Abs.	0.123	0.114
U.S. Naval Observatory, Flagstaff	nm0008	689	56	Rel.	0.027	0.071
U.S. Naval Observatory, Flagstaff	nm0009	689	28	Rel.	0.082	0.095
U.S. Naval Observatory, Flagstaff	nm0010	689	114	Rel.	0.039	0.075
U.S. Naval Observatory, Flagstaff	nm0077	689	874	Abs.	0.133	0.258
U.S. Naval Observatory, Flagstaff	nm0078	689	116	Abs.	0.156	0.233
Optec Observatory	nm0081	755	755	Abs.	0.129	0.013
Observatorio do Pico dos Dias, Itajuba	nm0003	874	433	Rel.	0.498	0.200
Observatorio do Pico dos Dias, Itajuba	nm0004	874	759	Rel.	0.167	0.212
Observatorio do Pico dos Dias, Itajuba	nm0013	874	66	Abs.	0.168	0.244

Table 3.5: Observations rejected by the outlier filter, grouped by dataset identifier.

Dataset ID	N_{rej}
673_nm0079	1
874_nm0003	10
874_nm0013	1




Figure 3.2: O-C Residuals NEP097 colored per file id, rejected observations are colored red and marked with x.




Figure 3.3: O-C Residuals with and without manual bias overlaid

Before any analysis results are discussed, a brief summary of the astrometric data used is presented in ??, where the observed-minus-computed (O–C) residuals with respect to the NEP097 kernel are described, along with the temporal spread of the data, the rejection algorithm, and the bias correction applied to a specific dataset.

3.4.1. Simulated Observations Analysis

In order to assess the limits of the dynamical model, an analysis based on simulated observations from the NEP097 kernel is performed. The simulated observations are direct three-dimensional positions of Triton with respect to Neptune, extracted from the SPICE kernel NEP097 at a cadence of 3 hours and placed homogeneously throughout the full estimation period. In essence, they constitute a complete reconstruction of the orbit of Triton from the NEP097 kernel at a resolution of one sample per 3 hours.

Feeding these positions as observations to the estimation procedure, without weights, biases, or added noise, should result in clear and rapid convergence. Unlike real data, which consist of two angular observables irregularly spaced in time and subject to biases and random noise, the simulated observations exhibit none of these complications. The estimator therefore only needs to account for differences between the two dynamical models, making it easier to identify discrepancies between the dynamics as formulated in Tudat and those implicit in the NEP097 kernel, and potentially revealing which dynamical parameters absorb most of the difference between the two solutions.

Additionally, the sensitivity of Triton’s orbit to various dynamical parameters can be assessed. In practice, real data may be insufficient to constrain certain parameters, either due to limited accuracy or incomplete temporal coverage. The simulated observations are, by construction, infinitely accurate and complete, allowing full orbital reconstruction. The dynamical parameters considered in this analysis are the Neptune pole model parameters as described in [subsection 3.1.2](#), the gravitational parameters of Triton and Neptune, and Neptune’s extended gravity field spherical harmonics (J_2 and J_4).

The analysis involves three distinct steps. First, the initial propagation (without estimation) is compared to NEP097. Second, an estimation of the initial state only is performed using the simulated observations. Third, multiple estimations of the initial state are conducted in combination with the dynamical parameters listed above. All estimations are compared against the kernel, which is treated as ground truth. The best solution, judged by the root-mean-square difference with respect to NEP097, is then carried forward for the analyses with real data.

Beyond summary statistics, RSW plots of promising or notable solutions are examined in detail, and the updated estimated parameters are compared to their initial values and to literature values to assess physical plausibility.

Two points are worth noting. First, two different pole models are used: IAU 2015 ([Archinal et al. \(2018\)](#)) and Jacobson 2009 ([Jacobson \(2009\)](#)). Second, high correlations are expected between certain parameter pairs, for example between pole model parameters, between spherical harmonic coefficients, and between the gravitational parameters of Triton and Neptune. Several estimations are deliberately designed to include highly correlated parameters, in order to investigate whether the estimator remains numerically stable, whether it converges, and whether the results are physically meaningful.

The results are presented in [chapter 4](#). The best-performing dynamical parameter combination is selected and referred to throughout as the `FitPole` initialisation, which serves as the starting point for the subsequent analyses.

3.4.2. Pole Estimation with Astrometric Data Analysis

The second analysis, presented in [chapter 5](#), applies the dynamical model and parameter set selected on the basis of the Simulated Observations Analysis ([subsection 3.4.1](#), results in [chapter 4](#)) together with the weighting scheme selected in the Weight Strategy Analysis ([subsection 3.4.3](#), results in [chapter 6](#)).

It is important to note the circular dependency between the Pole Estimation Analysis and the Weight Strategy Analysis. An accurate pole estimation requires a well-defined weighting scheme, while a properly selected weighting scheme requires formal errors that are consistent with the actual difference from NEP097, which in turn requires the uncertainties in the most sensitive dynamical parameters to be included in the estimation. The pole parameters, specifically the pole position and pole libration rate, are identified as the most influential in the Simulated Observations Analysis. The workflow proceeds as follows: the Weight Strategy Analysis is first performed with only the initial state of Triton estimated; the best-performing weighting strategy is selected; the Pole Estimation Analysis is then carried out; and finally the Weight Strategy Analysis is rerun with both the initial state of Triton and the Neptune pole parameters estimated. It is therefore decided to present the Pole Estimation Analysis before discussing the weight selection in detail.

Two initial conditions are considered. The first, named `IAUPole`, uses an initial state of Triton taken directly from the NEP097 kernel and the pole model from Archinal et al. (2018). The second, named `FitPole`, corresponds to the best-performing estimation from the Simulated Observations Analysis. The weights are constructed from the residuals of each respective initialisation.

Multiple estimations with different combinations of pole parameters as estimated quantities are performed and compared in terms of difference with the NEP097 kernel, formal errors, and residuals with respect to the astrometric data. Additional correlation and stability analyses of the estimated parameters are also presented. The full list of estimations in this analysis is given in Table 4.2.

The best-performing estimation, based on the comparison criteria described above, is selected as the final estimation of this report and further compared to literature values in chapter 7.

3.4.3. Weight Strategy Analysis

The third analysis, presented in chapter 6, investigates the effect of the weighting scheme on the estimation results using the full astrometric dataset. The three base weighting schemes and their derived hybrid variants, as described in subsection 3.2.2, are compared in terms of difference with the NEP097 kernel, formal error consistency, and astrometric residuals. The best weighting strategy is selected and is the one already used to produce the results in the Pole Estimation Analysis (chapter 5).

The results presented in chapter 6 consist of two separate sub-analyses. The first, named `WA-IAU`, is based on the initial propagation with Triton's initial state taken from the NEP097 kernel and the pole model from Archinal et al. (2018), estimating only the initial state of Triton. The second, named `WA-FIT-PL`, is initialised from the best-performing estimation of the Pole Estimation Analysis (chapter 5), which has a substantially different initial state of Triton and Neptune pole model. The weights for each sub-analysis are derived from the residuals of their respective initialisations and therefore differ slightly.

The bulk of the chapter discusses the differences between estimations of the `WA-FIT-PL` analysis, whose naming conventions are presented in Table 6.2.

4

Parameter Estimation with Simulated Observations

In this chapter the analysis of fitting the dynamical model to simulated observations from the benchmark NEP097 are shown and discussed. The simulated observations are 3D positions homogenously spaced in time every 3 hours for the full period of the estimation time directly taken from the NEP097 kernel without any added noise or bias. The precise methodology of this analysis is presented in [subsection 3.4.1](#). The naming convention for all estimations is described in the [Table 4.2](#). Two possible initializations are defined and their initial values are presented in [Table 4.1](#).

Table 4.1: Initial values of all estimated parameters used in the analysis. The Triton initial state is taken from the NEP097 ephemeris at 2006-10-01. Gravitational parameters (GM) and spherical-harmonic coefficients (\bar{C}_{20} , \bar{C}_{40}) are from the DE440 SPICE kernel and the [Jacobson \(2009\)](#), respectively; these values are identical for both rotation models. Pole parameters differ between the IAU 2015 ([Archinal et al. \(2018\)](#)) and [Jacobson \(2009\)](#) models. Degree-2 libration terms (α_2 , δ_2) are [Jacobson \(2009\)](#) specific and have no IAU 2015 counterpart.

Parameter	Group	Unit	IAU 2015	Jacobson 2009
X	Position	km	$2.744\,134 \times 10^5$	$2.744\,134 \times 10^5$
Y	Position	km	$-2.444\,959 \times 10^4$	$-2.444\,959 \times 10^4$
Z	Position	km	$-2.234\,963 \times 10^5$	$-2.234\,963 \times 10^5$
\dot{X}	Velocity	km s^{-1}	-1.875 074	-1.875 074
\dot{Y}	Velocity	km s^{-1}	-3.472 847	-3.472 847
\dot{Z}	Velocity	km s^{-1}	-1.922 285	-1.922 285
GM_N	Gravity	$\text{km}^3 \text{s}^{-2}$	$6.835\,103 \times 10^6$	$6.835\,103 \times 10^6$
GM_T	Gravity	$\text{km}^3 \text{s}^{-2}$	$1.428\,495 \times 10^3$	$1.428\,495 \times 10^3$
α_0	Pole Position	deg	$2.993\,600 \times 10^2$	$2.994\,609 \times 10^2$
δ_0	Pole Position	deg	$4.346\,000 \times 10^1$	$4.340\,393 \times 10^1$
$\dot{\alpha}_0$	Pole Rate	deg yr^{-1}	0.000 000	0.000 000
$\dot{\delta}_0$	Pole Rate	deg yr^{-1}	0.000 000	0.000 000
α_1	Pole Librations	deg	$7.000\,000 \times 10^{-1}$	$6.353\,980 \times 10^{-1}$
δ_1	Pole Librations	deg	$-5.100\,000 \times 10^{-1}$	$-4.616\,274 \times 10^{-1}$
\bar{C}_{20}	Spherical Harmonics	—	$-1.524\,296 \times 10^{-3}$	$-1.524\,296 \times 10^{-3}$
\bar{C}_{40}	Spherical Harmonics	—	$1.113\,297 \times 10^{-5}$	$1.113\,297 \times 10^{-5}$
α_2	Pole Librations	deg	—	$-2.420\,971 \times 10^{-3}$
δ_2	Pole Librations	deg	—	$8.794\,384 \times 10^{-4}$

Table 4.2: Estimations naming conventions used in figures for Simulated Observations Analysis.

Label	Description	Neptune Pole Model
prop. IAU	Pure propagation, no estimation	IAU 2015
prop. Jac.	Pure propagation, no estimation	Jacobson 2009
state	Triton initial state only ($x, y, z, \dot{x}, \dot{y}, \dot{z}$)	IAU 2015
state (Jac.)	Triton initial state only	Jacobson 2009
state+GM_T	Initial state + GM_{Triton}	IAU 2015
state+GM_N	Initial state + GM_{Neptune}	IAU 2015
state+GM_TN	Initial state + GM_{Triton} + GM_{Neptune}	IAU 2015
state+SH.	Initial state + spherical harmonics (C_{20}, C_{40})	IAU 2015
state+GM_TN+SH.	Initial state + GM_{Triton} + GM_{Neptune} + spherical harmonics (C_{20}, C_{40})	IAU 2015
state+pos.	Initial state + pole position (α_0, δ_0)	IAU 2015
state+rot.	Initial state + pole rotation rate ($\dot{\alpha}_0, \dot{\delta}_0$)	IAU 2015
state+pos.+rot.	Initial state + pole position + rotation rate	IAU 2015
state+lib.	Initial state + pole libration (α_1, δ_1)	IAU 2015
state+pos.+lib.	Initial state + pole position + libration	IAU 2015
state+full.	Initial state + all pole parameters (position + rotation rate + libration)	IAU 2015
state+pos. (Jac.)	Initial state + pole position (α_0, δ_0)	Jacobson 2009
state+rot. (Jac.)	Initial state + pole rotation rate	Jacobson 2009
state+lib1. (Jac.)	Initial state + pole libration, degree-1 terms	Jacobson 2009
state+lib2. (Jac.)	Initial state + pole libration, degree-2 terms	Jacobson 2009
state+pos.+lib1. (Jac.)	Initial state + pole position + libration (degree-1)	Jacobson 2009
state+pos.+lib2. (Jac.)	Initial state + pole position + libration (degree-2)	Jacobson 2009
state+full. (Jac.)	Initial state + all pole parameters, degree-2 terms	Jacobson 2009
state+SH.+full.	Initial state + spherical harmonics (C_{20}, C_{40}) + all pole parameters	IAU 2015
all	All parameters: state + GM_{Triton} + GM_{Neptune} + spherical harmonics + all pole parameters	IAU 2015

4.1. Initial Propagation vs. NEP097

The RSW differences relative to the NEP097 kernel for the initial propagation of the Triton dynamical model are presented in Figure 4.1. The figure compares two propagation cases, *prop. IAU* and *prop. Jac.*, as defined in Table 4.2. The specific RSW difference between these two pole models is further illustrated in Figure 4.2. To maximize observational data density, the initial state is set to October 1, 2006, with Triton’s state extracted directly from the NEP097 kernel to initialize the propagation.

Both propagations exhibit similar behavior relative to NEP097. A small, growing divergence in the radial (R) direction suggests a minor mismatch in gravitational parameters. The growth in the along-track (S) direction is consistent across both cases; this is expected, as even a slight mismatch in the semi-major axis accumulates over the orbital path. From the 2006 epoch back to the termination point in 1963, the accumulated error is approximately 10,000 km. While this represents a drift of roughly 0.7 km per day, it remains subtle: for perspective, a 1 arcsecond error in an Earth-based observation of Triton corresponds to approximately 22,000 km at Neptune’s average distance.

The most notable behaviour occurs in the cross-track (W) direction, where both models show significant growing differences compared to the NEP097 kernel. This suggests a mismatch in the satellite’s orbital inclination. As hypothesized in previous chapters, this likely results from uncertainties in Neptune’s pole model, which shifts the orientation of Neptune’s extended body gravity effect on Triton. Because the specific modeling and estimation choices for the NEP097 kernel are not publicly documented, it was initially assumed that the model followed the conventions described in Jacobson (2009). Surprisingly, the Jacobson (2009) pole model appears to have little influence on the cross-track divergence. The maximum difference between the two propagations is approximately 200 km—corresponding to roughly 10 mas in an Earth-based observation—which is too small to confidently assess given current observational accuracy.

In conclusion, these initial propagations demonstrate that the Jacobson (2009) and IAU 2015 (Archinal et al. (2018)) pole models differ only marginally. Furthermore, neither model perfectly aligns with the pole configuration used to generate the NEP097 kernel. A detailed investigation into the exact pole model used in NEP097 is provided in subsequent sections.

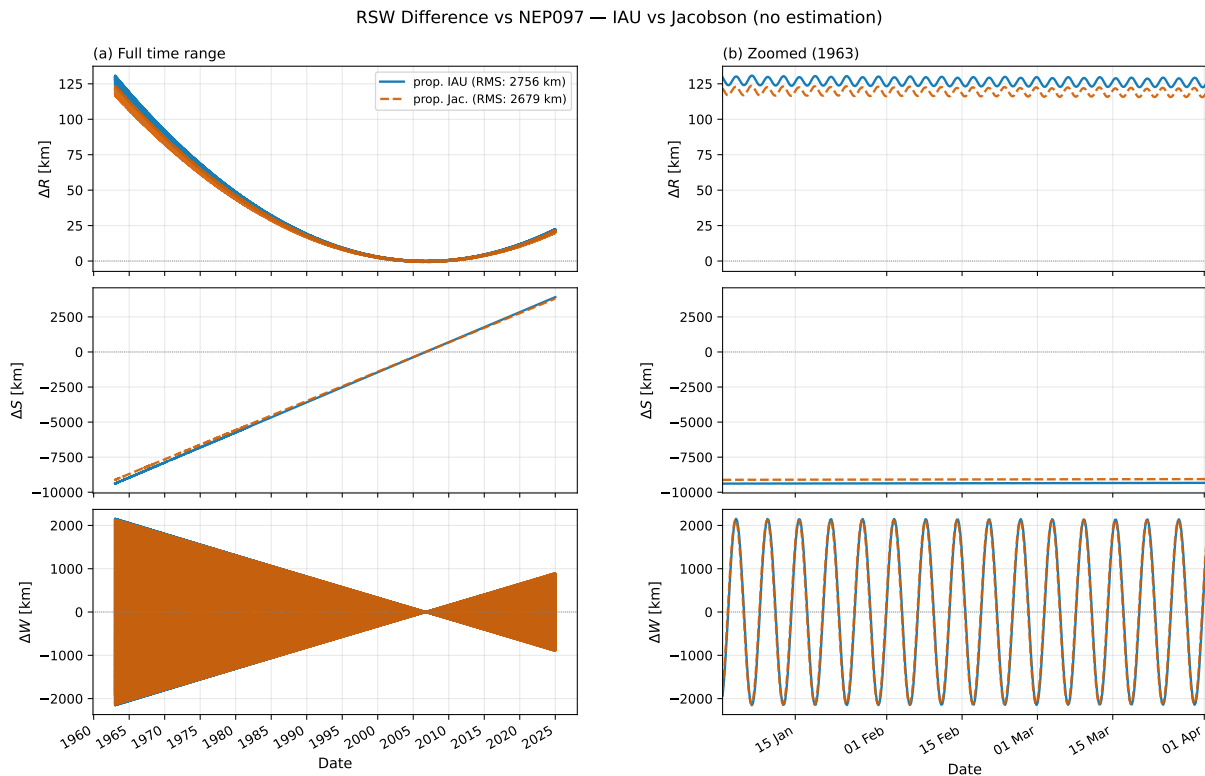


Figure 4.1: RSW Positional Difference with NEP097 kernel for initial propagations with Neptune pole models IAU 2015 and Jacobson as described in Table 4.2.




Figure 4.2: RSW Difference of the initial propagations with different Neptune pole models IAU and Jacobson as described in [Table 4.2](#).

4.2. Initial State Estimation

The first estimation case considers the adjustment of the initial state only. As previously discussed, the initial state is extracted from the NEP097 kernel; if this kernel is treated as the ground truth, any remaining divergence from NEP097 must be attributed to dynamical parameter mismatches between the two models.

As shown in [Figure 4.3](#), estimating the initial state significantly reduces the differences in the R and S directions compared to the initial propagation of the model (see [Figure 4.1](#)). There is also a notable reduction in the W direction, with the peaks decreasing by approximately 500 km relative to the non-estimated propagation.

There is statistically no significant difference between the two pole models under consideration. The resulting differences in R and S are at most a few kilometres, while the difference in W is less than 100 km; in terms of RMS, these differences represent only a fraction of a kilometre. It is important to reiterate that a 100 mas error in an Earth-based observation corresponds to a positional difference of roughly 2,200 km. Even assuming a highly optimistic average observational accuracy of 10 mas, a difference of 100 km remains well below the threshold of detection.

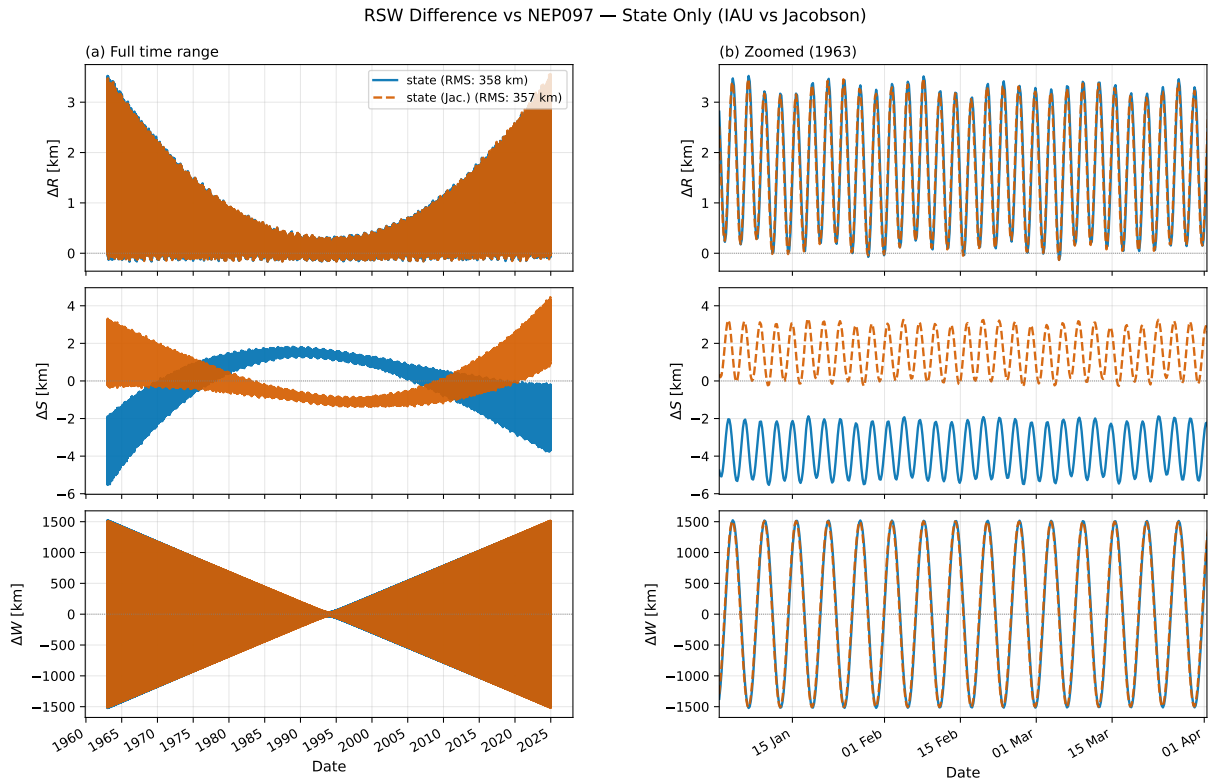


Figure 4.3: RSW Difference with NEP097 of initial state estimation with simulated observations for both pole models (naming convention defined in [Table 4.2](#))

The initial state adjustments in the RSW frame, resulting from the estimation process, are presented in [Figure 4.4](#). For both estimation cases, the updated state shows a correction of approximately 500 km in the W direction, while virtually no change is observed in the R and S positions. This strongly suggests that the primary source of uncertainty in the orbit is localized in the cross-track direction, and by extension, the orbital inclination.

It is evident that while the initial state is capable of absorbing dynamical mismatches in the R and S directions, it lacks the functional flexibility to fully compensate for the cross-track divergence. This supports the hypothesis that the remaining residuals are likely tied to the orientation of the primary's pole rather than a simple state error.

Finally, the correlation heatmaps for the state-only estimations are shown in [Figure 4.5](#). No excessively high correlations are observed in either estimation variant. The moderate correlations between specific position and velocity components are expected; in particular, the coupling between radial position and along-track velocity is a standard characteristic of orbital motion in such dynamical systems.

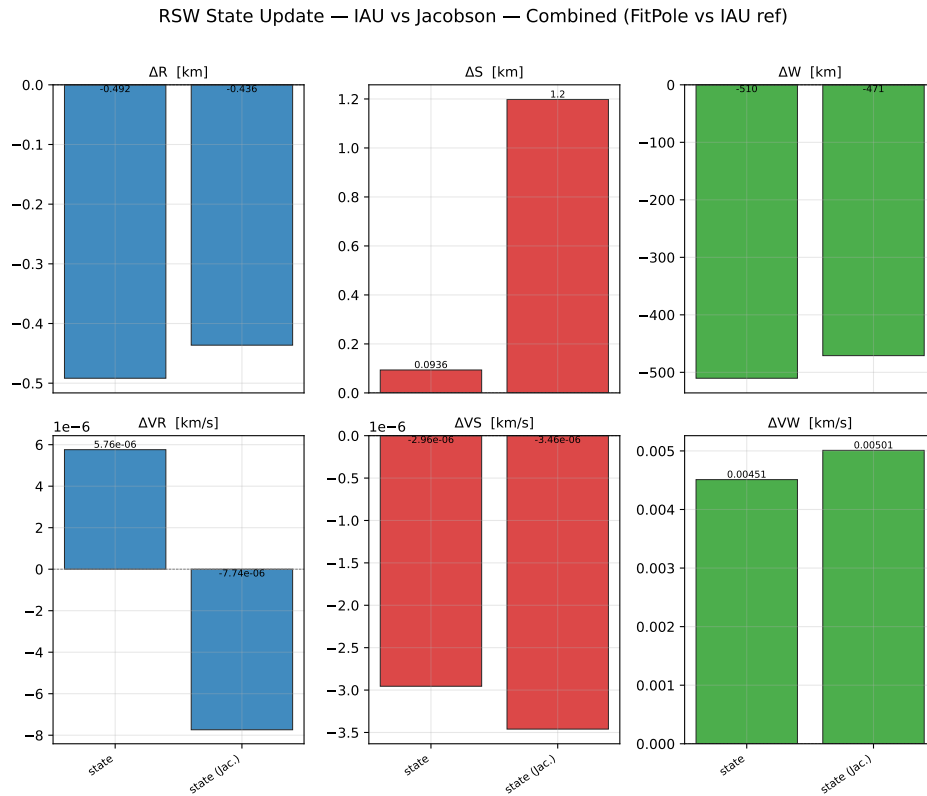




Figure 4.4: Parameter Update of Initial State Estimation

|Correlations| — State Only (IAU 2015)



(a) Neptune pole model: IAU 2015

|Correlations| — State Only (Jacobson 2009)



(b) Neptune pole model: Jacobson 2009

Figure 4.5: Correlation heatmaps for the initial state estimations comparing the (a) IAU 2015 and (b) Jacobson 2009 Neptune pole models.

4.3. Gravitational Parameter Estimations of Neptune and Triton

The gravitational parameters (GM) of Neptune and Triton were evaluated as potential estimation parameters. The resulting RMS differences with respect to the NEP097 kernel are shown in Figure 4.6, the updates to the Triton initial state in Figure 4.8a, and the gravitational parameter updates in Figure 4.8b.

Estimations that include gravitational parameters perform slightly worse in terms of RMS difference with respect to NEP097 compared to the state-only estimation *state*, and no clear explanation for this behaviour has been identified. The inclusion of GM does not lead to a reduction in residuals despite the estimator converging, suggesting that the gravitational parameters are not improving the fit in any meaningful way. The parameter updates in Table 4.3 show that when estimating GM_{Neptune} or GM_{Triton} alone, the estimator applies a shift of approximately $-3.1 \times 10^{12} \text{ m}^3\text{s}^{-2}$ to each respectively. When both are estimated simultaneously, this update is split exactly in half between the two parameters, and in all three cases the total change in system GM is identical. This strongly suggests that the data constrains only the total system GM rather than the individual contributions of Neptune and Triton separately. Notably, when estimating GM_{Triton} alone, the resulting value becomes negative, which is physically inadmissible.

The RSW time series difference with respect to NEP097 is presented in Figure 4.7. The solutions for all GM variants are indistinguishable from one another, confirming that all three cases converge to effectively the same mathematical solution. An interesting pattern is visible in the time series: the radial (R) direction exhibits an offset of approximately 60 km from the NEP097 kernel, while the along-track (S) direction shows a difference of only approximately 6 km. This suggests that the gravitational parameters are being used by the estimator to reduce the S direction error at the cost of a worse R direction fit. This trade-off is physically expected, as the radial and along-track directions are strongly coupled through the orbital dynamics. In the analytical two-body solution, the mean motion of a satellite depends only on the gravitational parameter and the semi-major axis, which in this case is effectively equivalent to the R direction given Triton's near-circular orbit. While the dynamical model employed here includes additional perturbations, their magnitudes are several orders lower than the point-mass gravitational attraction of Neptune, and the fundamental R-S coupling therefore remains dominant. The across-track (W) direction difference for all gravitational parameter estimations is indistinguishable from the *state* only case, which is equally expected as this direction is dynamically decoupled from the gravitational parameter.

The correlation heatmaps in Figure 4.9 show the correlations between the initial state and the gravitational parameters. A moderate correlation is present between the X positional component and the Y velocity component, consistent with the R-S coupling discussed above. The correlation matrix for *state+GM_TN* contains NaN values, therefore no figure is presented, indicating a numerical singularity arising from the extremely high correlation between the two gravitational parameters. This is physically expected and consistent with the equal partitioning of the GM update discussed previously. The other two estimation cases do not exhibit particularly high inter-parameter correlations; some positional and velocity components in fact show higher correlations than in the state-only case shown in Figure 4.5a.

In conclusion, the estimator converges when gravitational parameters are included, and three-dimensional positional data is sufficient to constrain the total system GM . However, the estimations do not outperform the *state* only case, and no explanation for this has been found. The update of approximately $-3.1 \times 10^{12} \text{ m}^3\text{s}^{-2}$, which corresponds to roughly 1.1 times the mass of Triton, does not necessarily imply a mismatch between the gravitational parameters of this solution and the NEP097 kernel. Rather, it suggests that the estimator is compensating for unmodelled perturbations or an imperfect dynamical model through the gravitational parameters. Estimating GM_{Triton} alone produces a negative mass, which is physically inadmissible, and estimating both simultaneously leads to a numerically singular estimation. GM_{Neptune} could in principle be included as an estimated parameter without causing numerical instability, but given that it does not improve the agreement with NEP097 and appears to compensate for other model deficiencies rather than reflecting a genuine physical signal, gravitational parameters are not included in the subsequent analyses.




Figure 4.6: RSW difference with NEP097: Statistical comparison of estimations involving gravitational parameters.

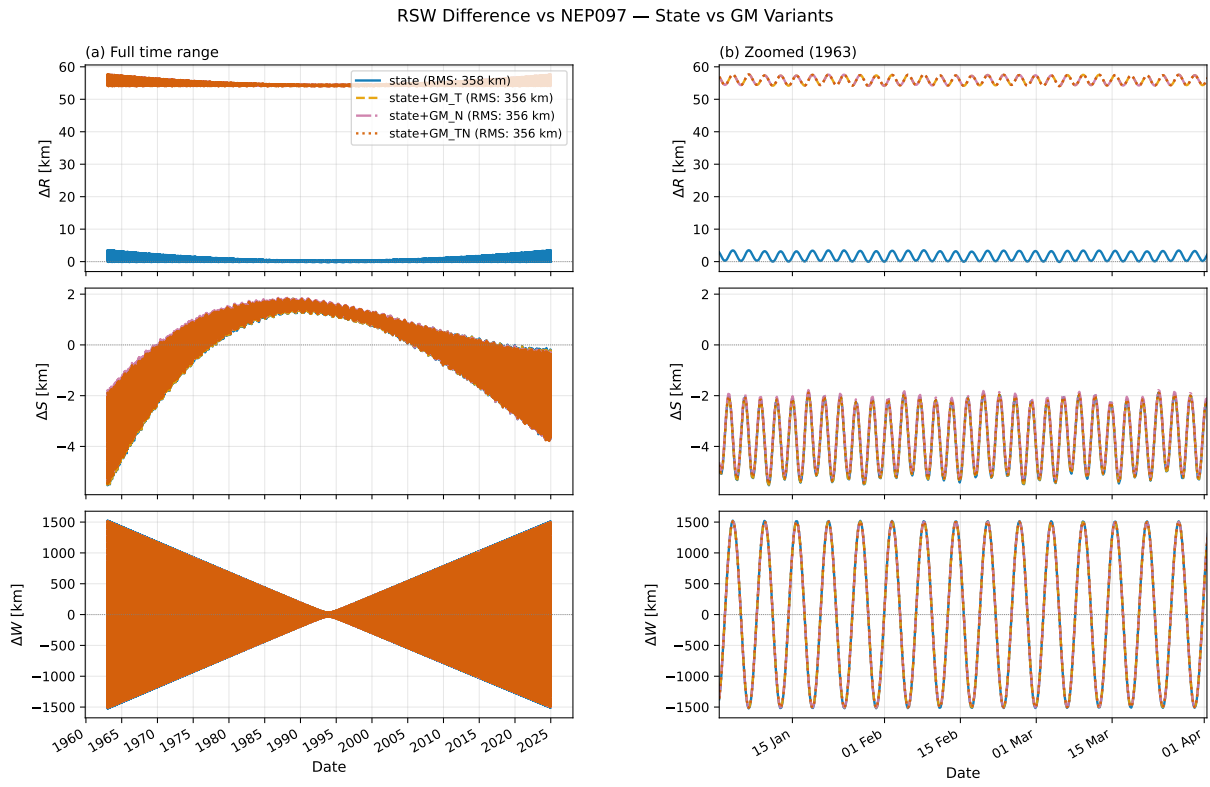


Figure 4.7: RSW difference with NEP097: Timeseries comparison of all GM estimation variants.





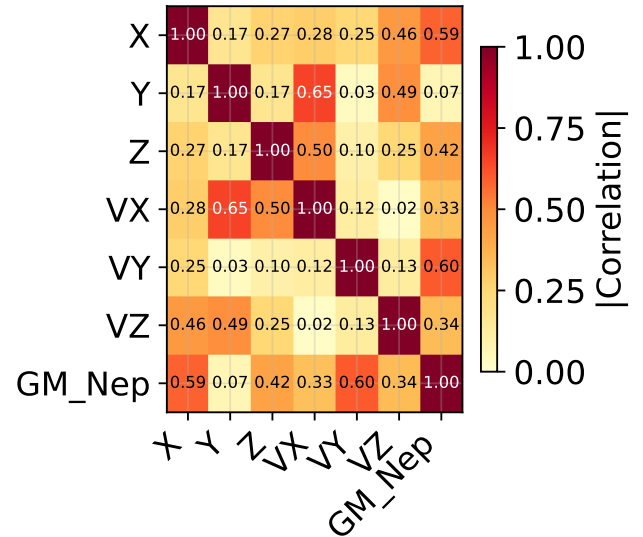
Figure 4.8: Comparison of parameter updates for the Triton initial state and gravitational parameters across GM variants.

[Correlations] — State + GM Triton (IAU)



(a) Correlations for state+GM_T.

[Correlations] — State + GM Neptune (IAU)



(b) Correlations for state+GM_N.

Figure 4.9: Correlation heatmaps for state estimations including (a) Triton’s gravitational parameter and (b) Neptune’s gravitational parameter.

Table 4.3: Estimated gravitational parameters for the GM simulation variants (IAU 2015 rotation model). Initial values are taken from the DE440 SPICE kernel. $\Delta = \text{Final} - \text{Initial}$; $\Delta[\%] = \Delta / |\text{Initial}| \times 100$.

Simulation	Parameter	Initial ($m^3 s^{-2}$)	Final ($m^3 s^{-2}$)	Δ	$\Delta[\%]$
GM_Triton_IAU	GM_T	$1.428\,495 \times 10^{12}$	$-1.706\,063 \times 10^{12}$	$-3.134\,559 \times 10^{12}$	-219.431%
GM_Neptune_IAU	GM_N	$6.835\,103 \times 10^{15}$	$6.831\,969 \times 10^{15}$	$-3.134\,558 \times 10^{12}$	-0.046%
GM_Both_IAU	GM_N	$6.835\,103 \times 10^{15}$	$6.833\,536 \times 10^{15}$	$-1.567\,273 \times 10^{12}$	-0.023%
	GM_T	$1.428\,495 \times 10^{12}$	$-1.387\,897 \times 10^{11}$	$-1.567\,285 \times 10^{12}$	-109.716%

4.4. High Level Analysis of Best Performing Estimations

Here are presented all estimations which performed well below the expected accuracy threshold of the observational data. The RMS difference relative to NEP097 is shown in Figure 4.10 for all estimations involving pole parameters and two estimations involving Neptune’s extended body gravitational parameters (J_2 and J_4). The naming convention

for these estimations is defined in [Table 4.2](#).

In both the IAU and Jacobson models, the initial pole rotation rates are defined as zero. While one might expect a free rotation rate to offer a similar dynamical degree of freedom to pole librations, this analysis shows that estimating these rates yields no measurable improvement over the initial state estimation. This indicates that a constant rotation rate is insufficient to model the observed perturbations in the pole orientation, whereas the periodic nature of the libration parameters provides a significantly better fit to the data. Consequently, estimations involving pole rotation parameters are excluded from further analysis and are provided in the Appendix for completeness.

Estimations including Neptune's extended body gravity (J_2 and J_4) are included because Neptune's pole position and its J_2/J_4 coefficients are expected to have similar effects on Triton's orbit. This coupling is evidenced by the fact that simultaneous estimation of the pole parameters and J_2/J_4 caused the estimator to fail to converge. While estimating J_2 and J_4 yields a similar fit to estimating pole parameters, it is slightly less accurate. The decision to prioritize the pole estimation over Neptune's extended gravity is based on the existence of direct J_2 and J_4 observations from the Voyager 2 mission [Jacobson \(2009\)](#); these values are assumed to be sufficiently accurate and will be treated as "considered" parameters in later stages.

Finally, comparing identical estimations using different initial pole models (IAU vs. Jacobson) revealed negligible differences. Therefore, the Jacobson-based estimations will not be pursued further, as they provide no functional advantage over the IAU model.




Figure 4.10: RMS difference relative to NEP097 for estimations including initial state, pole parameters, and Neptune's extended body gravity (SH).




Figure 4.11: RMS difference relative to NEP097 decomposed into RSW components for estimations including initial state, pole parameters, and Neptune’s extended body gravity (SH).

When analysing the RMS decomposed into RSW components (Figure 4.11), the `state + rot.` estimation performs similarly to the initial state estimation and notably worse than the other models across all directions. The choice of pole model (IAU vs. Jacobson) has no significant effect on the RMS in any direction. Furthermore, because the Jacobson Neptune pole model includes degree-2 librations, the RMS difference between estimating only degree-1 librations (`state + lib1. (Jac.)`) versus both degree-1 and 2 (`state + lib2. (Jac.)`) is negligible. This is consistent with the aggregate RMS analysis shown in Figure 4.10.

4.5. IAU Pole Parameter Estimations

This section evaluates pole parameter estimations using the IAU 2015 Neptune pole model. As established previously, rotation rate parameters are excluded from this analysis. The naming conventions for all estimation variants follow the definitions provided in Table 4.2.

4.5.1. RMS Analysis

The total and component-wise RSW RMS differences relative to NEP097 are presented in Figure 4.12 and Figure 4.13, respectively. Generally, the inclusion of additional pole parameters reduces the RMS residual. In the RSW decomposition, all estimations perform well in the radial (R) direction. However, the inclusion of pole librations involves a trade-off, slightly decreasing accuracy in the transverse (S) direction in exchange for a significant reduction in the normal (W) direction residuals. Notably, the `state + pos. + lib.` and `state + full` variants show negligible differences across all RMS metrics.




Figure 4.12: Total RMS difference relative to NEP097 for estimations including initial state and pole parameters.




Figure 4.13: RSW component decomposition relative to NEP097 for estimations including initial state and pole parameters.

4.5.2. Parameter Correlations and Conditioning


Correlation heatmaps for the various estimation variants are shown in [Figure 4.14](#) through [Figure 4.15b](#). Moderate correlations between position and velocity parameters are observed; this is expected due to the high coupling of radial and along-track components in this dynamical system.

When estimated separately, pole position and libration parameters do not exhibit high correlations. However, in the combined `state + pos. + lib.` estimation, the declination (δ_0) and the pole libration rate exhibit a correlation coefficient significantly close to one. Despite this, the condition numbers listed in [Table 4.4](#) indicate that the system remains numerically stable.

Table 4.4: Condition numbers of the correlation matrices for IAU 2015 variants.


Simulation	Estimated parameters	Condition number
state	state	4.7919×10^9
state+pos.	state, α_0, δ_0	5.4966×10^9
state+lib.	state, α_1, δ_1	5.9855×10^9
state+pos.+lib.	state, $\alpha_0, \delta_0, \alpha_1, \delta_1$	6.2704×10^9
state+full.	state, $\alpha_0, \delta_0, \dot{\alpha}_0, \dot{\delta}_0, \alpha_1, \delta_1$	1.6089×10^{10}

[Correlations] — State + Pole Position (IAU 2015)



(a) state + pos.

[Correlations] — State + Pole Librations (IAU 2015)



(b) state + lib.

Figure 4.14: Correlation heatmaps for separate pole position and libration estimations.

4.5.3. Residual Timeseries Analysis

The RSW residual timeseries relative to NEP097 for the state, state + pos. + lib., and state + full cases are presented in [Figure 4.16](#). A noticeable improvement is evident in all directions, particularly in across-track (W). This is expected, as the normal component is responsible for the orbital inclination, which is strongly influenced by the Neptune pole model via its extended body gravity.

The differences between the `state + pos. + lib.` and `state + full` variants are further inspected in [Figure 4.17](#). Given the expected noise levels of real observational data, these two high-fidelity models would be functionally indistinguishable.

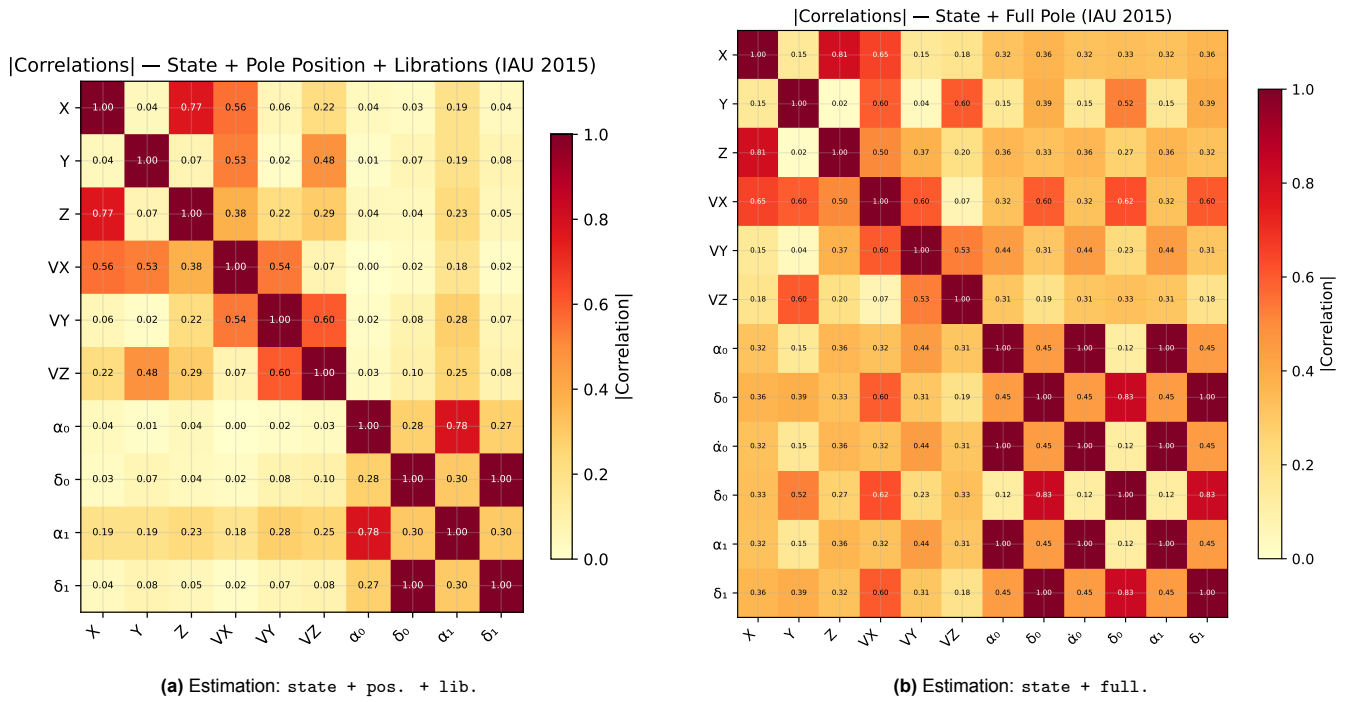


Figure 4.15: Correlation heatmaps for the combined IAU 2015 pole estimations, comparing the standard combined model (a) against the full model including rotation rates (b).

RSW Difference vs NEP097 — State, Pos.+Lib., Full (IAU 2015)

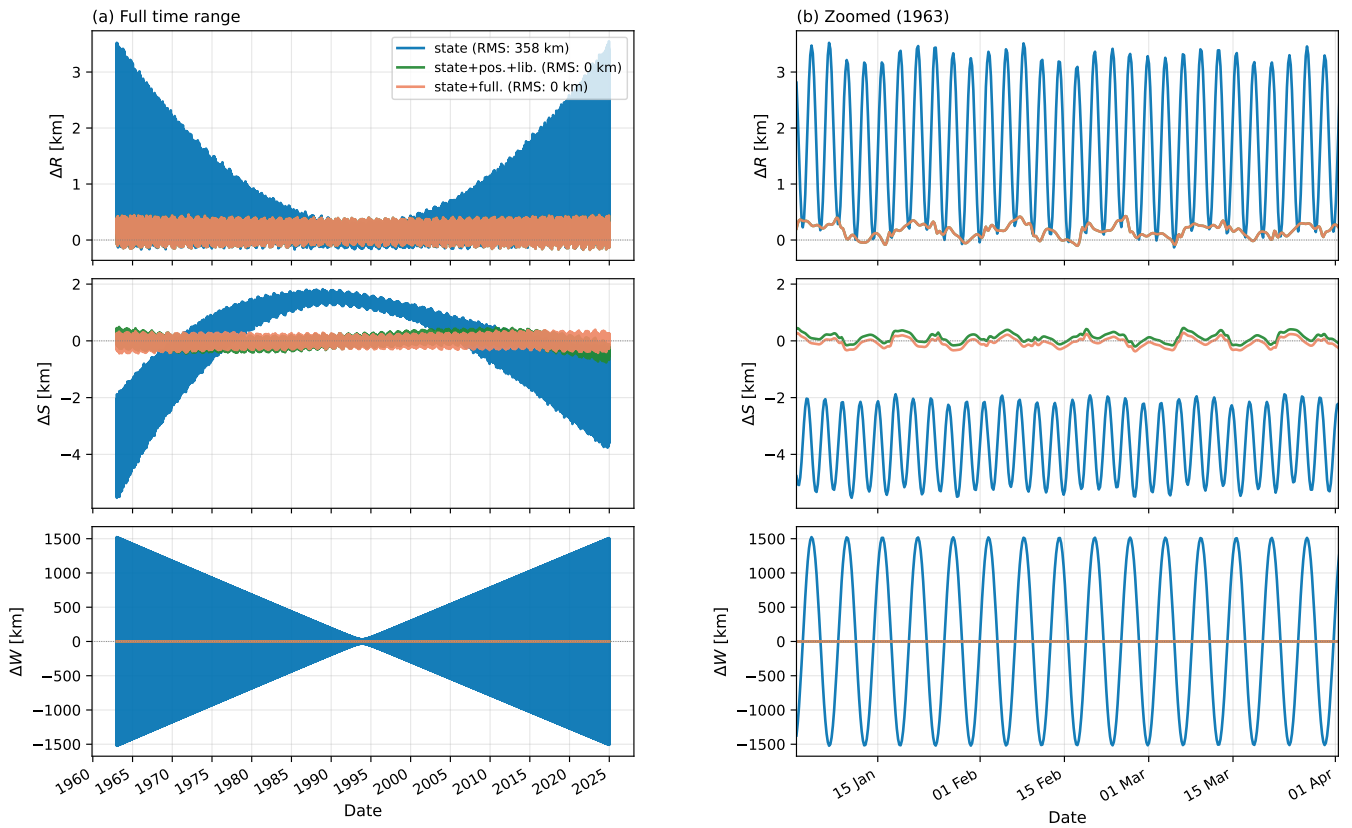


Figure 4.16: RSW Position difference with NEP097 for the baseline and top-performing pole estimations.

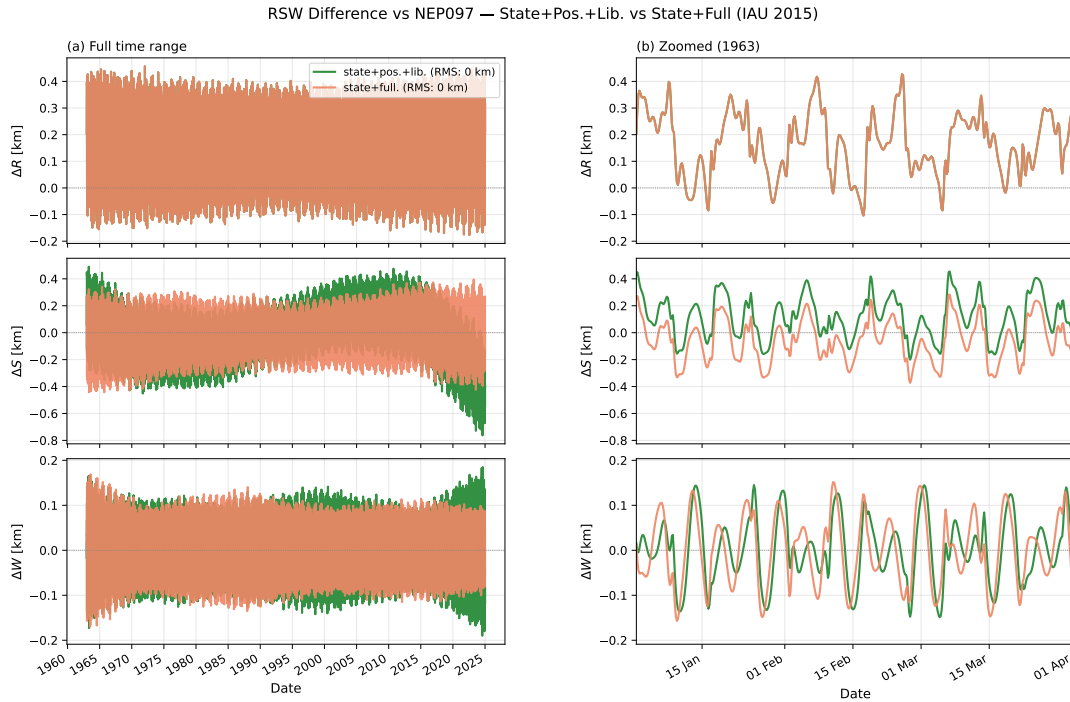


Figure 4.17: RSW Positional difference with NEP097 comparison between the state + pos. + lib. and state + full variants.

4.5.4. Parameter Updates and Physical Interpretation

Analyzing parameter updates ensures the solution remains physically plausible. [Figure 4.18](#) and its RSW decomposition ([Figure 4.19](#)) show that estimations including pole parameters require significantly smaller updates to the initial state than the state-only baseline. Crucially, the across-track (W) update is greatly reduced, as the pole model now accounts for residuals previously absorbed by the state.

Updates for pole position and libration are shown in [Figure 4.20](#) and [Figure 4.21](#). The top-performing variants show lower pole position updates by trading off an increase in libration parameters to achieve a better fit. As IAU 2015 does not report uncertainties, we compare our results to [Jacobson \(2009\)](#) in [Table 4.5](#).

Table 4.5: Comparison of estimated updates versus formal uncertainties from [Jacobson \(2009\)](#).

Parameter	Estimated Update (Δ)	Jacobson (2009) σ	Ratio (Δ/σ)
Pole Position (α_0, δ_0)	$\approx 0.13^\circ$	0.14°	≈ 0.93
Pole Libration (α_1, δ_1)	$\approx 1.80^\circ$	0.0001°	$\approx 18,000$

Our pole position update is remarkably consistent with Jacobson’s formal error ($\approx 1\sigma$). However, the libration update shows a massive mismatch. This may suggest that Jacobson’s libration formal errors are optimistic or that the NEP097 kernel assumes a significantly different dynamical model. Despite this, the parameters remain within a physically plausible range; as shown in [Figure 4.22](#), total pole movement remains under one degree.

The observed 65-year arc (1960–2025) represents a small fraction of the 688-year libration cycle. Sampling near the inflection point of the sine term creates an apparent linear trend in Right Ascension even without a secular rate ($\dot{\alpha}_0 = 0$).

State Update ($|\Delta\text{pos.}|$, $|\Delta\text{vel.}|$) — Pole Estimation Variants (IAU 2015) — Combined (FitPole vs IAU ref)

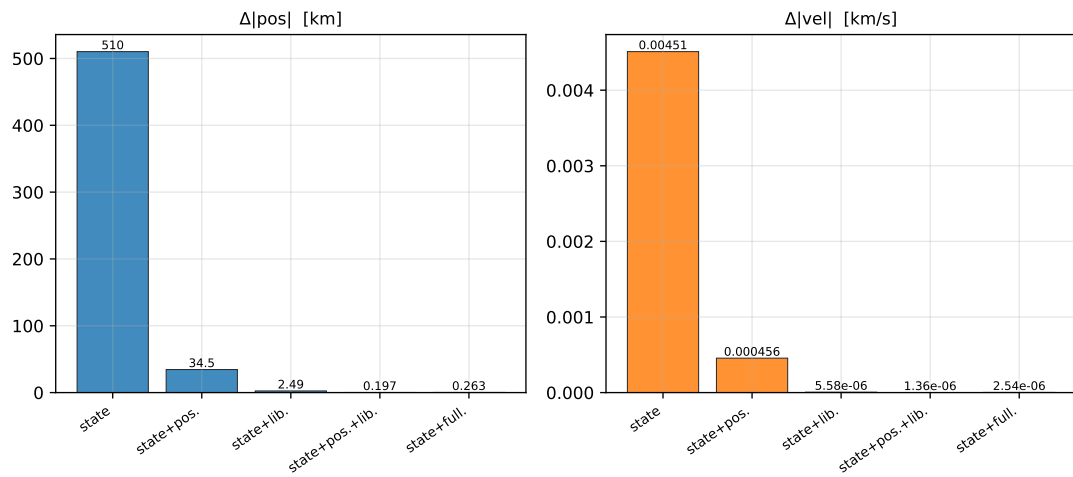


Figure 4.18: Estimated state updates for IAU pole cases.

RSW Position Update — Pole Estimation Variants (IAU 2015) — Combined (FitPole vs IAU ref)

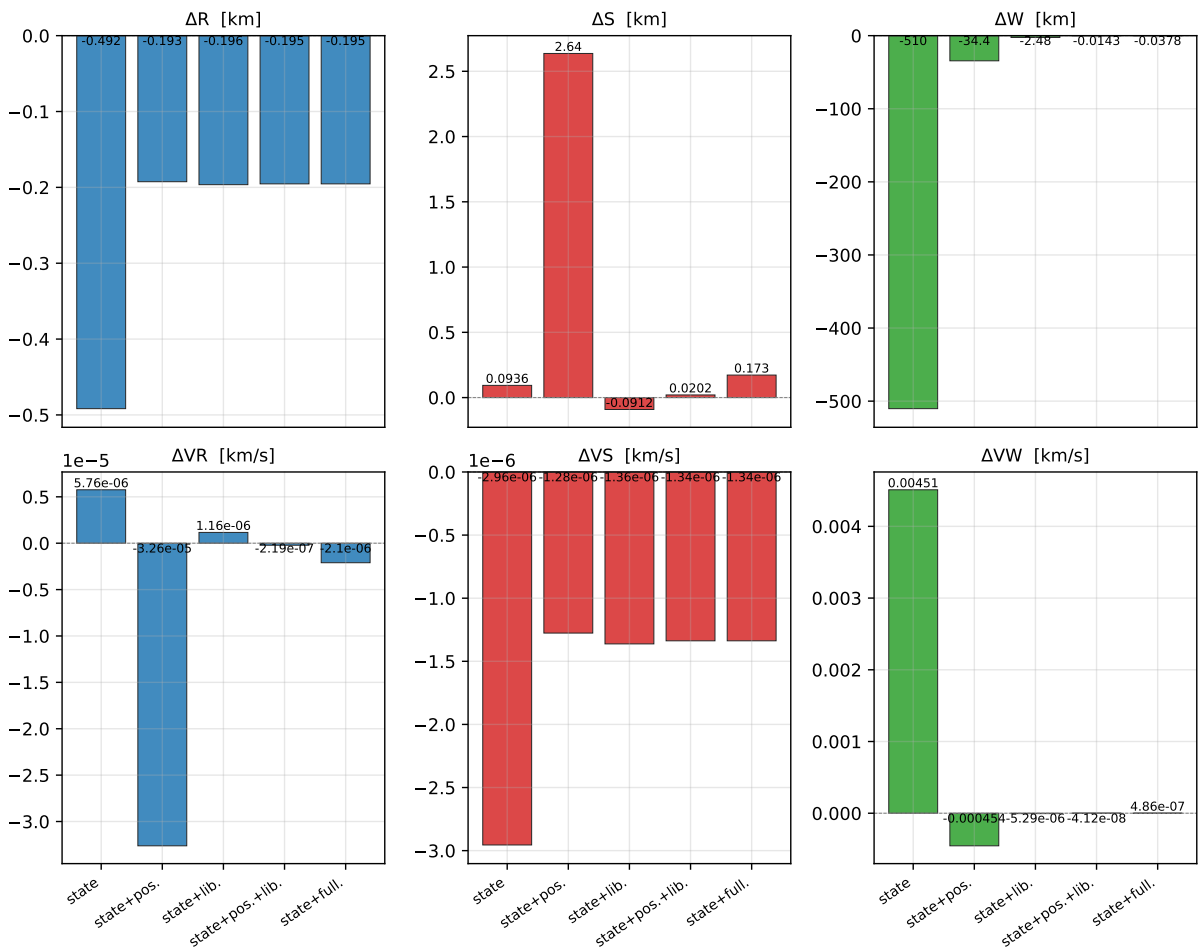


Figure 4.19: RSW decomposition of estimated state updates.



Pole Position Update (α_0, δ_0) — IAU 2015Figure 4.20: Estimated pole position (α_0, δ_0) updates.Pole Libration Update (α_1, δ_1) — IAU 2015

Figure 4.21: Estimated pole libration updates.




Figure 4.22: Neptune pole position evolution relative to initial IAU 2015 values.

4.5.5. Pole Model Conclusion

The state + pos. + lib. estimation achieves the best fit to the NEP097 kernel. While the inclusion of rotation rates (state + full) is numerically stable, its marginal improvement does not justify deviating from standard literature models. Thus, state + pos. + lib. is selected for subsequent real-data analysis. Final values are summarized in Table 4.6.

Table 4.6: Final estimated parameter values for best IAU 2015 variants.

Param.	Group	Unit	Initial Models		state + pos. + lib.		state + full	
			IAU 2015	Jac. 2009	Value	Δ	Value	Δ
X	Pos.	km	2.7441×10^5	2.7441×10^5	2.7441×10^5	-0.1531	2.7441×10^5	-0.2072
Y	Pos.	km	-2.4450×10^4	-2.4450×10^4	-2.4450×10^4	-0.0113	-2.4450×10^4	-0.1461
Z	Pos.	km	-2.2350×10^5	-2.2350×10^5	-2.2350×10^5	0.1234	-2.2350×10^5	0.0718
\dot{X}	Vel.	km s^{-1}	-1.8751	-1.8751	-1.8751	4.21×10^{-7}	-1.8751	-1.28×10^{-6}
\dot{Y}	Vel.	km s^{-1}	-3.4728	-3.4728	-3.4728	1.05×10^{-6}	-3.4728	1.50×10^{-6}
\dot{Z}	Vel.	km s^{-1}	-1.9223	-1.9223	-1.9223	7.50×10^{-7}	-1.9223	1.60×10^{-6}
α_0	Pole	deg	299.36	299.46	299.42	0.0551	299.41	0.0524
δ_0	Pole	deg	43.46	43.40	43.33	-0.1290	43.34	-0.1167
$\dot{\alpha}_0$	Rate	deg yr^{-1}	0.0	0.0	0.0	0.0	6.30×10^{-4}	6.30×10^{-4}
$\dot{\delta}_0$	Rate	deg yr^{-1}	0.0	0.0	0.0	0.0	1.30×10^{-5}	1.30×10^{-5}
α_1	Lib.	deg	0.700	0.635	-1.072	-1.772	-1.142	-1.842
δ_1	Lib.	deg	-0.510	-0.462	0.794	1.304	0.781	1.291

4.6. Extended Body Gravity Estimations of Neptune

The gravitational potential coefficients C_{20} and C_{40} of Neptune were also considered as estimated parameters. The total RMS differences relative to NEP097 are presented in Figure 4.23, with the RSW component decomposition shown in Figure 4.24.

Overall, the RMS values for the `state + SH.` estimation are comparable to the previously selected `state + pos. + lib.` variant. While this model shows a negligible improvement in the radial (R) direction, it performs slightly worse in the cross-track (W) direction. The RSW difference timeseries relative to NEP097 for both the `state + SH.` and the `state + pos. + lib.` cases are compared in Figure 4.25.

The estimation results and numerical properties are summarized in Table 4.7. As shown in the correlation heatmap in Figure 4.26, there is a high correlation between the two extended body gravity parameters. This coupling is expected as both parameters serve a similar dynamical function, with C_{40} providing higher-resolution, lower-amplitude corrections to the gravity field. Despite this correlation, the condition number remains within a stable range.

However, the physical validity of the solution is questionable. While the C_{20} update is approximately 5%, well within an expected range, the C_{40} parameter requires an update of over two orders of magnitude, or roughly 11,000%. Given this unreasonably high update for C_{40} and the fact that the RSW residuals are slightly larger than those of the `state + pos. + lib.` case, the extended body gravity parameters are not considered further in this study.

It is possible that better performance could be achieved if the parameters were estimated separately or if a strict a priori covariance were employed to penalize such large deviations. Nevertheless, based on the current results, they are not included in the final model configuration.

Table 4.7: Summary of spherical harmonics estimation for the `state+SH.` variant. $\kappa(\mathbf{C})$ denotes the condition number of the correlation matrix; $\Delta\bar{C}_{20}$ and $\Delta\bar{C}_{40}$ represent the estimated updates to the normalized gravitational coefficients.

Simulation	Estimated parameters	$\kappa(\mathbf{C})$	$\Delta\bar{C}_{20}$	$\Delta\bar{C}_{20}$ [%]	$\Delta\bar{C}_{40}$	$\Delta\bar{C}_{40}$ [%]
<code>state+SH.</code>	<code>state, $\bar{C}_{20}, \bar{C}_{40}$</code>	2.2412×10^{10}	-7.5643×10^{-5}	-4.96	-1.2335×10^{-3}	-11079.40




Figure 4.23: Total RMS difference with NEP097: Comparison between the `state + SH.` variant, the baseline `state` estimation, and the `state + pos. + lib.` model.

R/S/W RMS vs NEP097 — SH vs Benchmarks




Figure 4.24: RSW component RMS difference with NEP097: Comparison between state + SH., baseline state, and state + pos. + lib.




Figure 4.25: RSW difference timeseries with NEP097: Comparison of the state + SH. estimation against the selected state + pos. + lib. model.




Figure 4.26: Correlation heatmap for the state + SH. estimation variant.

4.7. Conclusion

This chapter presented a systematic evaluation of the dynamical model’s ability to reproduce the NEP097 kernel using simulated three-dimensional positional observations of Triton. The analysis progressed from a baseline propagation through increasingly complex parameter sets, culminating in the selection of a preferred estimation configuration for use with real astrometric data.

The initial propagation confirmed that neither the IAU 2015 nor the Jacobson 2009 pole model fully reproduces the pole configuration implicit in the NEP097 kernel. The most significant divergence occurs in the cross-track (W) direction, indicating a mismatch in orbital inclination that cannot be resolved by adjusting the initial state alone. The two pole models differ only marginally from each other, and this difference is below the threshold of detectability given current observational accuracy.

Estimating the initial state of Triton significantly reduces the R and S direction residuals but provides only a partial improvement in the W direction. The primary correction applied by the estimator is a cross-track shift of approximately 500 km, confirming that the orbital inclination mismatch is the dominant source of divergence from NEP097 and that it cannot be fully resolved through state estimation alone.

The inclusion of gravitational parameters does not improve the fit with respect to NEP097 and, in some cases, produces physically inadmissible results. The data constrains only the total system GM rather than the individual contributions of Neptune and Triton, and the estimator compensates for unmodelled perturbations through the gravitational parameters rather than improving the dynamical fit. Gravitational parameters are therefore excluded from subsequent analyses.

The inclusion of Neptune’s pole parameters substantially improves the W direction agreement with NEP097. Among the pole parameter configurations evaluated, the `state+pos.+lib.` estimation achieves the best overall fit. The addition of pole rotation rates in the `state+full.` variant provides no meaningful improvement over `state+pos.+lib.` at the accuracy level of the available observational data, and is therefore not pursued further. The Jacobson 2009 pole model offers no functional advantage over the IAU 2015 model for any of the estimation variants considered, and is similarly excluded from subsequent analyses.

The pole libration updates required to match the NEP097 kernel are substantially larger than the formal errors reported by [Jacobson \(2009\)](#), with a ratio of approximately 18 000. This likely reflects either overly optimistic formal errors in [Jacobson \(2009\)](#) or a fundamentally different dynamical model assumed in the construction of the NEP097 kernel. Despite this, the estimated pole parameters remain within a physically plausible range, with total pole movement remaining below one degree over the estimation arc.

Neptune’s extended body gravity coefficients C_{20} and C_{40} were also evaluated. While the C_{20} update is of a physically reasonable magnitude, the C_{40} update exceeds its initial value by more than two orders of magnitude. Given this result and the fact that the residuals are marginally larger than those of the `state+pos.+lib.` variant, estimation of extended body gravity parameters is not included in the final model configuration.

In summary, the `state+pos.+lib.` estimation with the IAU 2015 pole model and the NEP097-fitted initial values is selected as the preferred configuration for the subsequent real astrometric data analysis. The final estimated parameter values are provided in [Table 4.6](#) and serve as the `FitPole` initialisation for all subsequent chapters. The key limitation identified in this chapter is that the dynamical model, even when fitted to perfect three-dimensional positional data, does not fully converge to the NEP097 kernel in the W direction. This residual discrepancy is expected to be partially alleviated by the inclusion of high-precision datasets such as Voyager 2, Hubble, and Gaia in future work.

5

Pole Estimation with Astrometric Data

In this chapter, the pole estimation results using the full set of astrometric observations described in ?? are presented, including the manual bias correction applied to observation 689 nm0077 as discussed therein. The methodology of this chapter is described in detail in [subsection 3.4.2](#).

Two distinct groups of estimations are considered. The first group is initialised with the IAU 2015 pole model ([Archinal et al. \(2018\)](#)) and the initial state of Triton sourced from the NEP097 kernel. The second group, designated by the suffix (Fit.), uses the NEP097-fitted pole model selected as the best-performing estimation from [chapter 4](#). The initial state of Triton at epoch 2006, 10, 1 and the pole parameters for both groups are listed in [Table 5.2](#). Each estimation is propagated forwards to 2025, 1, 1 and backwards to 1963, 1, 1.

As discussed in [chapter 4](#), the correlation between the pole position and pole libration parameters is significant, making the use of a priori covariance necessary. In this chapter, an a priori covariance on the pole position and pole libration parameters is set to twice the change presented in [Table 4.6](#), such that the IAU 2015 pole model ([Archinal et al. \(2018\)](#)) lies within the uncertainty boundary.

It should be noted that a circular dependency exists between this chapter and the weight analysis presented in [chapter 6](#): a proper weight analysis should account for known dynamical modelling errors, which in turn depend on the estimation results presented here. The justification for the weighting scheme adopted throughout this chapter is therefore deferred to [chapter 6](#). All estimations here use the `scaled per file` weighting scheme described in [chapter 3](#). The weights are initialised by first fitting a state-only estimation to the observational data without weights, and then using the resulting residuals to compute the weights. As a consequence, the weights are identical across all estimations within each initialisation group, that is, all IAUPole estimations share the same weights and all FitPole estimations share the same weights.

As the number of pole estimation simulations is significant, the naming conventions used throughout this chapter are defined in [Table 5.1](#).

The final result of this chapter is the selected best estimation for constraining the motion of Triton and Neptune's pole, based on the minimisation of the observation residuals and the agreement of the estimation with the NEP097 kernel and formal errors.

5.1. Single Value Fit Metrics

The most direct way to compare the various pole estimation results, both with respect to the NEP097 kernel and with respect to each other, is through a single-value fit metric. The root mean square (RMS) difference is employed here as the primary metric. In this section, both the RMS between the estimation results and the NEP097 kernel (positional difference) and the RMS of the observation residuals (angular difference) are presented and discussed.

When examining the RMS difference with respect to NEP097 ([Figure 5.1](#)), the formal errors ([Figure 5.2](#)), and their ratio ([Figure 5.3](#)), the `state (Fit.)` estimation exhibits the smallest overall difference with respect to NEP097, while `state+pos.+lib. (Fit.)` exhibits the highest formal errors. The ratio of RMS difference to formal error is closest to unity for `state+lib. (Fit.)`, making it the preferred solution despite not achieving the lowest absolute

Table 5.1: Weight-scheme naming conventions used in figures for Pole Estimation Analysis.

Label	Description
state	Triton initial state only $(x, y, z, \dot{x}, \dot{y}, \dot{z})$, IAU 2015 pole model
state+pos.	Initial state + pole pos. (α_0, δ_0) , IAU 2015 pole model
state+lib.	Initial state + pole lib. (α_1, δ_1) , IAU 2015 pole model
state+pos.+lib.	Initial state + pole pos. + lib. $(\alpha_0, \delta_0, \alpha_1, \delta_1)$, IAU 2015 pole model
state (Fit.)	Triton initial state only $(x, y, z, \dot{x}, \dot{y}, \dot{z})$, NEP097-fitted pole model
state+pos. (Fit.)	Initial state + pole pos. (α_0, δ_0) , NEP097-fitted pole model
state+lib. (Fit.)	Initial state + pole lib. (α_1, δ_1) , NEP097-fitted pole model
state+pos.+lib. (Fit.)	Initial state + pole pos. + lib. $(\alpha_0, \delta_0, \alpha_1, \delta_1)$, NEP097-fitted pole model

Table 5.2: Initial parameter values for the IAU and FitPole rotation models. $\Delta = p_{\text{FitPole},0} - p_{\text{IAU},0}$ in the listed unit; $\Delta[\%] = \Delta / |p_{\text{IAU},0}| \times 100$.

Parameter	Group	Unit	IAU initial	FitPole initial	Δ	$\Delta[\%]$
X	Position	km	$2.744\,13 \times 10^5$	$2.744\,13 \times 10^5$	$-1.530\,51 \times 10^{-1}$	-0.00
Y	Position	km	$-2.444\,96 \times 10^4$	$-2.444\,96 \times 10^4$	$-1.126\,04 \times 10^{-2}$	-0.00
Z	Position	km	$-2.234\,96 \times 10^5$	$-2.234\,96 \times 10^5$	$1.234\,27 \times 10^{-1}$	0.00
VX	Velocity	km/s	-1.875 07	-1.875 07	$4.214\,51 \times 10^{-7}$	0.00
VY	Velocity	km/s	-3.472 85	-3.472 85	$1.048\,25 \times 10^{-6}$	0.00
VZ	Velocity	km/s	-1.922 29	-1.922 28	$7.499\,13 \times 10^{-7}$	0.00
α_0	Pole Position	deg	$2.993\,60 \times 10^2$	$2.994\,15 \times 10^2$	$5.505\,81 \times 10^{-2}$	0.02
δ_0	Pole Position	deg	$4.346\,00 \times 10^1$	$4.333\,10 \times 10^1$	$-1.290\,35 \times 10^{-1}$	-0.30
α_1	Pole Librations	deg	$7.000\,00 \times 10^{-1}$	-1.071 80	-1.771 80	-253.11
δ_1	Pole Librations	deg	$-5.100\,00 \times 10^{-1}$	$7.936\,10 \times 10^{-1}$	1.303 61	255.61

RMS difference. The unexpectedly strong performance of `state (Fit.)` in terms of absolute RMS difference is investigated further in the next section.

Considering the two initialisations, `IAUPole` and `FitPole`, the `state+lib.` estimations consistently exhibit the closest ratio to unity and generally perform well in terms of RMS difference. The `FitPole` initialisation performs consistently better than the `IAUPole` initialisation across all metrics.

While the 3-parameter pole model performed optimally in the simulation environment presented in [chapter 4](#), the results presented thus far suggest that the 2-parameter `state+lib.` estimation achieves a better fit to the real astrometric data. This implies that even with a priori constraints, the simultaneous estimation of all pole parameters does not yield the expected improvement when evaluated against the NEP097 benchmark.




Figure 5.1: RMS of the difference relative to NEP097 for pole estimations using astrometric data.




Figure 5.2: RMS of formal errors for the pole estimation variants.




Figure 5.3: Ratio of the RMS difference relative to NEP097 to the RMS of the formal errors.

Three goodness-of-fit metrics relative to the observational data are presented in [Figure 5.4](#): the RMS of the residuals, the weighted RMS, and the cost function. The initial values for each metric are consistent across estimations within the `IAUPole` and `FitPole` groups respectively, as the weighting schemes and initial conditions are identical within each group.

The initial values for the `FitPole` estimations are notably lower than those for the `IAUPole` estimations. This is expected, as the `FitPole` initialisation is derived from the `state+pos.+lib.` variant of the simulated observation

analysis (chapter 4), which was already fitted to the NEP097 kernel.

Across all three goodness-of-fit metrics, all estimation variants perform similarly, with differences on the order of fractions of a milliarcsecond for the unweighted RMS and only a few milliarcseconds for the weighted RMS. The observation residuals ($O-C$) for the `state+lib. (Fit.)` estimation are shown in Figure 5.5. The residuals remain flat and exhibit no discernible trends in either right ascension or declination across the full observed arc.

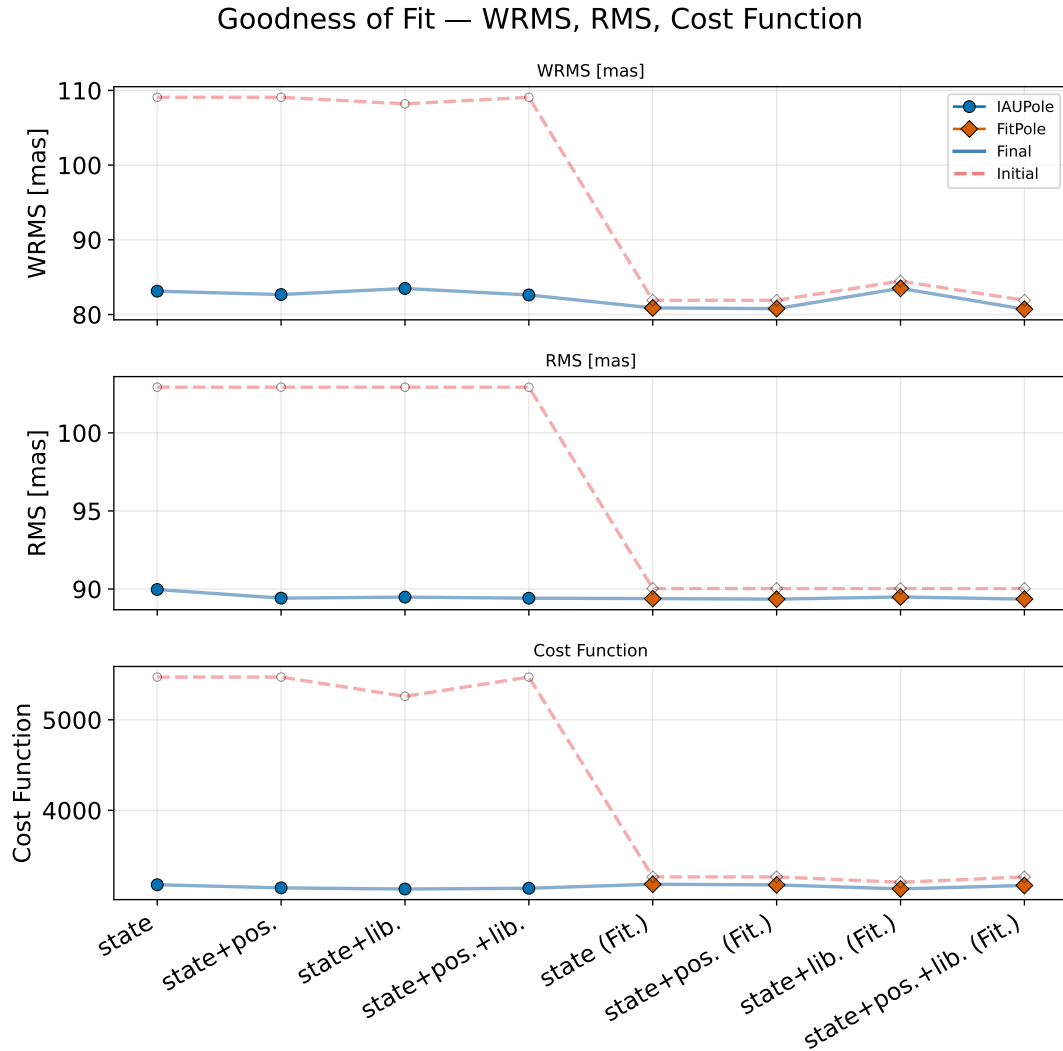


Figure 5.4: Goodness-of-fit metrics derived from observation residuals for the pole estimation variants.




Figure 5.5: Observation residuals categorised by data source for the *state+lib.* (Fit.) estimation.

In conclusion, no meaningful difference between the estimation variants can be discerned from the observational metrics alone, as all variants perform similarly across all three goodness-of-fit measures. Therefore, in the absence of an external benchmark such as NEP097, identifying a superior estimation would be virtually impossible without relying heavily on formal errors, an approach that has been shown to be unreliable in [chapter 6](#). The availability of the high-precision NEP097 benchmark is therefore essential to the conclusions drawn in this study.

5.2. RSW Statistics with Respect to NEP097

A more detailed comparison can be made by examining the per-direction RSW statistics of the difference with respect to NEP097 and the formal errors, which are summarised in [Figure 5.6](#).

In the R direction, both the *IAUPole* and *FitPole* initialisations produce similar results in terms of both RMS difference and formal errors across all estimation variants. The only consistent trend is that the *state+lib.* estimation yields the most coherent results: the difference with respect to NEP097 is the lowest, the formal errors are well-calibrated, and the ratio of difference to formal error is approximately 0.8. A ratio slightly below unity is considered desirable, as it indicates that the formal errors are marginally more conservative than the actual difference, which is preferable to the opposite case where formal errors underestimate the true error.

In the S direction, the general trend is that the inclusion of pole parameters reduces the ratio of difference to formal errors for both initialisations. The lowest difference with respect to NEP097 is observed for the *state+lib.* estimation, where the *IAUPole* initialised variant achieves an RMS difference approximately 30 km lower than the *FitPole* initialised variant, with ratios of 1.6 and 1.87 respectively. The *state+pos.+lib.* estimations exhibit a lower ratio of 1.1 and 0.9 for *IAUPole* and *FitPole* respectively, however their RMS difference with respect to NEP097 is approximately double that of the *state+lib.* estimations, making them less favourable overall.

In the W direction the behaviour differs significantly between the two initialisations. For the *IAUPole* initialisation, the inclusion of pole position or pole libration separately reduces the difference with respect to NEP097 significantly compared to the state-only estimation. Notably, *state+pos.* achieves an RMS difference approximately 40 km lower than *state+lib.*, while *state+pos.+lib.* produces a higher RMS difference than even the state-only estimation. For the *FitPole* initialisation, the inclusion of any pole parameters, whether separately or in combination, increases the difference with respect to NEP097. The lowest difference is observed for the *state* (Fit.) estimation, which is most likely a consequence of the initial state being unable to significantly alter the across-track position of Triton,

combined with the fact that the `FitPole` initialisation already fits the *W* direction well. This estimation also achieves the lowest ratio of difference to formal error, at approximately 0.44. For the preferred solution `state+lib. (Fit.)`, the RMS difference with respect to NEP097 in the *W* direction is substantial at approximately 375 km, with a ratio of difference to formal error of 1.68, which is considered acceptable.

In conclusion, the estimation of the initial state together with pole libration parameters improves the *R* and *S* direction agreement with NEP097 compared to the state-only estimation for both initialisations. The *W* direction is more difficult to assess, as improvements in the *R* and *S* directions appear to come at the cost of a larger *W* direction difference. Nevertheless, the ratio of difference to formal errors remains most consistent for the `state+lib.` estimations. Further analysis of the parameter updates and RSW time series differences with respect to NEP097 are presented in the following sections to determine whether the `FitPole` initialised `state+lib.` estimation is preferable.

The lower RMS difference with NEP097 for the `state (Fit.)` estimation is interesting to investigate. As discussed in [chapter 4](#) the initial state has little effect on the *W* direction difference with NEP097. In contrast the pole model has large effect on the *W* direction difference, as the `FitPole` model initialization has a significantly small difference with NEP097, therefore the `state (Fit.)` estimation has such low RMS as it does not have the ability to affect the *W* direction.




Figure 5.6: RSW statistics of the difference with respect to NEP097, formal errors, and their ratio for the pole estimation variants using astrometric data.

5.3. Correlation Plots

The correlation matrices for the estimations `state`, `state+lib.`, and `state+pos.+lib.` are shown in [Figure 5.9](#), [Figure 5.8](#), and [Figure 5.7](#) respectively.

All off-diagonal correlations are in the low-to-mid range, with some expected coupling between positions and velocities arising from the system dynamics and the nature of the observational data. The initial state position and velocity correlations are similar across all estimations, with only slight coupling between certain position and velocity components, as expected from the dynamics of the Triton system where the radial and along-track directions are known to influence each other significantly.

The condition numbers of all correlation matrices are similar at approximately 10^9 , even for `state+pos.+lib.`, which

exhibits a large correlation between the declination of the pole position δ_0 and the pole libration rate δ_1 . While condition numbers of this magnitude may appear large, they remain within a range that indicates dynamical stability of the estimation.

The high correlation between δ_0 and δ_1 is a consequence of the construction of the pole model, as described in [chapter 3](#) and [section 5.6](#), and is amplified by the specific time period considered. This correlation implies that the estimation cannot uniquely resolve the individual contributions of these two parameters. While this ambiguity could be a concern for long-term propagations of Triton's orbit, it is not expected to affect the solution significantly over shorter periods of approximately 100 to 200 years, as the combined effect of the two parameters on the trajectory remains consistent within that timeframe.

In conclusion, while the coupling between the pole position declination and the pole libration rate is significant, it is not severe enough to drive the condition number to a value that would render the estimation untrustworthy. Nevertheless, the results presented in the preceding sections indicate that simultaneously estimating both parameters leads to unsatisfactory outcomes. This may be attributable to incorrectly specified a priori covariances, insufficient observational data to jointly constrain these parameters, or both. It is worth noting that [Jacobson \(2009\)](#) enforces a dynamically derived pole model, which may help reduce the coupling between pole position and libration parameters and could be a worthwhile avenue to explore in future work.




Figure 5.7: Correlations between estimated parameters for FitPole estimation for pole position and pole libration rate




Figure 5.8: Correlations between estimated parameters for FitPole estimation for pole libration rate




Figure 5.9: Correlations between estimated parameters for FitPole estimation for initial state only

5.4. RSW Timeseries Analysis

The RSW difference time series with respect to NEP097 for the state-only and state+lib. estimations are shown in Figure 5.10 and Figure 5.11 respectively, alongside their formal errors in Figure 5.12. Several trends are apparent across both estimation variants.

In the radial (R) and cross-track (W) directions, both the IAU and fitted variants exhibit oscillatory behaviour at the orbital period, while the along-track (S) direction captures the secular drift of the solution relative to NEP097.

The state-only estimation reveals a subtle but important effect. As demonstrated in chapter 4 and the preceding sections, the initial state alone is incapable of constraining the W direction difference with respect to NEP097, meaning the cross-track component retains a large error of approximately 2000 km for the IAUPole variant throughout the full time range. The state (Fit.) variant, having been initialised with pole parameters adjusted to match NEP097, shows negligible W direction residuals not because the estimation has improved the fit in that direction, but precisely because it lacks the sensitivity to deviate from the initialisation. This suggests that the observational data drives the solution away from the NEP097 kernel in directions the estimation can resolve, while the W direction remains anchored by the fitted initial conditions rather than by genuine observational constraint.

For the state+lib. estimation, the IAUPole and FitPole variants perform comparably overall, with the FitPole variant achieving a marginally lower total RMS (372 km versus 350 km). This improvement is primarily driven by greater consistency in the W direction, which is likely attributable to the FitPole initial pole parameters being constrained to match the NEP097 kernel rather than a genuine improvement in the dynamical fit. This comes at the cost of a slightly larger S direction offset, illustrating the trade-off inherent to the fitted initialisation: some along-track agreement with NEP097 is sacrificed in exchange for better cross-track consistency.

Compared to the state-only variants, the state+lib. estimations exhibit a similar magnitude and period in the R direction, a substantially lower S direction difference, and a significant improvement in the W direction relative to the IAUPole initialisation. The same improvement is not observed for the FitPole initialisation, whose propagation behaviour was discussed in chapter 4 and presented in Figure 4.17. Unlike the simulated observation estimations, the real data estimations move away from the NEP097 kernel, suggesting that the real observational data is not sufficient to reproduce the NEP097 solution.

Considering the formal errors shown in Figure 5.12, the difference with respect to NEP097 and the formal errors are in reasonable agreement in the S direction. In the W direction a mismatch of approximately a factor of two is present, and in the R direction the mismatch is approximately a factor of five. The R direction mismatch is particularly

concerning, as it suggests the presence of systematic biases in the observational data that skew the estimated semi-major axis away from the true value. Correcting such a bias would be expected to partially alleviate the S direction offset as well, given the dynamical coupling between the two directions. In the W direction, the simulated observation results in [chapter 4](#) demonstrate that the dynamical model is capable of fitting to the NEP097 kernel, and the weight analysis in [chapter 6](#) supports the conclusion that the weighting strategy is reasonably well chosen given the available data. The remaining possibilities are therefore either the presence of systematic biases in the observational dataset, or that the quantity and precision of the available data are insufficient to constrain the solution to the accuracy of the NEP097 kernel. Without access to the full documentation of the NEP097 solution, including its data sources, weighting strategy, and dynamical model assumptions, it is difficult to assess precisely where the differences originate.

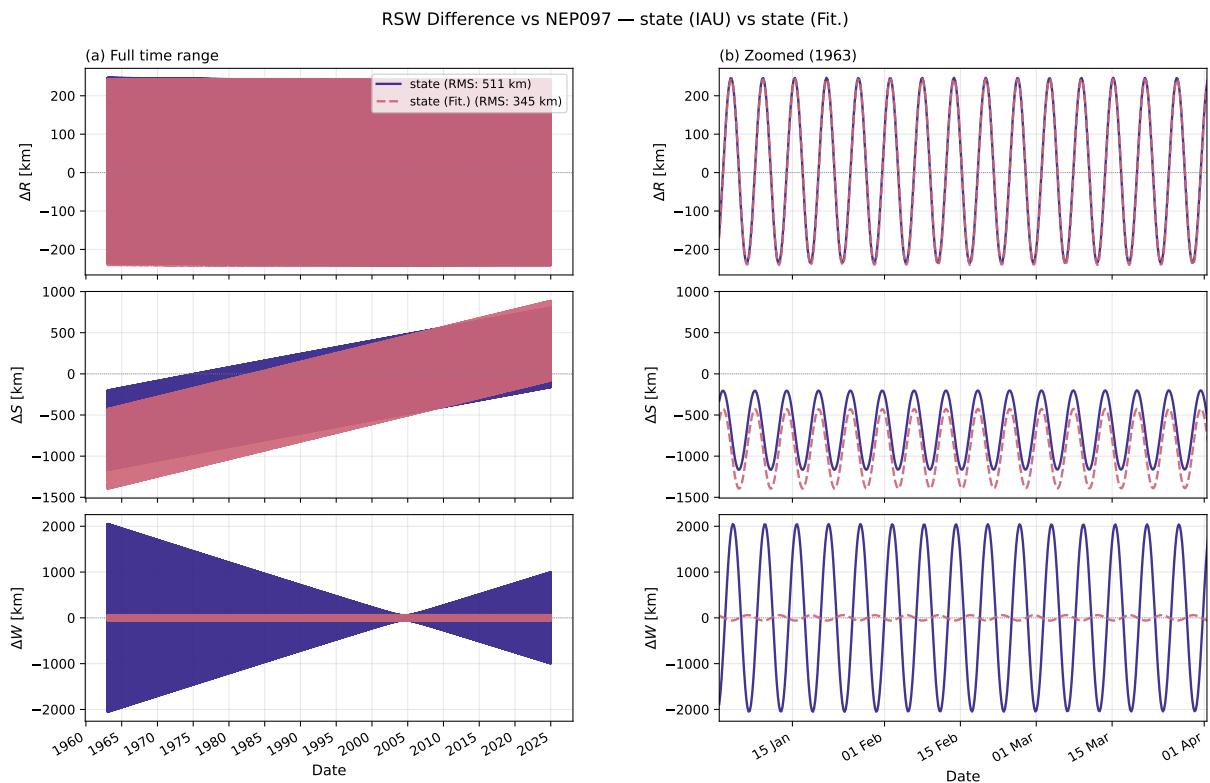


Figure 5.10: RSW Difference with NEP097 for estimations state and state (Fit.)




Figure 5.11: RSW Difference with NEP097 for estimations state + lib. and state + lib. (Fit.)




Figure 5.12: RSW Formal Errors for estimations state + lib. and state + lib. (Fit.)

5.5. Parameter Updates

Having established that the two `state+lib.` variants produce comparable agreement with NEP097 and similarly low correlations, this section examines the underlying parameter updates to determine which solution is physically more plausible.

The total initial state update of Triton in position and velocity for the two `state+lib.` estimations and the benchmark `state` estimation are presented in [Figure 5.14](#), with the per-component RSW decomposition shown in [Figure 5.13](#). Note that both the `IAUPole` and `FitPole` updates are expressed relative to the `IAUPole` initial propagation, meaning it is not the raw estimator corrections that are compared here but rather the total departure from the common `IAUPole` reference. The actual estimator correction for `state+lib. (Fit.)` is given in [Table 5.3](#).

In terms of total initial state update, both `state+lib.` and `state+lib. (Fit.)` produce larger corrections than their `state` counterparts. This is unexpected, as the pole libration parameters were anticipated to absorb part of the signal that the initial state was capturing in the `state` estimations, thereby reducing the required state correction.

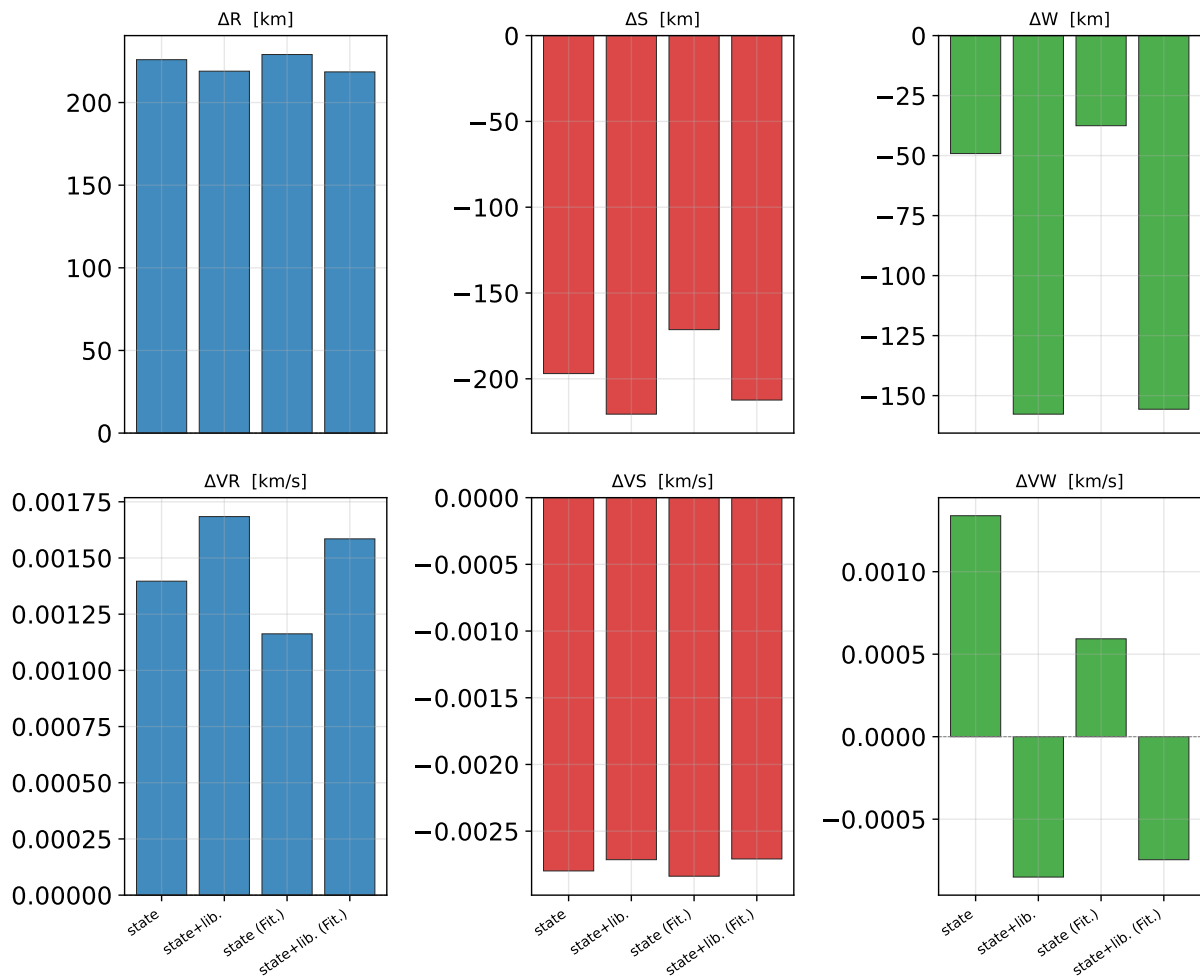


Figure 5.13: RSW-decomposed initial position and velocity update of Triton for `state+lib.` under both `IAUPole` and `FitPole` initialisations, and the benchmark `state` estimation.




Figure 5.14: Total initial position and velocity update of Triton for `state+lib.` under both `IAUPole` and `FitPole` initialisations, and the benchmark `state` estimation.

In the RSW decomposition, the `state` and `state (Fit.)` estimations update primarily the radial and along-track directions by approximately 200 km each, while the cross-track direction receives comparatively little correction. This stands in stark contrast to the findings of [chapter 4](#), specifically [Figure 4.4](#), where the update is dominated by the W direction with a correction of approximately 500 km, while the R and S directions receive corrections of at most 1 km. This discrepancy is also reflected in the RSW differences with respect to NEP097 shown in [Figure 5.10](#): despite large R and S initial state corrections, the radial and along-track residuals with respect to NEP097 remain considerably larger than those obtained with simulated observations. This is most likely a consequence of the observational dataset employed here, which lacks both the coverage and the three-dimensional positional information of the homogeneously spaced simulated data. Furthermore, the absence of the highly precise Voyager 2 data included in [Jacobson \(2009\)](#), combined with the influence of older micrometric observations, likely pulls the solution in a different direction compared to the NEP097 kernel. An additional contributing factor may be an inconsistency between the weighting scheme employed here and that used in the construction of NEP097, though this is investigated further in [chapter 6](#).

In the `state+lib.` estimations, the initial state corrections are significant in all RSW directions for both position and velocity. The R and S direction position updates are comparable in magnitude to those of the `state` estimations, while the W direction update is nearly three times larger. This result is counter-intuitive: the W direction is dynamically decoupled from the R and S directions, and as demonstrated in [chapter 4](#), initial state corrections in this direction have little effect on the W direction difference with respect to NEP097. It would therefore be expected that the pole libration parameters absorb the W direction signal, leaving the state correction in that direction near zero. The fact that this does not occur suggests that the observational data is driving the state update in an unexpected manner, possibly due to the limited coverage or systematic biases discussed above.

The pole libration parameter updates are summarised in [Table 5.3](#). The libration amplitude corrections $\Delta\alpha_1$ and $\Delta\delta_1$ are considerable. For `state+lib.`, the estimator updates α_1 by $+5.04^\circ$ from the `IAU0` reference, while `state+lib. (Fit.)` updates it by $+4.59^\circ$ from the same reference, both representing substantial departures from the IAU 2015 values. Expressed relative to the `FitPole` reference `FP0`, the shifts are even larger, reaching $+6.81^\circ$ and $+6.36^\circ$ respectively. As established in [chapter 4](#), libration updates of this magnitude are not unexpected given the large ratio between the estimated correction and the formal errors of [Jacobson \(2009\)](#), which was found to reach values on the order of approximately 18 000 even in the simulated observation case. This mismatch likely reflects overly optimistic formal errors in [Jacobson \(2009\)](#), or a significantly different dynamical model assumed in the construction of NEP097, rather than an unphysical estimation result.

In summary, both `state+lib.` variants make unexpectedly large adjustments to both the initial state and the pole libration parameters, contrary to the expectation that libration would be the primary degree of freedom absorbed by the estimator. The magnitude of the libration corrections, which substantially exceed the [Jacobson \(2009\)](#) formal errors for both initialisations, warrants caution in interpreting the results. The solution may be sensitive to the obser-

Table 5.3: Key parameter estimates for selected simulations. IAU_0 and FP_0 are the initial values for the IAU 2015 and NEP097-fitted pole models respectively. $\Delta_{\text{IAU}} = \text{Est.} - \text{IAU}_0$; $\Delta_{\text{FP}} = \text{Est.} - \text{FP}_0$.

Parameter	Unit	IAU_0	FP_0	state+lib.		state+lib. (Fit.)	
				Δ_{IAU}	Δ_{FP}	Δ_{IAU}	Δ_{FP}
X	km	2.74413×10^5	2.74413×10^5	3.37461×10^2	3.37614×10^2	3.32655×10^2	3.32808×10^2
Y	km	-2.44496×10^4	-2.44496×10^4	6.35588×10^1	6.35701×10^1	5.83314×10^1	5.83427×10^1
Z	km	-2.23496×10^5	-2.23496×10^5	5.98045×10^1	5.96811×10^1	5.51603×10^1	5.50368×10^1
V_X	km/s	-1.87507	-1.87507	2.86029×10^{-3}	2.85986×10^{-3}	2.73204×10^{-3}	2.73162×10^{-3}
V_Y	km/s	-3.47285	-3.47285	1.51374×10^{-3}	1.51270×10^{-3}	1.58012×10^{-3}	1.57907×10^{-3}
V_Z	km/s	-1.92229	-1.92228	6.73022×10^{-4}	6.72272×10^{-4}	6.65649×10^{-4}	6.64899×10^{-4}
α_1	deg	7.00000×10^{-1}	-1.07180	5.04165	6.81346	4.58996	6.36176
δ_1	deg	-5.10000×10^{-1}	7.93610×10^{-1}	1.06138	-2.42234×10^{-1}	1.19008	-1.13535×10^{-1}

vational dataset and data weighting strategy rather than reflecting a robust improvement over the prior pole model. The following section investigates the pole orientation trajectory implied by these parameter updates in more detail.

5.6. Pole Movement

The Neptune pole trajectory implied by the estimated libration parameters for `state + lib.` and `state + lib. (Fit.)` is shown in [Figure 5.15](#), alongside the `IAUPole` and `FitPole` nominal trajectories for reference. The pole orientation is governed by:

$$\alpha = \alpha_0 + \dot{\alpha}_0 \cdot t + \alpha_1 \cdot \sin(N) \quad (5.1)$$

$$\delta = \delta_0 + \dot{\delta}_0 \cdot t + \delta_1 \cdot \cos(N) \quad (5.2)$$

where N is the longitude of Neptune's ascending node, advancing at 52.316° per century, giving a libration period of approximately 688 years.

In the declination panel, both estimations appear primarily as a constant offset relative to the IAU nominal trajectory rather than exhibiting a visually distinct oscillatory trend. This is physically expected: over the 62-year observation window (1963–2025), the libration phase N sweeps only approximately 32° , from $\sim 338^\circ$ to $\sim 11^\circ$, over which $\cos(N)$ varies from ~ 0.930 to ~ 0.982 , a change of less than 5%. The term $\delta_1 \cdot \cos(N)$ therefore appears nearly constant throughout the observation arc, rendering δ_1 and δ_0 practically indistinguishable from one another over this window. This parameter degeneracy explains the correlation between the libration amplitude and the pole position observed in the previous section, and means the estimator effectively uses δ_1 to absorb any systematic offset in declination. Both estimated trajectories lie between the `IAUPole` and `FitPole` nominal values in declination, and notably closer to `FitPole` than to `IAUPole`. This is a physically reassuring result, as it indicates the estimations are not moving away from the NEP097-consistent `FitPole` model but rather converging toward an intermediate solution.

The right ascension panel tells a different story. Here, $\sin(N)$ varies from ~ -0.37 to $\sim +0.19$ over the observation arc, a much more significant change that makes α_1 genuinely resolvable. Both estimations produce a positive gradient in α that is steeper than either nominal model, suggesting the estimated α_1 is larger in magnitude than both the IAU 2015 and `FitPole` values, consistent with the large $\Delta\alpha_1$ updates reported in [Table 5.3](#). Importantly, the sign of the gradient is consistent with the `IAUPole` model for both estimations, whereas `FitPole` exhibits the opposite trend. Both `state + lib.` and `state + lib. (Fit.)` converge to a similar right ascension trajectory despite their different initialisations, which further supports the robustness of the estimated α_1 correction.




Figure 5.15: Neptune pole trajectory for state + lib. and state + lib. (Fit.) compared to the IAU Pole and FitPole nominal trajectories.

5.7. Final Solution

The observation residuals for the final estimation `state + lib. (Fit.)` are presented in [Table 5.4](#). The residuals serve primarily as a confirmation of acceptable absolute fit quality across all datasets rather than as a discriminator between estimations; as established in the goodness-of-fit analysis, the total observational RMS is virtually indistinguishable across all estimation variants. The residuals are generally consistent with the expected astrometric precision of each dataset, with the higher-precision relative astrometry datasets from the U.S. Naval Observatory achieving sub-0.1'' RMS in both coordinates. These residuals can be compared to the NEP097 kernel residuals for the same dataset in ??.

The estimated parameters and their formal errors are presented in [Table 5.5](#). The formal errors on the initial state are on the order of 60 to 95 km in position and $\sim 10^{-3}$ km/s in velocity, evaluated at the estimation epoch `DateTime(2006, 10, 1)`, while the differences with the NEP097 kernel are on the order of 300 km in the X direction and approximately 60 km in the Y and Z directions. Although these values are physically plausible, the positional updates in combination with their formal errors do not encompass the NEP097 solution. This is most likely caused by the initial state absorbing residual signal from parameters that are not estimated, such as Neptune's gravitational zonal harmonics.

For the pole libration parameters, the estimated corrections relative to the IAU 2015 initial values are $\Delta\alpha_1 = +4.59^\circ$ and $\Delta\delta_1 = +1.19^\circ$, consistent in sign with the findings from [chapter 4](#) reported in [Table 4.6](#), though larger in magnitude. The formal error on α_1 is $\pm 1.71^\circ$, placing the correction at approximately 2.7σ ; this is suggestive but not strongly statistically significant. By contrast, the formal error on δ_1 is $\pm 0.19^\circ$, placing the correction at approximately 6.3σ and making it statistically significant. This discrepancy is noteworthy: geometrically, δ_1 enters the pole model as $\delta_1 \cdot \cos(N)$, where $\cos(N)$ varies by less than 5% over the observation arc, rendering δ_1 nearly degenerate with the pole position offset δ_0 . One would therefore expect δ_1 to be poorly determined, yet the formal error suggests the opposite. Conversely, α_1 enters through $\sin(N)$, which varies more substantially over the arc and should in principle

Table 5.4: RMS observed minus computed (O–C) residuals per observation dataset using post-estimation residuals (state + lib. (Fit.)). MPC Code is the three-digit Minor Planet Center observatory code. N_{obs} is the number of astrometric observations in the dataset. Obs. Type indicates relative (Rel.) or absolute (Abs.) astrometry. All RMS values in arcseconds [″].

Observatory	NSDC Listing	MPC Code	N_{obs}	Obs. Type	RMS O–C Final [″]	
					RA	Dec
Holosivskiy district-Kyiv	nm0083	83	1	—	0.818	0.205
Abastuman	nm0017	119	54	—	0.413	0.417
Majdanak	nm0016	188	5	—	0.143	0.328
Majdanak	nm0083	188	9	—	0.571	0.527
Yunnan Observatory	nm0084	286	1135	—	0.033	0.033
Yunnan Observatory	nm0091	286	1918	—	0.028	0.034
Peking Observatory, Xinglong Station	nm0019	327	139	—	0.105	0.119
Peking Observatory, Xinglong Station	nm0082	327	34	—	0.060	0.066
Sheshan, formerly Zo-Se	nm0015	337	943	—	0.077	0.053
Sheshan, formerly Zo-Se	nm0019	337	956	—	0.058	0.058
Sheshan, formerly Zo-Se	nm0085	337	2299	—	0.050	0.050
Sheshan, formerly Zo-Se	nm0086	337	604	—	0.071	0.049
Sheshan, formerly Zo-Se	nm0088	337	1006	—	0.042	0.039
Table Mountain Observatory, Wrightwood	nm0079	673	200	—	0.059	0.053
U.S. Naval Observatory, Flagstaff	nm0007	689	124	—	0.116	0.116
U.S. Naval Observatory, Flagstaff	nm0008	689	56	Rel.	0.019	0.067
U.S. Naval Observatory, Flagstaff	nm0009	689	28	Rel.	0.067	0.092
U.S. Naval Observatory, Flagstaff	nm0010	689	114	Rel.	0.034	0.073
U.S. Naval Observatory, Flagstaff	nm0077	689	874	—	0.130	0.150
U.S. Naval Observatory, Flagstaff	nm0078	689	116	—	0.144	0.243
Optec Observatory	nm0081	755	755	—	0.114	0.051
Observatorio do Pico dos Dias, Itajuba	nm0003	874	423	Rel.	0.103	0.152
Observatorio do Pico dos Dias, Itajuba	nm0004	874	759	Rel.	0.191	0.189
Observatorio do Pico dos Dias, Itajuba	nm0013	874	65	—	0.159	0.098

be better resolved, yet it carries a much larger formal error. The cause of this inconsistency is not fully understood and may be related to the structure of the a priori covariance applied to the libration parameters.

Table 5.5: Estimated parameters for simulation state + lib. (Fit.). IAU_0 and FP_0 are the initial values for the IAU 2015 and NEP097-fitted pole models, respectively. $\pm 1\sigma$ is the formal error from the post-fit covariance matrix (square root of the diagonal element).

$$\Delta_{\text{IAU}} = \text{Final} - \text{IAU}_0; \Delta_{\text{FP}} = \text{Final} - \text{FP}_0.$$

Parameter	Group	Unit	IAU_0	FP_0	Final	$\pm 1\sigma$	Δ_{IAU}	Δ_{FP}
X	Position	km	2.74413×10^5	2.74413×10^5	2.74746×10^5	6.11948×10^1	3.32655×10^2	3.32808×10^2
Y	Position	km	-2.44496×10^4	-2.44496×10^4	-2.43913×10^4	9.54534×10^1	5.83314×10^1	5.83427×10^1
Z	Position	km	-2.23496×10^5	-2.23496×10^5	-2.23441×10^5	7.73294×10^1	5.51603×10^1	5.50368×10^1
VX	Velocity	km/s	-1.87507	-1.87507	-1.87234	1.08690×10^{-3}	2.73204×10^{-3}	2.73162×10^{-3}
VY	Velocity	km/s	-3.47285	-3.47285	-3.47127	9.47272×10^{-4}	1.58012×10^{-3}	1.57907×10^{-3}
VZ	Velocity	km/s	-1.92229	-1.92228	-1.92162	8.62159×10^{-4}	6.65649×10^{-4}	6.64899×10^{-4}
α_1	Pole Librations	deg	7.00000×10^{-1}	-1.07180	5.28996	1.70523	4.58996	6.36176
δ_1	Pole Librations	deg	-5.10000×10^{-1}	7.93610×10^{-1}	6.80075×10^{-1}	1.87152×10^{-1}	1.19008	-1.13535×10^{-1}

The RSW difference with NEP097 and the propagated formal errors for the final estimation are shown in Figure 5.16. In the radial direction, the formal error is significantly smaller than the difference with NEP097, indicating a systematic mismatch that the formal errors do not capture. In the along-track direction, the formal error is comparable to the difference near the estimation epoch but the difference grows to exceed the formal error boundary toward the edges of the arc, particularly toward 2025, while in the backward direction toward 1963 the two remain comparable. This asymmetry suggests overly optimistic weights assigned to the newer, more precise datasets, contrasted with better-calibrated weights for the older, less precise data. Modern astrometric data is extremely precise, meaning systematic biases become non-negligible; given that this analysis has addressed only a single such bias through manual subtraction, which is consequently not reflected in the formal error analysis, this represents a point worthy of further investigation. In the cross-track direction, the RSW difference remains consistent with the formal error envelope throughout most of the arc but increases substantially prior to 1985, which coincides with a period of sparser data coverage and markedly different data quality compared to later years.




Figure 5.16: RSW difference with NEP097 and propagated formal errors (1σ) for the final estimation state + lib. (Fit.).

In summary, state + lib. (Fit.) is selected as the final estimation for constraining the motion of Triton and the orientation of Neptune's pole. The observation residuals confirm an acceptable fit across all datasets. The libration amplitude corrections are substantial: the δ_1 correction of $+1.19^\circ$ is statistically significant at approximately

6.3σ , while the α_1 correction of $+4.59^\circ$, despite its larger magnitude, is only moderately significant at 2.7σ due to its correspondingly larger formal error. The cause of this asymmetry warrants further investigation. The RSW formal errors are broadly consistent with the NEP097 differences in the along-track and cross-track directions, with the exception of the period prior to 1985 in the cross-track direction, while the radial direction reflects the fundamental limitations of the dataset.

6

Weighting Strategy Comparison Across Estimation Analyses

This chapter presents the analysis of the observational weighting strategy applied in the orbit and pole estimation of Triton. The full methodology of the weighting strategy is presented in [subsection 3.4.3](#). The goal is to assess which weighting scheme produces the most accurate and reliable orbit determination solution, where accuracy is measured with respect to the NEP097 benchmark kernel and reliability is measured by the consistency between the formal errors and the actual solution error.

Assessing whether the weights accurately represent the true observation uncertainty is inherently difficult. If the exact error of each observation were known, it could be corrected for directly, leaving no residual error. Since this is not the case, two broad categories of metrics are used in this research. The first category measures the agreement between the produced solution and the NEP097 benchmark kernel, which is significantly more precise than the solutions produced here and therefore serves as a reliable reference. The second category measures the agreement between the formal errors of the solution and the actual solution error. The actual solution error is not directly known, but it can be approximated by the difference between the solution and the NEP097 benchmark, under the assumption that NEP097 is sufficiently more accurate than the solutions investigated here. Additionally, the estimations are compared based on the observational data residuals. However, all weighting schemes perform indistinguishably with respect to the observational residuals, with RMS differences of at most a couple of mas, from which no concrete evidence of a better performing estimation can be derived.

Three metrics are defined for each category, as derived in detail in [chapter 3](#). For the comparison with NEP097, the metrics are the RMS of the positional difference with NEP097, the RMS of the formal errors, and the ratio between the two. A ratio close to 1 indicates that the formal errors are well-calibrated with respect to the actual solution accuracy, while a large ratio indicates overconfident formal errors. For the observational residuals, the metrics are the RMS of the residuals, the weighted RMS of the residuals, and the cost function value. These residual metrics are not differentiated by direction, as the total values show little variation across estimations.

It has been found that the outcome of the weighting strategy analysis is highly sensitive to the reference trajectory used as the basis for constructing the residuals, and more specifically which weighting scheme achieves the highest agreement with the NEP097 benchmark. For this reason, two distinct analyses are performed and compared, with the most complete one discussed in detail in this chapter and the remaining one presented in [Appendix C](#). The two analyses differ only in their initial conditions. The first is based on a propagation using the IAU Neptune pole model and the NEP097-derived position of Neptune, named IAU_0 , as discussed in [chapter 4](#) with the exact values presented in [Table 5.5](#). The second is performed using the NEP097-fitted pole model and initial state concluded as best in the Pole Estimation Analysis [chapter 5](#), with initial values also presented in [Table 5.5](#). The naming conventions adopted for the two analyses and the individual weighting scheme estimations are summarised in [Table 6.1](#) and [Table 6.2](#) respectively.

The structure of this chapter is as follows. First, the single value metrics for all weighting schemes are presented across both analyses, giving an initial overview of the relative performance of each scheme and allowing conclusions that are robust to the choice of initial conditions. Following this, a more detailed investigation of the WA-FIT-PL

analysis is performed, split into two parts: first the standard weighting schemes are compared with respect to the reference solution, and then the hybrid weighting schemes are evaluated against each other and the best performing standard scheme. For each group, the NEP097 comparison metrics are first examined per RSW direction to identify whether any particular direction dominates the total RMS difference, after which the observation weights and the full RSW positional difference and formal error timeseries are presented and discussed. The analogous results for the second analysis are provided in [section C.3](#). The chapter closes with concluding thoughts on the most appropriate weighting strategy for this estimation problem.

Table 6.1: Naming convention for the two weighting strategy analyses.

Name	Description
WA-IAU	IAU pole model, SPICE Neptune position, initial state estimation only, named IAU_0 in Table 5.5
WA-FIT-PL	NEP097-fitted pole and initial state, initial state and pole libration estimation, named Final in Table 5.5

Table 6.2: Weight-scheme naming conventions used in figures for WA-FIT-PL, analogues for WA-IAU without `ref`, derivation of weighting schemes presented in [chapter 3](#).

Label	Description
<code>ref</code>	Reference run: SimPole rotation model, estimating state + pole librations, scaled per file weight scheme
<code>per file</code>	Per-observation-file RMSE weights
<code>scaled per file</code>	Per-file RMSE weights, scaled by $\sqrt{N_f/T_f}$
<code>per timeframe</code>	Per-timeframe RMSE weights
<code>per timeframe free</code>	Per-timeframe RMSE weights, no upper limit on weight magnitude
<code>scaled hybrid geom.</code>	geometric average weights of scaled per file and per timeframe
<code>scaled hybrid arith.</code>	arithmetic average weights of scaled per file and per timeframe

6.1. Comparison of the Two Analyses

The RMS difference with respect to NEP097 (here assumed to represent the true error) and the ratio of the RMS difference with respect to NEP097 to the RMS of the formal errors are shown in [Figure 6.1](#), for both analyses WA-IAU and WA-FIT-PL. The two analyses differ in three key aspects: the estimated parameters (initial state only versus initial state and pole libration rates), the initial Neptune pole model and initial state (where WA-FIT-PL is already fitted to the observations), and consequently the residuals used to compute the weights, which differ between the two analyses as a result of the different initial conditions.

Before discussing the weighting schemes, it is worth noting the `ref` point visible only in WA-FIT-PL. It has the lowest RMS difference with NEP097 of all schemes, and a true to formal error ratio close to 1. The `ref` run represents the final solution from the Pole Estimation chapter, estimating the initial state together with pole libration rates using real astrometric observations, with initial conditions taken from the Simulated Observations chapter. Those initial conditions were themselves derived by fitting to synthetic three-dimensional positions sampled directly from the NEP097 kernel, meaning the initial pole model and state are already well-aligned with NEP097 before the estimation begins. As the estimation proceeds under different weighting schemes, the solution is pulled away from these initial conditions by the real observational data, which explains why all other WA-FIT-PL schemes produce higher RMS differences with NEP097 despite sharing the same starting point.

Considering the RMS difference with NEP097 shown in [Figure 6.1a](#), the best performing weighting scheme differs between the two analyses. In WA-IAU, the `scaled per file` scheme achieves the lowest RMS difference, whereas in WA-FIT-PL it produces the highest. This reversal suggests that the RMS difference with NEP097 is sensitive to initial conditions, making it a difficult metric for a fair comparison across analyses. Nonetheless, it remains an important metric as it serves as the best available analogue for true error in this research.

The behaviour of `per timeframe` and `per timeframe free` is worth examining more closely. The two schemes are identical in construction, with the only difference being that `per timeframe free` imposes no upper limit on individual timeframe weights, as summarised in Table 6.2. In both analyses, the two schemes produce very similar RMS differences with NEP097, yet their true to formal error ratios diverge dramatically, with `per timeframe free` reaching the highest ratio of all schemes. Removing the weight cap allows certain timeframes to receive extremely high weights, skewing the solution towards those observations. This happens to bring the solution closer to NEP097 in terms of RMS, but simultaneously causes the formal errors to collapse, inflating the ratio significantly. This behaviour is discussed further in the residual analysis of the following sections.

The true to formal error ratio, shown in Figure 6.1b, proves to be a considerably more stable metric across the two analyses. The `per file` scheme, identified in chapter 3 as the most widely used scheme in the literature, consistently produces overconfident formal errors, with ratios well above 1 in both analyses. The `scaled per file` scheme substantially improves on this, bringing the ratio much closer to 1 while requiring minimal additional setup, it differs from `per file` only by the scaling factor $\sqrt{N_f/T_f}$ described in chapter 3. The two hybrid schemes, `scaled hybrid geom.` and `scaled hybrid arith.`, perform similarly to `scaled per file` in terms of the true to formal error ratio, and additionally exhibit lower RMS differences with NEP097 in WA-FIT-PL, making them interesting candidates for further investigation.

In summary, the true to formal error ratio is the most robust metric for comparing weighting schemes across different initial conditions and estimated parameters. The `per file` scheme produces overconfident formal errors in both analyses, while `scaled per file` offers a simple and effective correction. The hybrid schemes are competitive on both metrics and warrant closer examination.

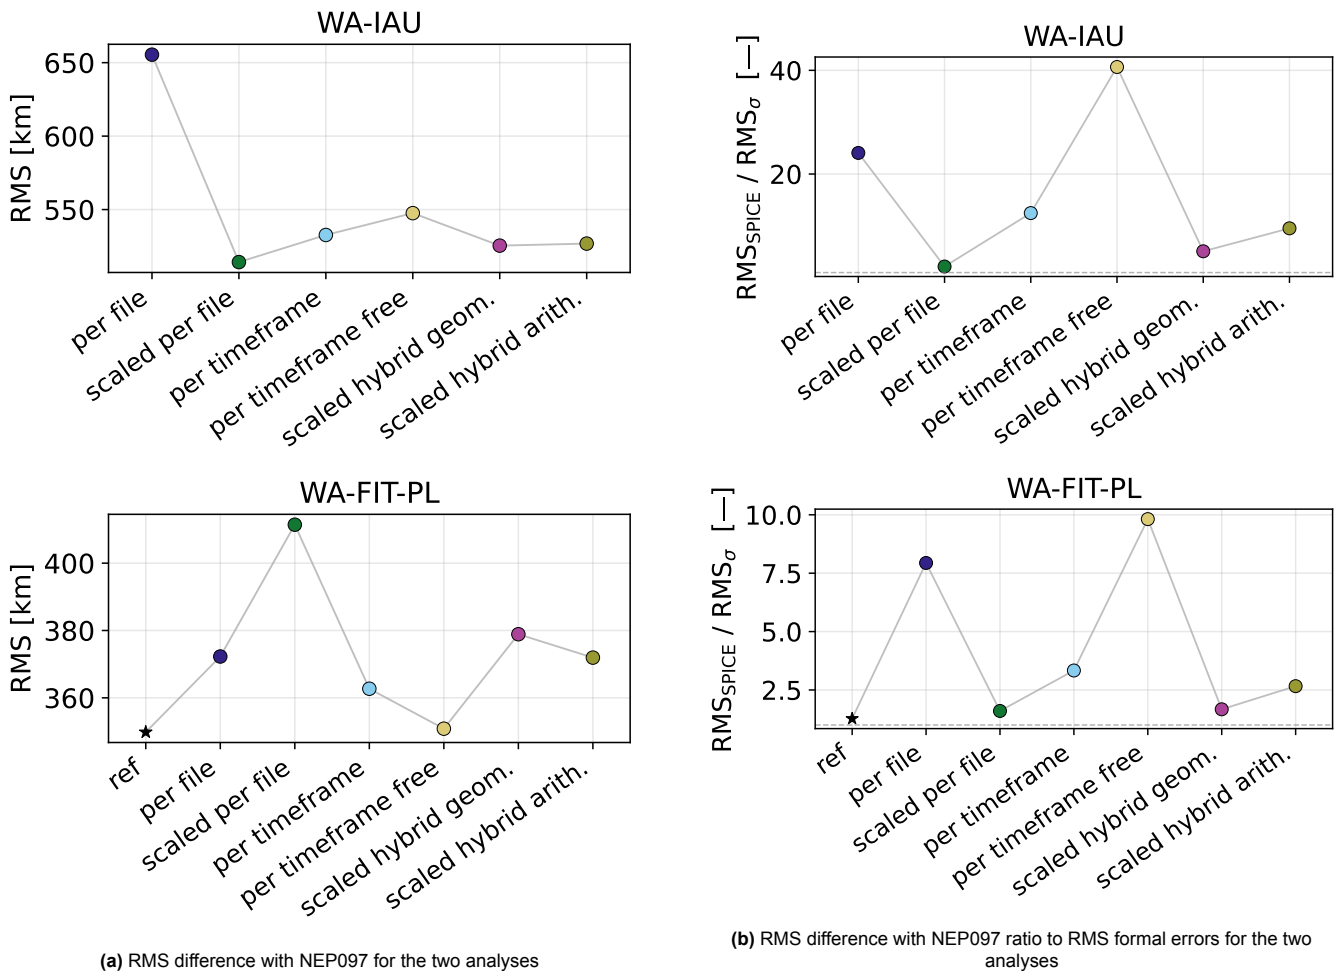


Figure 6.1: Comparison of RMS derived true and formal errors for the two weighting strategy analyses

Besides comparison to the benchmark NEP097, it is important to assess how well each weighting scheme fits the

observational data. Three metrics are presented: the RMS of the post-fit residuals in ??, the weighted RMS in Figure 6.2b, and the cost function in Figure 6.2c. The derivation of these metrics is presented in chapter 3.

The RMS of the residuals is the most interpretable of the three metrics, precisely because it is independent of the weighting scheme. As a consequence, all estimations within a single analysis share the same initial RMS, which is clearly visible in ??. For WA-IAU, the final RMS is reduced by approximately 12 mas relative to the initial value, which is the expected behaviour of a well-posed estimation converging on the observational data. For WA-FIT-PL, the initial and final RMS are essentially identical, consistent with the fact that the initial conditions are already well fitted to the observations. Critically, the difference in final RMS between any two weighting schemes within either analysis is negligible, well below 1 mas. A difference of this magnitude is not conclusive, and no meaningful distinction between the weighting schemes can be drawn on this basis alone.

The weighted RMS and cost function tell a very different story, but not because the underlying solutions differ meaningfully. As shown in Figure 6.2b and Figure 6.2c, both metrics vary enormously across weighting schemes, with initial values spanning orders of magnitude within the same analysis. This spread is entirely a product of the weighting scheme itself: each scheme assigns different weights to the same residuals, so the weighted quantities are not comparable across schemes. Comparing the weighted RMS or cost function between two different weighting schemes is therefore analogous to comparing quantities expressed in different units. The convergence behaviour is equally misleading. For per timeframe free in WA-IAU, the initial cost function value is the largest of all schemes by far, and decreases substantially over the estimation. For per file in WA-FIT-PL, the cost function starts at a similarly large value but barely decreases at all. Yet in both cases the unweighted RMS confirms that the underlying solution has converged normally. The apparent difference in convergence behaviour is therefore a weighting artefact, not a reflection of solution quality.

In conclusion, the observational metrics on their own are insufficient to discriminate between weighting schemes. All schemes converge to solutions that are indistinguishable in terms of unweighted residuals, and the weighted metrics are inherently scheme-dependent and therefore not directly comparable. Without an external benchmark, it is virtually impossible to assess which weighting scheme produces the most reliable solution. Even with a benchmark, one must be careful about how closely that benchmark reflects the true problem, as discussed in the previous section. This motivates the use of the true to formal error ratio as the primary comparison metric, since it relates the solution quality to the internal uncertainty estimates rather than to the weighted residuals alone.

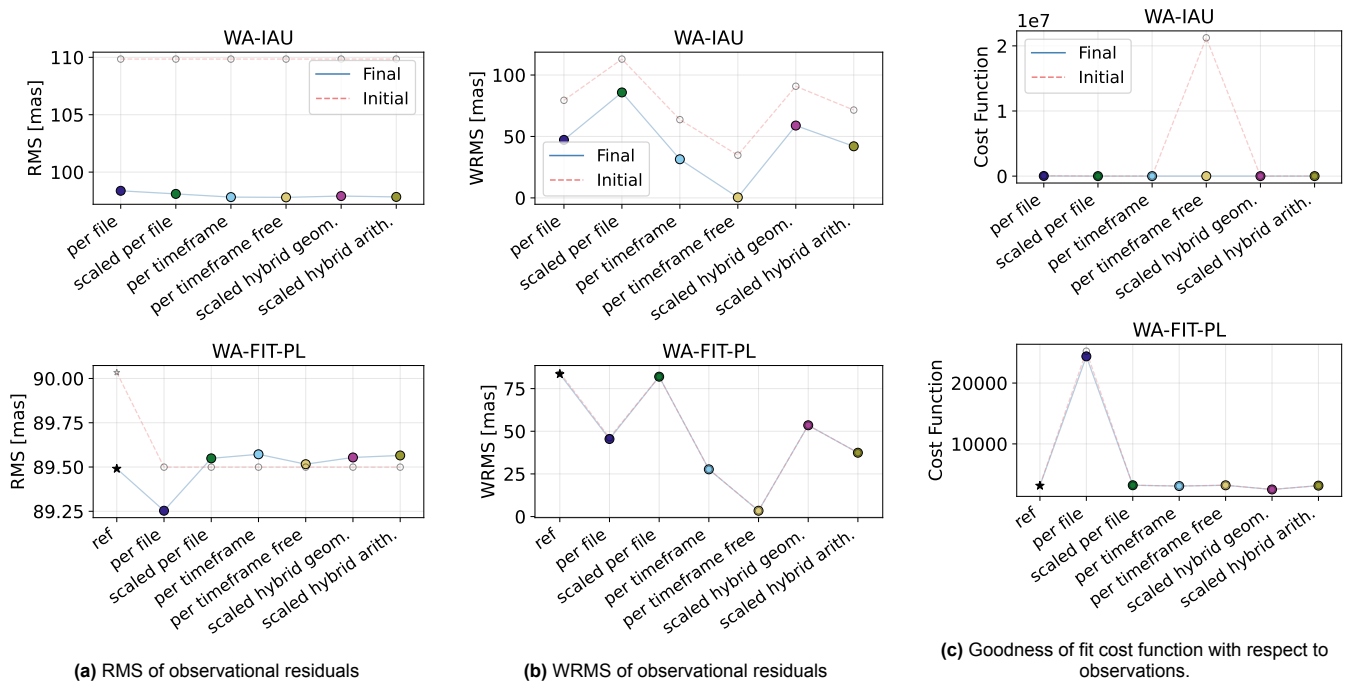


Figure 6.2: Comparison of the observational goodness of fit metrics for the weighting strategy analyses.

6.2. Analysis WA-FIT-PL

In this section the effects of the weighting scheme will be investigated for the analysis WA-FIT-PL. An analogous analysis is made in [Appendix C](#) for the analysis WA-IAU.

This analysis includes, the presentation and discussion of all weights for all weighting schemes. The RSW Statistics analysis and the RSW Timeseries Analysis.

6.2.1. Weights of Observations

The only difference between the estimations considered in this section are the inputted weights for the observations. Therefore, it makes sense to present them here for reference. As the weights themselves are difficult to intuitively grasp ($1/\text{rad}^2$ units with values ranging from 10^{10} to 10^{17}), a more intuitive metric is the uncertainty, defined in [chapter 3](#) as

$$\sigma_{\alpha,i} = \frac{1}{\sqrt{w_{\alpha,i}}} \quad (6.1)$$

and analogously for the declination (δ). The units of this metric are milliarcseconds. The figures below show how each scheme distributes uncertainty across the observations, which reflects the relative trust placed in different data sources or time periods. The standard and hybrid schemes are presented separately, as the hybrid schemes combine elements of both per file and per timeframe weighting and are therefore discussed together.

Standard Weighting Schemes Uncertainties

All uncertainties for the per file and scaled per file weighting schemes are presented in [Figure 6.3](#) and [Figure 6.4](#) respectively. The uncertainties are shown in blue, with a light grey boxplot in the background whose right axis represents the number of observations in each timeframe. As discussed previously, both schemes assign weights on a per-file basis, and individual observation files appear as horizontal lines sharing a common uncertainty. The uncertainty in right ascension and declination within a single file is not always identical, though the order of magnitude is generally consistent between them.

For the per file scheme, several newer observation files between 2010 and 2025 exhibit notably low uncertainties of approximately 10 to 20 mas, as visible in [Figure 6.3](#), while the overall range spans from 10 to approximately 800 mas. The scaled per file scheme presents a markedly different picture: the same low-uncertainty files now exhibit uncertainties of approximately 500 mas, while files covering the period from roughly 2000 to 2013 show comparatively little change.

The difference between the two schemes is governed by a simple scaling factor of $\sqrt{N_f/T_f}$ per file, as defined in [chapter 3](#), where N_f is the number of observations in a file and T_f is the number of timeframes. This ratio is shown in the upper panel of [Figure 6.5](#). Some files exhibit ratios above 80, while larger files typically fall between 40 and 50, corresponding to an uncertainty increase of up to a factor of seven. Files with a ratio close to unity are affected negligibly by the scaling, consistent with what was already observed in the uncertainty figures.

The scaling factor $\sqrt{N_f/T_f}$ effectively acts as a deweighting correction: for dense observation files, where many measurements fall within the same timeframe, correlated noise and systematic biases can accumulate and are suppressed by the increased uncertainty. For sparse files, where observations are spread widely in time and the ratio approaches unity, such temporal correlation is less of a concern and the correction is correspondingly small. However, sparse files introduce a different difficulty: with only one or a few observations per timeframe, the statistical assumptions underlying the RMS-based uncertainty calculation become questionable, as individual observations with atypically low or high residuals can disproportionately influence the assigned uncertainty. The per file schemes partially mitigate this problem by averaging over the full file, whereas the per timeframe schemes are considerably more sensitive to it, as will be discussed below.

The substantial increase in uncertainty for the best-performing observation files between 2015 and 2025 has a clear effect on the ratio of formal to true errors, as shown in [Figure 6.1b](#), while its effect on the difference with the NEP097 benchmark is more ambiguous and dependent on the starting propagation, as shown in [Figure 6.1a](#).

In summary, the scaled per file scheme assigns considerably lower trust to the denser observation files between 2015 and 2025, while leaving files with a ratio N_f/T_f close to unity largely unchanged relative to the per file scheme. This approach already demonstrates favourable results in terms of the formal to true error ratio, as shown in [Figure 6.1b](#), and a more detailed investigation of the RSW time series difference with respect to NEP097 and the formal error evolution is presented in the following sections.




Figure 6.3: Uncertainties derived from weights of all observations for the per file weighting scheme.




Figure 6.4: Uncertainties derived from weights of all observations for the scaled per file weighting scheme.




Figure 6.5: Ratio of number of observations to number of timeframes (top) and total number of observations (bottom) per file.

The uncertainties for the per timeframe and per timeframe free weighting schemes are presented in [Figure 6.6](#) and [Figure 6.7](#) respectively. The only structural difference between them is that per timeframe enforces a minimum uncertainty of 10 mas per timeframe, while per timeframe free imposes no such lower bound.

The resulting uncertainties differ substantially from those of the per file schemes. Individual timeframes exhibit a wide spread, ranging from the enforced lower bound of 10 mas to as high as 2000 mas in some cases. The bulk of observations from approximately 2005 to 2007 cluster around uncertainties of 100 to 600 mas, as indicated by the more saturated blue spread in the figure. The observation files from 2015 to 2025, by contrast, exhibit considerably

more compact uncertainty distributions. This is likely a consequence of those files containing a large number of observations per timeframe, for which the RMS-based uncertainty calculation is statistically well founded. For files with a low N_f/T_f ratio, the per timeframe scheme produces highly volatile uncertainties, with large spread between adjacent timeframes.

The effect of removing the lower bound is clearly visible in the `per timeframe free` scheme, where some timeframes exhibit uncertainties well below 1 mas, a value that is physically implausible for ground-based astrometry and would result in those timeframes dominating the estimation entirely. As shown in [Figure 6.1b](#), the absence of a lower bound drives the formal to true error ratio to unrealistically high values.

In summary, the `per timeframe` scheme offers a physically motivated representation of how observational accuracy evolves over time, but its reliability is fundamentally limited by the number of observations available per timeframe. A lower bound on the per timeframe uncertainty is necessary to prevent sparse timeframes from producing implausible weights.




Figure 6.6: Uncertainties derived from weights of all observations for the `per timeframe` weighting scheme.

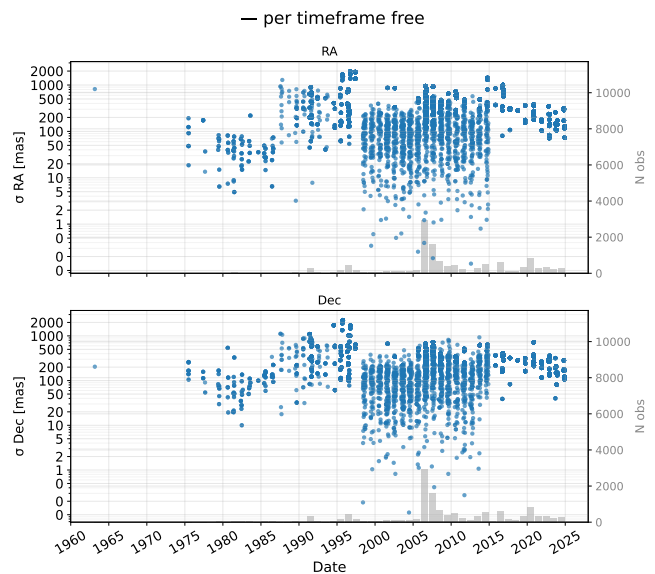


Figure 6.7: Uncertainties derived from weights of all observations for the `per timeframe free` weighting scheme.

In conclusion, the weighting schemes produce sufficiently different uncertainty distributions that the differences are clearly visible in the figures presented above. The key distinction lies in how each scheme handles dense observation files. The `per file` scheme does not correct for temporal correlation or systematic biases within a file, resulting in low uncertainties for the densest modern datasets. The scaled `per file` scheme addresses this by applying the $\sqrt{N_f/T_f}$ correction, substantially downweighting those files while leaving sparse files largely unaffected. The `per timeframe` scheme provides the most temporally resolved picture of observational accuracy, but requires a carefully chosen lower bound to remain statistically meaningful. The complementary strengths and limitations of these schemes motivate the hybrid approach discussed in the following section.

Hybrid Weighting Schemes Uncertainties

The two hybrid weighting schemes `scaled hybrid arith.` and `scaled hybrid geom.` have their uncertainties presented in [Figure 6.8](#) and [Figure 6.9](#), respectively. The figures look similar to the `per timeframe` weighting scheme, with the difference that the largest uncertainties of around 2000 mas seen near 1995 in [Figure 6.6](#) are largely reduced to about 1000 mas for both hybrid schemes. There also seems to be an overall clumping of the uncertainties for timeframes between 2000 and 2015, compared to the `per timeframe` weighting scheme, with values now ranging from approximately 20 to 150 mas, compared to 10 to 400 mas for `per timeframe`. This clumping is much more apparent for the `scaled hybrid geom.` weighting scheme, where the lower bound is lifted to about 30 mas and the uncertainties are more tightly clustered overall.

In summary, both hybrid weighting schemes produce noticeable differences in the computed uncertainties compared to `per timeframe`, which consequently affects the estimation results, formal errors, and convergence behaviour. The

reduction of peak uncertainties from 2000 mas to roughly 1000 mas, and the elevation of the lower bound from 10 mas to approximately 30 mas, are encouraging signs that the hybrid schemes address the shortcomings of the `per timeframe` weighting scheme discussed previously.




Figure 6.8: Uncertainties derived from weights of all observations for the scaled hybrid arith. weighting scheme.




Figure 6.9: Uncertainties derived from weights of all observations for the scaled hybrid geom. weighting scheme.

6.2.2. RSW Statistics

The RMS difference with NEP097, the formal error RMS, and their ratio for all weighting schemes are presented per RSW direction in [Figure 6.10](#). The reference solution, marked with a star, serves as the baseline for comparison throughout. A ratio close to 1 indicates well-calibrated formal errors, while a high ratio indicates overconfident formal errors that do not reflect the true solution uncertainty.

Considering the R and S directions, the weighting schemes fall broadly into three behavioural groups. The first group consists of `per file`, which exhibits the largest RMS difference with NEP097 in both directions, and a ratio well above 1, indicating overconfident formal errors. The second group consists of `per timeframe`, which achieves the lowest RMS difference with NEP097 in R and S, coming closest to the benchmark among all schemes. However, its ratio is elevated compared to `scaled per file` and the hybrid schemes, and as discussed previously, its uncapped variant `per timeframe free` achieves a similar RMS difference at the cost of an extremely inflated ratio, confirming that the good RMS performance of this group comes with reliability concerns. The third group consists of `scaled per file` and the two hybrid schemes `scaled hybrid geom.` and `scaled hybrid arith.`, all of which cluster near the reference solution in RMS difference and produce ratios close to 1, making them the most balanced performers in R and S.

The W direction presents a notably different picture and warrants separate discussion. Here, `per file` achieves the lowest RMS difference with NEP097, while `scaled per file` performs worse in this metric, sitting further from the benchmark. The ratio behaviour is also more moderate across all schemes in W, with no single scheme standing out as an extreme outlier. This is consistent with the finding in previous chapters that the W direction carries the largest uncertainty due to the Neptune pole model, making it less sensitive to the choice of weighting scheme overall. The hybrid schemes again perform comparably to the reference in W in terms of ratio.

Considering all three directions together, `scaled per file` and the two hybrid schemes produce the most consistent performance across both metrics. The hybrid schemes perform similarly to `scaled per file` in terms of ratio, and are competitive in RMS difference. However, the hybrid schemes depend on the `per timeframe` weighting scheme as a component, inheriting its complications, most notably the need for a carefully chosen weight cap to avoid the pathological behaviour demonstrated by `per timeframe free`. The `scaled per file` scheme offers comparable performance without this dependency, requiring only the application of the scaling factor $\sqrt{N_f/T_f}$ to the standard `per file` weights. On this basis, `scaled per file` is selected as the weighting scheme for the final estimation.

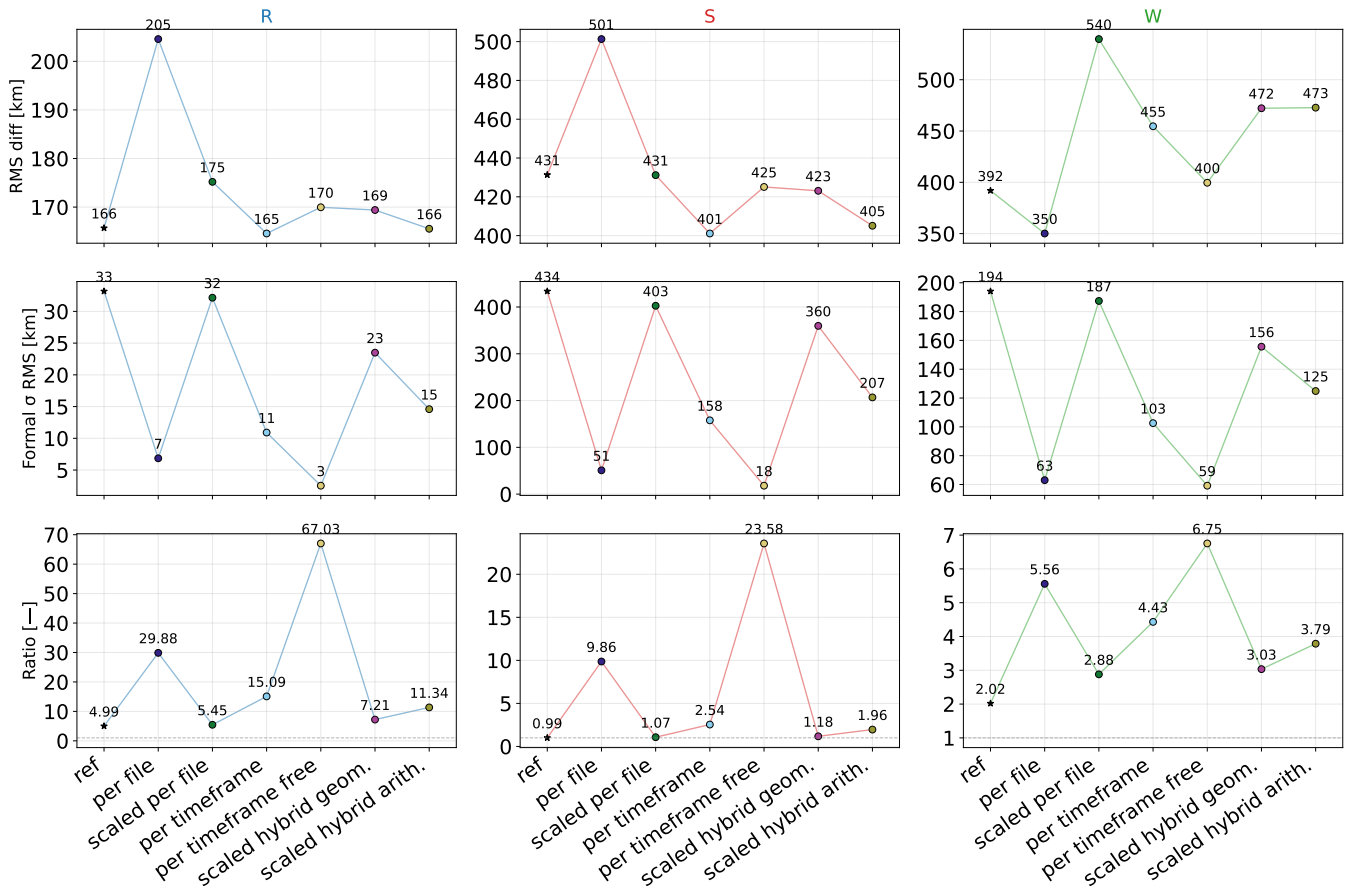


Figure 6.10: RMS difference with NEP097 and formal errors per RSW direction

6.2.3. RSW Timeseries Analysis

This subsection presents the RSW position differences with respect to the benchmark NEP097 and the associated formal errors as timeseries spanning the full estimation range (1960–2025). The subsection is split into two parts: first the base weighting schemes per file, scaled per file, per timeframe, and per timeframe free are compared, followed by the hybrid schemes scaled hybrid geom. and scaled hybrid arith..

Base Weighting Scheme Analysis

The base weighting schemes are compared in three pairings. The comparison between per file and scaled per file is shown in Figure 6.11, the comparison between scaled per file and per timeframe in Figure 6.12, and the comparison between per timeframe and per timeframe free in Figure 6.13.

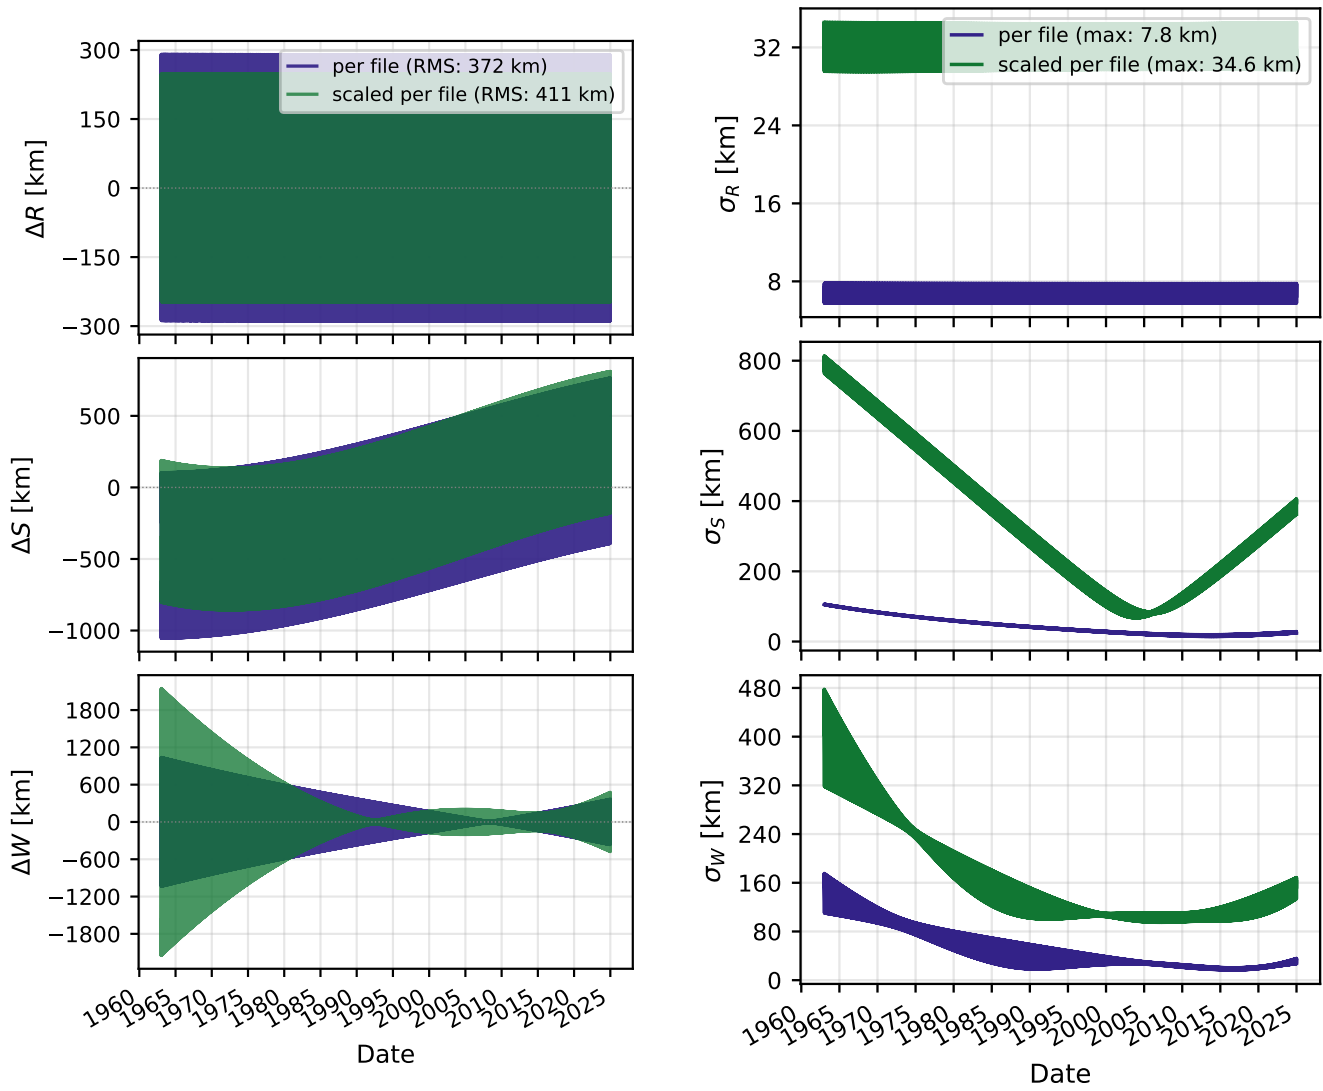
Comparing per file and scaled per file in Figure 6.11, the most striking difference is in the cross-track (W) direction, where per file produces a substantially lower RMS difference with NEP097 of approximately 1000 km, roughly half that of scaled per file. In the along-track (S) direction, however, the two solutions overlap almost entirely in terms of difference with NEP097, despite their formal errors differing by nearly an order of magnitude. The formal errors in the R direction are flat for both schemes, which is physically expected given the strong dynamical and observational constraint in the radial direction. However, the magnitude differs substantially: per file produces formal errors of approximately 7.8 km while scaled per file reaches up to 34.6 km. The same pattern holds across all directions, with per file consistently producing much lower formal errors. The reduction in true error in the W direction does not justify the much lower formal errors across all directions.

Comparing scaled per file and per timeframe in Figure 6.12, the two solutions are nearly identical in terms of difference with NEP097 in the R direction, with the lines lying almost on top of each other throughout the full time range. In the S and W directions, per timeframe performs slightly better, consistent with the marginally lower RMS values reported in Figure 6.10. The formal errors for per timeframe are lower than those of scaled per file across all directions, particularly in W where the difference is most pronounced, though it narrows toward 2025. As

established previously, the reduction in true error achieved by `per timeframe` does not justify the increase in the true to formal error ratio.

Focusing on the `scaled per file` solution in more detail, the R direction shows a difference with NEP097 that remains essentially constant throughout the estimation, with only a small periodic term at the orbital frequency. This suggests either a gravitational parameter mismatch between the dynamical model and the NEP097 kernel, or an uncorrected systematic bias in the dynamical model or observations. The S direction exhibits an almost linear trend from 1960 to 2025, with a slight kink near 1963 most likely caused by the sparse observational coverage in the pre-1985 period, where only a single datapoint exists near 1963. The formal errors in S closely follow the shape of the true error throughout the estimation, and even overestimate it at the endpoints near 1960 and 2025. In the W direction, the difference with NEP097 has a complex shape with two local minima, one near 1990 and another inflection point near 2015, remaining below approximately 600 km between 1990 and 2025 but growing to over 1900 km toward 1960. This growth is again linked to the sparse observational coverage in the pre-1985 period. The formal errors in W broadly follow the shape of the true error but are consistently lower in magnitude, with a ratio of approximately 4 near 1963 and approximately 5 for the period 1995 to 2025.

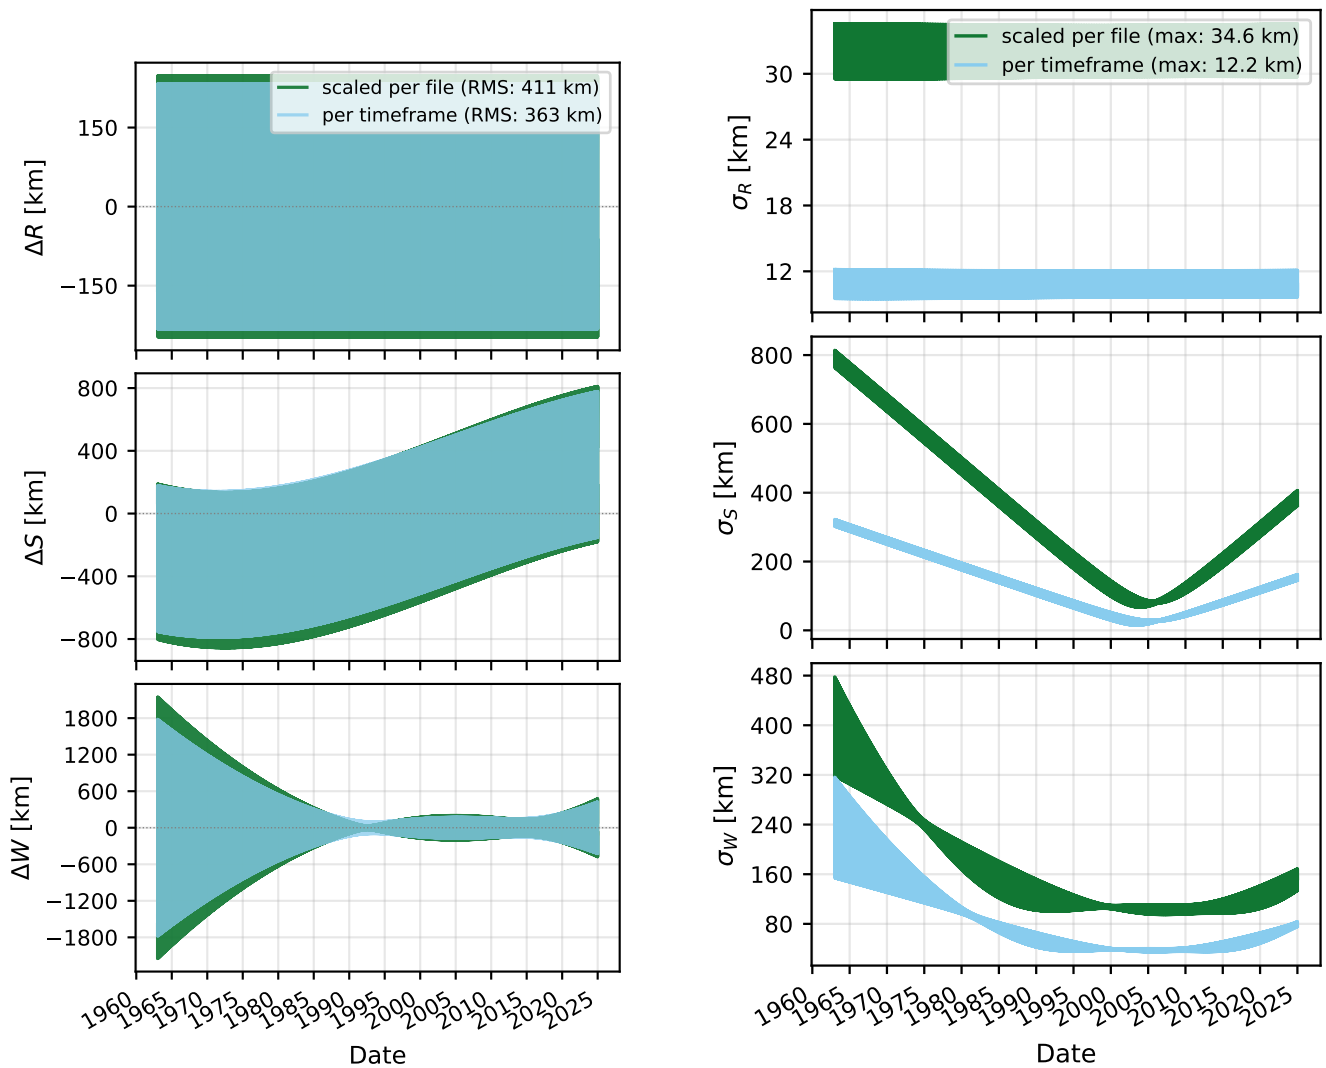
In conclusion, `scaled per file` produces slightly larger true errors than `per file` in the W direction and slightly larger than `per timeframe` across all directions, but it consistently provides the most conservative and therefore most reliable formal errors among the base schemes. The persistent offset in R and the growing divergence toward 1960 in W suggest that the dominant sources of error are systematic rather than statistical, pointing to gravitational parameter mismatches and the inclusion of pole parameters and extended body gravity of Neptune as the most likely candidates for improvement.



(a) RSW position difference with respect to NEP097 for per file and scaled per file weighting schemes.

(b) RSW formal errors for per file and scaled per file weighting schemes.

Figure 6.11: Comparison of RSW position differences with respect to NEP097 (left) and formal errors (right) for per file and scaled per file weighting schemes. From top to bottom: radial (R), along-track (S), and cross-track (W) components. The full time range (1960–2025) is shown.



(a) RSW position difference with respect to NEP07 for scaled per file and per timeframe weighting schemes.

(b) RSW formal errors for scaled per file and per timeframe weighting schemes.

Figure 6.12: Comparison of RSW position differences with respect to NEP07 (left) and formal errors (right) for scaled per file and per timeframe weighting schemes. From top to bottom: radial (R), along-track (S), and cross-track (W) components. The full time range (1960–2025) is shown.

The comparison between per timeframe and per timeframe free is shown in [Figure 6.13](#). The two solutions are nearly identical in terms of RSW difference with NEP07, with per timeframe free achieving only a marginally lower RMS across all three directions. The formal errors, however, differ dramatically: per timeframe free produces formal errors of approximately 3.1 km in R compared to 12.2 km for per timeframe, and the same pattern holds in S and W. This confirms that removing the weight cap does not meaningfully improve the solution, but causes the formal errors to collapse, making per timeframe free an unreliable scheme despite its marginally better true error performance.

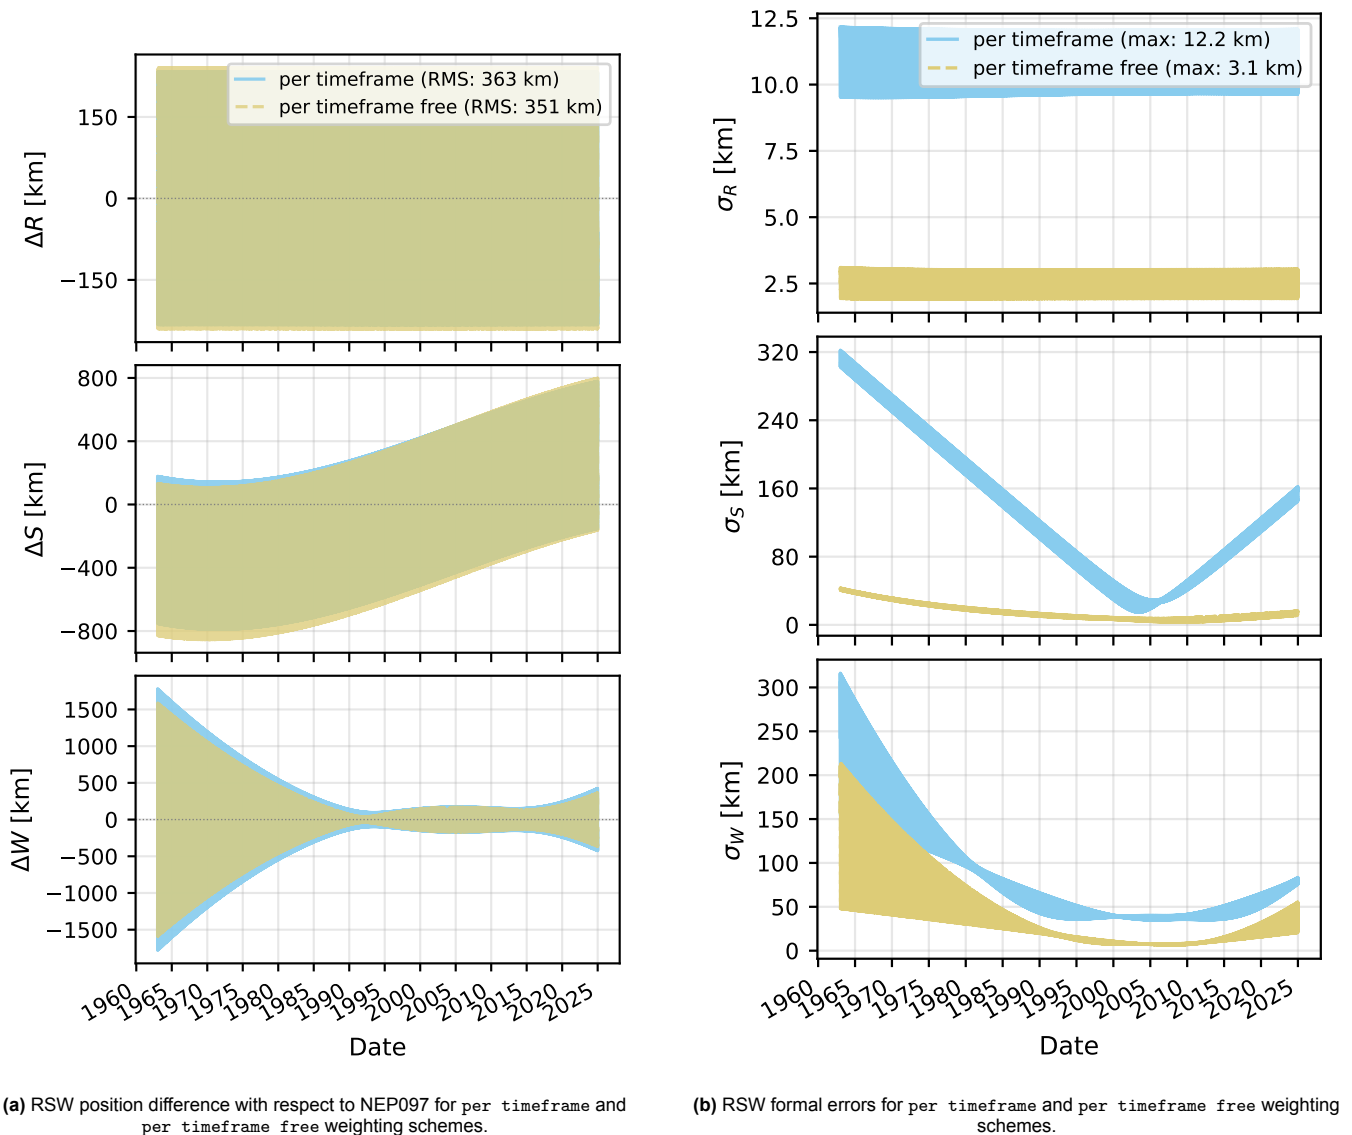


Figure 6.13: Comparison of RSW position differences with respect to NEP097 (left) and formal errors (right) for per timeframe and per timeframe free weighting schemes. From top to bottom: radial (R), along-track (S), and cross-track (W) components. The full time range (1960–2025) is shown.

Hybrid Weighting Scheme Analysis

The comparison between scaled hybrid geom. and scaled hybrid arith. is shown in Figure 6.14 and Figure 6.15. The two schemes produce virtually identical RSW differences with NEP097 across all three directions, with RMS values of 379 km and 372 km respectively, consistent with the results reported in Figure 6.10. The solutions also share the same general temporal shape as scaled per file and per timeframe, confirming that the hybrid averaging does not introduce qualitatively different solution behaviour.

The formal errors, however, differ meaningfully between the two schemes. scaled hybrid arith. produces lower formal errors than scaled hybrid geom. across all directions, most prominently in S and W. In the S direction, scaled hybrid geom. reaches up to approximately 700 km at the endpoints while scaled hybrid arith. stays around 400 km. In R, both schemes produce flat formal errors as expected, at 25.6 km and 16.2 km respectively. This difference is a direct consequence of the averaging method: the geometric mean is more sensitive to large weights from the per timeframe component, pulling the formal errors upward compared to the arithmetic mean. Both schemes sit between scaled per file and per timeframe in terms of formal error magnitude, which is expected given that they are constructed as averages of those two schemes.

As established in Figure 6.10, both hybrid schemes produce true to formal error ratios close to those of scaled per file, while offering no meaningful improvement in true error. Combined with the added complexity of depending on the per timeframe component and the need to carefully manage its weight cap, the hybrid schemes do not present a compelling case over scaled per file as the primary weighting scheme.




Figure 6.14: RSW position difference with respect to NEP097 for scaled hybrid geom. and scaled hybrid arith. weighting schemes.




Figure 6.15: RSW formal errors for scaled hybrid geom. and scaled hybrid arith. weighting schemes.

6.3. Conclusion

Across all evaluated metrics, the weighting schemes produce broadly similar solutions when assessed on the basis of observational residuals alone, confirming the finding of section 6.1 that these metrics are insufficient to discriminate between schemes. The RSW differences with NEP097 are comparable across all schemes, with per file achieving the lowest cross-track (W) difference, per timeframe achieving the lowest radial (R) and marginally lower along-track (S) difference, and the hybrid schemes sitting between per timeframe and scaled per file as expected from their construction.

The most important distinctions emerge when formal errors and the ratio of true to formal errors are considered. The per file scheme produces significantly overconfident formal errors across all directions and both analyses, with ratios well above 1. This is a notable finding, as the per file scheme appears to be the most widely adopted approach in the literature on this topic, though the weighting methodology is often not stated explicitly (Jacobson (2009), Wang

et al. (2023) and Dolgakov and Pavlov (2025)). The `scaled per file` scheme corrects for this by applying the $\sqrt{N_f/T_f}$ scaling factor, producing much more conservative and reliable formal errors at minimal additional cost.

The `per timeframe` scheme requires careful management of a hard cap on individual timeframe weights. Without this cap, timeframes with few observations receive extremely high weights, which skews the solution toward those observations and causes the formal errors to collapse, as demonstrated by `per timeframe free`. The need to tune this cap introduces additional complexity and a degree of arbitrariness that must be carefully justified.

The hybrid weighting schemes `scaled hybrid geom.` and `scaled hybrid arith.` represent an interesting trade-off, combining the more conservative formal errors of `scaled per file` with the temporally resolved weights of `per timeframe`. In practice, however, the improvement in true error over `scaled per file` is negligible, and the formal errors are similar. The added complexity of tuning the `per timeframe` component to construct the hybrid weights is therefore difficult to justify given the marginal gains. Nonetheless, the hybrid approach is a worthwhile direction for future investigation, particularly if the `per timeframe` component can be made more robust.

In summary, `scaled per file` is selected as the weighting scheme for the final estimation. It offers the best compromise between solution accuracy, formal error reliability, and simplicity of setup, and addresses the overconfidence identified in the widely used `per file` scheme with only a minor modification.

7

Comparison with Published Studies and Limitations

Compared to previous studies, this work differs in several respects. The most comprehensive reference, [Jacobson \(2009\)](#), estimates a significantly larger set of parameters, including the gravitational parameters of the Neptune system, the extended gravity field of Neptune, a more detailed pole model for Neptune, and additional observational bias corrections. That study also incorporates Voyager 2 imaging and radio data, the latter of which is not publicly available. As a result, the formal errors reported in [Jacobson \(2009\)](#) are considerably better constrained, with results showing good agreement with the predecessor solution of [Jacobson \(1991\)](#).

Regarding observational biases, [Jacobson \(2009\)](#) identifies a per-night systematic bias for the data from (Qiao et al. 2007), in this study designated as 337 nm0015. Such a bias has not been detected in the current study, as no per-night bias correction has been attempted. Conversely, a declination bias of approximately 200 milliarcseconds has been found in the present work for the data from (Stone & Harris 2000; Stone 2000, 2001), here designated as 689 nm0077. [Jacobson \(2009\)](#) does not mention any declination bias for that dataset. This discrepancy could be caused by differences in the observational model or may indicate that the bias was not identified in that study.


Across all papers to which this study is compared, namely [Jacobson \(2009\)](#), [Wang et al. \(2023\)](#), and [Dolgakov and Pavlov \(2025\)](#), the observational weighting strategy is only briefly described. Based on the limited information available, the weighting approach in all three studies appears consistent with the per-file weighting strategy derived in [chapter 3](#) and investigated in [chapter 6](#). This scheme was almost certainly applied to ground-based astrometry in all cases, while the Voyager 2 data appears to have been weighted separately. In the present study, this weighting scheme was found to produce overly optimistic formal errors.

In [Wang et al. \(2023\)](#), the Voyager 2 data is identified as having the largest effect on constraining Triton's orbit. Their results show close agreement with the NEP097 kernel, with RSW differences of a few kilometres in R, approximately 200 km in S, and around 300 km at most in W. By contrast, the final estimation of the present work yields differences of approximately 200 km in R, 500 km in S, and up to 1 500 km in W. This degradation is attributable primarily to the exclusion of older micrometric datasets and the absence of Voyager 2 data.

The study by [Dolgakov and Pavlov \(2025\)](#) additionally incorporates Voyager 2 and Gaia observations, yielding results in even closer agreement with NEP097, with RSW differences of approximately 20 km in R, 50 km in S, and 80 km in W. That work also notes that Gaia data improves precision in right ascension but is less accurate in declination compared to ground-based datasets. The parameter set estimated in [Dolgakov and Pavlov \(2025\)](#) is broad, encompassing the initial state of Triton, the J_2 coefficient of Neptune, the gravitational parameters of both Neptune and Triton, and Neptune pole orientation parameters. While the present study found that including the gravitational parameters of Neptune and Triton leads to unphysical results when fitting to simulated observations from the NEP097 kernel, as discussed in [chapter 4](#), it is possible that the inclusion of the highly accurate Voyager 2 data alleviates this issue.

Overall, the comparisons highlight that the absence of Voyager 2 data is the dominant limiting factor of the present study. All three reference works that incorporate Voyager 2 observations achieve substantially better agreement with NEP097, with RSW differences roughly an order of magnitude smaller than those obtained here. The progressive

improvement from [Wang et al. \(2023\)](#) to [Dolgakov and Pavlov \(2025\)](#) further demonstrates that the combination of Voyager 2 and Gaia data, together with a richer parameter set, yields the most constrained solution currently available. The present work nonetheless contributes a systematic investigation of weighting strategies and pole model parametrisations using only publicly available ground-based astrometry, and identifies the per-file weighting scheme as a source of optimistic formal errors that likely affects all compared studies to some degree. Future work incorporating Voyager 2 data, Gaia astrometry, and a more complete parameter set would be the natural next step toward a solution competitive with [Jacobson \(2009\)](#) and [Dolgakov and Pavlov \(2025\)](#).



Conclusions and Recommendations

This study investigated the orbit reconstruction of Triton using Earth-based astrometric data, focusing on the sensitivity to dynamical parameters, the estimation of Neptune's pole orientation, and the design of robust observation weighting strategies. The answers to the research questions defined in [chapter 1](#) are presented in [section 8.1](#), a more detailed summary of the findings is given in [section 8.2](#), and recommendations for future work on the estimation of Triton's orbit are presented in [section 8.3](#).

8.1. Answers to Research Questions

RQ1: The Simulated Observations Analysis ([chapter 4](#)) identifies Neptune's pole model parameters and gravitational zonal harmonics (J_2 , J_4) as the dynamical parameters to which Triton's orbit is most sensitive, consistent with the findings of [Jacobson \(2009\)](#), [Wang et al. \(2023\)](#), and [Dolgakov and Pavlov \(2025\)](#). A simulated observations estimation including Triton's initial state, Neptune's pole position, and pole libration parameters achieves sub-kilometre agreement with the NEP097 kernel, demonstrating that the dynamics implemented in Tudat are sufficient to reproduce the reference ephemeris within the estimation timespan. The gravitational parameters of Neptune and Triton could not be reliably estimated, as their inclusion caused the estimation to diverge and produce unphysical values.

RQ2: The Pole Estimation Analysis with Astrometric Data ([chapter 5](#)) demonstrates that including Neptune's pole libration parameters as estimated quantities noticeably reduces the difference with the NEP097 kernel while increasing the formal errors to values more closely matching the actual solution discrepancy. The final solution achieves cross-track differences below 100 km over the modern arc (1990 to 2025), with formal errors that are at most a factor of two smaller than the actual differences in the along-track and cross-track directions.

RQ3: The Weight Scheme Analysis ([chapter 6](#)) derives and compares three base and two hybrid weighting strategies. The conventional per-file scheme, believed to be standard practice in the literature, is found to produce optimistically high weights and too low formal errors. The proposed `scaled per file (ID v2)` weighting strategy, which applies per-timeframe deweighting while assigning weights at the file level, produces formal errors that are consistently coherent with the actual solution accuracy across two independent initialisations.

RQ4: The principal limitation of this work is the exclusive use of Earth-based astrometric data, which lacks the constraining power of the Voyager 2 imaging and radio data incorporated in the NEP097 kernel. The remaining discrepancies with NEP097, particularly the growth of cross-track differences toward the edges of the observational arc and the factor of seven underestimation of radial formal errors, are attributed to this limitation. Additionally, the inability to reliably estimate the system gravitational parameters suggests that either the current observational data are insufficient to constrain them.

8.2. Key Contributions

The aim of this study is to investigate the estimation procedure for accurate orbit reconstruction of Triton using the available astrometric data. This involves estimation of the initial state of Triton and dynamical parameters of the system.

The NEP097 kernel is used as a benchmark to assess the accuracy of the solutions produced in this study. In

all aspects, this kernel is expected to be considerably more accurate than the current results, as it incorporates substantially more observational data, most importantly the Voyager 2 imaging and radio data. The former is believed to be the most constraining dataset for Triton's orbit (Jacobson, 2009; Wang et al., 2023; Tang et al., 2020), while the latter is not publicly available. Additionally, no supporting paper for the NEP097 kernel exists to the knowledge of the author; therefore, all differences between the NEP097 kernel and the results of this research are interpreted as assumptions about the underlying model differences. Nevertheless, the observational residuals are also part of the assessment of the different estimations. For almost all estimations presented in this study, however, the residual differences are negligible, meaning that the estimations are indistinguishable based on the limited observational data used in this research. The residuals between the initial propagation and any estimation with coherent estimatable parameters do produce flat residuals, confirming convergence.

The observational data used in this study are presented in ?? and chapter 2, and consist exclusively of Earth-based astrometric observations from roughly 1963 to 2025. A systematic bias is believed to exist for the dataset (Stone & Harris 2000; Stone 2000, 2001), here designated as 689 nm0077, in the declination direction of approximately 200 milliarseconds. This bias is not mentioned in other papers that employ the same data (Jacobson, 2009; Dolgakov and Pavlov, 2025).

Two Neptune pole models are presented in chapter 3: the Archinal et al. (2018) pole model and the Jacobson (2009) pole model. The Archinal et al. (2018) paper acknowledges the existence of Jacobson's Neptune pole model but states that no effort has been made to compare the two. The effect of both pole models on the orbit of Triton is presented in chapter 4 and more specifically in Figure 4.1. The difference is negligible, with at most 200 km in the cross-track direction. This conclusion is unexpected, as initial estimations (whether with simulated observations or real astrometric data) result in large discrepancies in the cross-track direction with respect to NEP097. This suggests that there is a significant Neptune pole model and/or gravitational zonal harmonics mismatch between the estimation in Jacobson (2009) and the NEP097 kernel, or a reference frame mismatch between the current analysis and those of Jacobson (2009) and the NEP097 kernel, though for the latter no evidence has been found.

The first step of the study is to investigate which dynamical parameters the orbit of Triton is most sensitive to. This analysis is presented in the Simulated Observations Analysis (chapter 4, definition in subsection 3.4.1), where estimations are performed using so-called simulated observations: direct three-dimensional positions of Triton extracted from the NEP097 kernel. The most influential dynamical parameters are found to be the pole model parameters (in this study the conventional IAU pole model representation of Archinal et al. (2018) is used) and the zonal harmonics of the gravitational field of Neptune, which is consistent with the literature (Jacobson, 2009; Wang et al., 2023; Dolgakov and Pavlov, 2025). The effects of the point-mass gravitational parameters of Neptune and Triton have been found to affect the estimation in an unexpected way: the estimation fails to produce meaningful results when it includes the initial state of Triton together with any gravitational parameter (be it Neptune, Triton, or both), producing differences with NEP097 much larger than an estimation with the initial state of Triton only. Additionally, it yields unphysical values for the gravitational parameters, specifically a negative mass for Triton. This behaviour is inconsistent with theory, as Triton's orbit should be most influenced by the gravitational parameter of the system. While Wang et al. (2025) does briefly mention a weak sensitivity of the estimation to Neptune's point-mass gravitational parameter, Jacobson (2009) and Dolgakov and Pavlov (2025) successfully include gravitational parameters of Triton and the Neptunian system as estimated parameters.

The conclusion of the Simulated Observations Analysis is that an estimation with Triton's initial state, Neptune's pole position, and pole libration parameters as estimated quantities produces a solution with sub-kilometre-level differences with the NEP097 kernel. The dynamics included in this research are therefore sufficient to match the NEP097 kernel within the estimation timespan (1963 to 2025). The Triton initial state update is approximately 100 metres in total, which is believed to be well below the accuracy of the NEP097 kernel. The pole position update, on the other hand, is 0.05 and -0.13 degrees in right ascension and declination, while the pole libration updates are -1.77 and 1.3 degrees respectively, as summarised in Table 4.6. These pole parameter updates, specifically the pole libration values, are significant, exceeding the uncertainties presented in Jacobson (2009) by a factor of over 18 000. Two potential causes are proposed: first, the uncertainties presented in Jacobson (2009) are extremely optimistic; second, the pole libration parameters are absorbing effects caused by other dynamical parameters (for example Neptune's gravitational zonal harmonics or the system point-mass gravitational parameters). Additionally, a high correlation between Neptune's pole declination position and libration is found, caused by the short estimation period compared to the precession rate of Neptune's pole; the pole movement is presented in Figure 4.22. This correlation is not problematic for the current study, as the exact values of pole position and libration are less relevant than the shape of the pole movement itself, which is what affects Triton's orbit. Furthermore, this correlation does not render the

estimation numerically unstable. It is suspected that the dynamical constraints employed in [Jacobson \(2009\)](#), which are not used in this study, help alleviate this correlation. The final estimation of the Simulated Observations Analysis is therefore used as the initialisation for the subsequent estimations with real astrometric data, as it captures the mismatch in the cross-track direction through the updated pole model.

As the most influential dynamical parameters are found to be Neptune’s pole model parameters, an analysis including these parameters in the estimation is presented in the Pole Estimation Analysis with Astrometric Data ([chapter 5](#)). The conclusion is that the inclusion of pole libration parameters noticeably decreases the difference with the NEP097 kernel while also increasing the formal errors to values more closely matching the actual differences. This means that the final solution, presented in [Table 5.5](#) and [Figure 5.16](#) uses the available information from the observational data to its full potential. The radial differences with NEP097 are on the order of 200 km, the along-track differences reach values of roughly 600 km, while the cross-track differences are well bounded below 100 km between 1990 and 2025 but increase rapidly from 1990 back to 1963, reaching approximately 1 500 km as presented in [Figure 5.16](#). The formal errors are at most a factor of two smaller than the actual differences with NEP097 in the along-track and cross-track directions, where for the most part they follow the differences closely; in the radial direction, however, the differences are nearly seven times larger than the formal errors. These deviations are much larger than those of the initialisation propagation (which achieved sub-kilometre-level agreement with the NEP097 kernel). While the observational data residuals do not change by more than 1 mas, the estimation appears to be pulled away from the orbit closely resembling the NEP097 kernel by the observational data. The updates on the initial state of Triton and the pole libration parameters are also significant, with over 300 km total positional change and 4.6 and 1.2 degrees change in pole libration right ascension and declination. The largest positional change is in the cross-track direction, which is attributed to the astrometric data, as this was not found to be the case in the Simulated Observations Analysis ([chapter 4](#)).

The Weight Scheme Analysis ([chapter 6](#)) is the most original contribution of this report. All papers consulted during this research only briefly describe the weighting strategy used to perform their estimations; this research aims to shed light on different weighting strategies and to assess their performance. Three base and two composite weighting strategies are mathematically derived and discussed in [chapter 3](#), and the resulting estimations with real astrometric data are presented and compared in [chapter 6](#). The benchmark weighting strategy, named `per file (ID)`, is believed to be the weighting strategy used for Earth-based astrometric data by most previous papers, specifically [Jacobson \(2009\)](#), [Wang et al. \(2023\)](#), [Tang et al. \(2020\)](#), and [Dolgakov and Pavlov \(2025\)](#). The findings suggest that this weighting strategy produces optimistically high weights, steering the solution toward observational data that may not be as accurate as estimated by this scheme, and additionally producing extremely low formal errors. The `per timeframe (TF)` weighting scheme tends to produce moderately low formal errors but is unstable for some datasets, specifically those with a low ratio of observations to timeframes N_f/T_f , where certain timeframes with few observations and low residuals with respect to the initial propagation receive extremely high weights. The `scaled per file (ID v2)` weighting strategy employs a per-timeframe deweighting procedure while assigning weights at the file level. This combination ensures stability even for files with low N_f/T_f ratios, while overall assigning lower weights and therefore producing coherent formal errors. Additionally, it has been found that the generation of weights (from the residuals of the initial propagation) is significantly affected by the initial propagation itself; the analysis is therefore conducted twice for two different initial propagations (and consequently slightly different weights). The weighting strategy with the lowest difference with respect to NEP097 is debatable and highly dependent on the initial propagation, while the formal error coherence with the actual differences is stable across both sets of estimations, suggesting that the `scaled per file` weighting strategy is sufficiently robust for this and similar estimation problems.

Additional weighting strategies named `hybrid weights` are constructed as arithmetic or geometric means of the two most prominent schemes: `scaled per file` and `per timeframe`. The weights are therefore assigned per timeframe, but the inherent instability of the `per timeframe` weights is partially alleviated. Specifically, the geometric mean variant (`hybrid G+v2`) appears to perform well. While further investigation of this weighting scheme has not been conducted, it shows promising results for cases where a more detailed per-timeframe analysis of the observational data is pursued.

In summary, this study demonstrates that accurate orbit reconstruction of Triton from Earth-based astrometric data alone is feasible, achieving cross-track agreement with the NEP097 kernel below 100 km over the modern observational arc (1990 to 2025). The inclusion of Neptune’s pole libration parameters as estimated quantities is found to be essential for capturing the dominant dynamical mismatch, while the proposed `scaled per file` weighting strategy provides a robust and well-calibrated alternative to the conventional per-file approach used in the literature. The formal errors produced by this combination are statistically consistent with the actual solution accuracy, giving confi-

dence that the reported uncertainties are meaningful. Nevertheless, the absence of Voyager 2 data and the inability to reliably estimate the gravitational parameters of the system remain the principal limitations of this work.

8.3. Recommendations for Future Work

The research questions identified during the initial phase of this study ([Appendix A](#)) encompassed a substantially broader scope than could be addressed within the available time. Several of the unaddressed or partially addressed subtopics represent promising directions for future work, organised below by theme.

Observational Data

The most impactful improvement to Triton's orbit reconstruction would be the inclusion of Voyager 2 imaging data, which is believed to provide the strongest constraints on the orbit ([Jacobson, 2009](#); [Wang et al., 2023](#); [Tang et al., 2020](#)). The reduced astrometric measurements from [Jacobson \(1991\)](#) are publicly available and their incorporation would directly address the principal limitation identified in this study (RQ4). Additionally, observations from the Hubble Space Telescope ([Showalter et al., 2019](#)) and Gaia ([Tanga et al., 2023](#)) could further strengthen the solution, though the latter may suffer from brightness-dependent astrometric degradation, also reported by [Dolgakov and Pavlov \(2025\)](#) (Question 1e). A systematic investigation of biases in historical datasets (Question 1d), specifically micrometric observations, is also an open topic, as papers that include such data ([Jacobson, 2009](#); [Wang et al., 2025](#)) tend to only briefly explain their assumptions and weighting strategies. The sensitivity of the solution to the inclusion or exclusion of specific datasets (Questions 1a, 1c) also remains unexplored and could inform future observation campaign planning, as proposed by [Emelyanov et al. \(2023\)](#), albeit for different moons of Neptune and other celestial bodies.

Dynamical Parameters

The inability to reliably estimate the point-mass gravitational parameters of Neptune and Triton is the most unexpected finding of this study and warrants further investigation. Possible avenues include reformulating the estimation to use the system gravitational parameter GM_{sys} rather than individual masses, tightening the a priori constraints, or investigating whether the inclusion of Voyager 2 data resolves the divergence. The estimation of tidal dissipation parameters (Question 2d), following the approach of [Wang et al. \(2025\)](#), would provide constraints on Neptune's interior structure but likely requires the additional data types mentioned above. The gravitational influence of Neptune's inner satellites on Triton's orbit (Question 2e) remains an open question; even a negative result would help quantify upper bounds on their perturbative effect.

Modelling Framework

The current study uses a single planetary ephemeris (NEP097 kernel). An assessment of the sensitivity to the choice of ephemeris (Question 3a), for example comparing DE421 and DE440 or solutions proposed by [Wang et al. \(2023\)](#), would quantify a source of systematic uncertainty that is currently unaccounted for. The implementation of corrections specific to older observation types, such as photographic plate scale factors and reference star catalogue biases (Question 3c), as well as improved Earth rotation and atmospheric refraction models for micrometric data (Question 3d), would enable the reliable use of pre-1963 observations and extend the temporal baseline of the estimation.

Analysis and Uncertainty Quantification

Several analysis techniques identified in the initial research plan remain unimplemented. A formal consider parameter analysis (Question 4b) would allow the uncertainties from unadjusted dynamical parameters, such as the gravitational zonal harmonics, to be propagated into the final covariance without requiring their direct estimation. Linearised covariance propagation (Question 4d) could be used to forecast the growth of orbital uncertainty over the coming decades (Question 4e), which is directly relevant for planning future observations and potential mission support. The prediction of mutual events or stellar occultations (Question 4f) and a study of Triton's long-term orbital evolution (Question 4g) represent scientifically valuable extensions that become feasible once the orbital solution is sufficiently accurate.

Weighting Strategy


The `scaled per file` weighting strategy proposed in this work could be further developed in several ways. A more rigorous per-timeframe analysis of observation quality, informed by metadata such as seeing conditions, instrument type, and reference catalogue version, would allow the weights to reflect known sources of heterogeneity rather than relying solely on residual statistics. The hybrid geometric mean variant (`hybrid G+v2`), which showed promising

results but was not investigated in depth, merits further study as it offers a natural extension toward per-timeframe weighting without the associated instability. Finally, the dependence of the weights on the initial propagation, identified in this study, suggests that an iterative reweighting procedure, in which the weights are recomputed from updated residuals after each estimation, could improve both the robustness and the formal error calibration of the solution.

Bibliography

- B. A. Archinal, C. H. Acton, M. F. A'Hearn, A. Conrad, G. J. Consolmagno, T. Duxbury, D. Hestroffer, J. L. Hilton, R. L. Kirk, S. A. Klioner, D. McCarthy, K. Meech, J. Oberst, J. Ping, P. K. Seidelmann, D. J. Tholen, P. C. Thomas, and I. P. Williams. Report of the IAU working group on cartographic coordinates and rotational elements: 2015. *Celestial Mechanics and Dynamical Astronomy*, 130(22), 2018. doi: 10.1007/s10569-017-9805-5.
- Marina Brozović and Robert A. Jacobson. Orbits of the irregular satellites of uranus and neptune. *Astronomical Journal*, 2022. doi: 10.3847/1538-3881/ac4e9c.
- Christopher F. Chyba and Cynthia B. Phillips. Europa as an abode of life. *Origins of Life and Evolution of the Biosphere*, 32(1):47–67, 2002. doi: 10.1023/A:1013958519734.
- Fabien Dahmani. *Orbit Estimation of Small Jovian Satellites*. PhD thesis, TU Delft / Tudat, 2024. Master's thesis.
- R. Deakin. *Notes on Least Squares*. Melbourne, Australia, 2005.
- Dominic Dirx, Marie Fayolle, Geoffrey Garrett, Miguel Avillez, Kevin Cowan, Sean Cowan, João Encarnacao, Carlos Fortuny Lombrana, Jérémie Gaffarel, Jonas Hener, Xuanyu Hu, Maarten van Nistelrooij, Filippo Oggioni, and Michael Plumaris. The open-source astrodynamics tudatpy software: Overview for planetary mission design and science analysis. In *16th Europlanet Science Congress (EPSC)*, 2022.
- Ivan Dolgakov and Dmitry Pavlov. Improving the ephemerides of Neptune's satellites with Gaia, Voyager 2, and ground-based data. *Celestial Mechanics and Dynamical Astronomy*, 137(4):20, 2025. doi: 10.1007/s10569-025-10250-6.
- N. V. Emelyanov and M. Yu. Samorodov. Analytical theory of motion and new ephemeris of Triton from observations. , 454(2):2205–2215, December 2015. doi: 10.1093/mnras/stv2116.
- Nikolay V. Emelyanov. *The Dynamics of Natural Satellites of the Planets*. Elsevier, Amsterdam, Netherlands, 2021. ISBN 9780128227046. doi: 10.1016/B978-0-12-822704-6.00006-6.
- Nikolay V. Emelyanov, Mikhail Kovalev, and Maxim I. Varfolomeev. The orbits of outer planetary satellites using the gaia data. *Monthly Notices of the Royal Astronomical Society*, 522:165–179, 2023. doi: 10.1093/mnras/stac223.
- Sam Fayolle. *Natural satellites ephemerides: The Galilean moons' dynamics in the JUICE-Europa Clipper era*. Phd thesis, Delft University of Technology, 2025. URL <https://doi.org/10.4233/uuid:667192b7-5555-49df-b2c6-2479942fc5f2>.
- Gaia Collaboration, S. A. Klioner, L. Lindegren, F. Mignard, J. Hernández, M. Ramos-Lerate, U. Bastian, M. Biermann, A. Bombrun, A. de Torres, E. Gerlach, R. Geyer, T. Hilger, D. Hobbs, U. L. Lammers, P. J. McMillan, H. Steidelmüller, D. Teyssier, C. M. Raiteri, S. Bartolomé, M. Bernet, J. Castañeda, Y. Frémat, and ... Gaia early data release 3: The celestial reference frame (gaia□crf3). *Astronomy & Astrophysics*, 667:A148, 2022. doi: 10.1051/0004-6361/202243483.
- Johann Galle and Heinrich Louis d'Arrest. Observation of the planet neptune based on le verrier's prediction. *Astronomische Nachrichten*, ??(??):??, 1846. discovery of Neptune following Le Verrier's calculation.
- Asaph Hall. Observations made with 26-inch refractor of the u.s. naval observatory at washington. *Astronomische Nachrichten*, 90:161, 1877.
- Institute of Celestial Mechanics and Computation of Ephemerides (IMCCE). Natural satellite database: Bnepomae. <https://nsdb.imcce.fr/obspos/bnepomae.htm>, 2025. Accessed: 2025-08-01.
- R. A. Jacobson. The orbits of triton and nereid from spacecraft and earth-based observations. NASA Technical Report JPL D-10096, Jet Propulsion Laboratory, California Institute of Technology, 1991. URL <https://ntrs.nasa.gov/citations/19920030016>. NASA Technical Memorandum 19920030016.

- R. A. Jacobson. The orbits of the neptunian satellites and the orientation of the pole of neptune. *Astronomical Journal*, 137:4322–4329, 2009. doi: 10.1088/0004-6256/137/5/4322.
- V. Lainey. Quantification of tidal parameters from solar system data. *Celestial Mechanics and Dynamical Astronomy*, 2016. URL <https://arxiv.org/abs/1604.04184>. Overview of k_2 and Q estimation methods.
- William Lassell. Discovery of triton, a satellite of neptune. *Monthly Notices of the Royal Astronomical Society*, ?? (??):??, 1846. first detection of Neptune’s moon Triton.
- Mikulski Archive for Space Telescopes (MAST). Mikulski archive for space telescopes. <https://archive.stsci.edu/>, 2025. Accessed: 2025-08-01.
- Oliver Montenbruck and Eberhard Gill. *Satellite Orbits: Models, Methods and Applications*. Springer, 2000. ISBN 354067280X.
- C. D. Murray and S. F. Dermott. *Solar System Dynamics*. Cambridge University Press, Cambridge, UK, first edition, 1999.
- Suely Oliveira and David E. Stewart. *Writing Scientific Software: A Guide to Good Style*. Cambridge University Press, 2006. ISBN 9780521675956.
- Vincent Robert, D. Pascu, V. Lainey, and J.-E. Arlot. Naroo program: Analysis of usno galilean observations 1967–1998. *Icarus*, page 116344, 2025. doi: 10.1016/J.ICARUS.2024.116344.
- M. R. Showalter, I. de Pater, J. J. Lissauer, and R. S. French. The seventh inner moon of neptune. *Nature*, 566(7744): 350–353, Feb 2019. doi: 10.1038/s41586-019-0909-9. URL <https://doi.org/10.1038/s41586-019-0909-9>.
- B. A. Smith et al. Voyager imaging science subsystem: Narrow and wide angle television cameras. *Space Science Reviews*, 21:103–127, 1977. Instrument and mission overview for Voyager ISS cameras.
- Bradford A. Smith, Laurence A. Soderblom, D. Banfield, C. Barnet, A. T. Basilevsky, R. A. Beebe, K. Bollinger, T. V. Johnson, R. L. Kirk, D. Morrison, S. A. Nelson, C. C. Porco, J. B. Pollack, C. Sagan, E. M. Shoemaker, M. P. Sulkaneen, R. G. Strom, V. E. Suomi, and J. N. Veverka. Voyager 2 at neptune: Imaging science results. *Science*, 246(4936):1422–1449, 1989. doi: 10.1126/science.246.4936.1422.
- Frank Sohl et al. Subsurface oceans and deep interiors of medium-sized outer planet satellites and large trans-neptunian objects. *Icarus*, ??(??):??, 2002. Modeling of internal structure and oceans in icy bodies.
- K. Tang, Y.-Z. Song, et al. The orbit of triton with new precise observations and the inpop19a ephemeris. *Astronomy & Astrophysics*, 641:A108, 2020. doi: 10.1051/0004-6361/202038556.
- P. Tanga, T. Pauwels, F. Mignard, K. Muinonen, A. Cellino, P. David, D. Hestroffer, F. Spoto, J. Berthier, J. Guiraud, W. Roux, B. Carry, M. Delbo, A. Dell’Oro, C. Fournon, and L. Galluccio. Gaia data release 3: The solar system survey. *Astronomy & Astrophysics*, 674:A12, 2023. doi: 10.1051/0004-6361/202243796.
- Robert H. Tyler. Strong ocean tidal flow and heating on moons of the outer planets. *Nature*, 456(7223):770–772, 2008. doi: 10.1038/nature07571.
- Sean E. Urban and P. Kenneth Seidelmann, editors. *Explanatory Supplement to the Astronomical Almanac*. University Science Books, Mill Valley, CA, 3rd edition, 2012. ISBN 9781891389856.
- Bo Wang, Jianguo Yan, Wutong Gao, Ye Yuan, Shangbiao Sun, Mao Ye, and Jean-Pierre Barriot. The neptunian gravity estimated from the motion of triton based on astrometric observations. *Astronomy & Astrophysics*, 671: A70, 2023. doi: 10.1051/0004-6361/202243796.
- Bo Wang, Xi Lu, Jianguo Yan, Jiawen Wang, Xiaowen Duan, and Wutong Gao. A plausible minimum value of the neptunian tidal dissipation factor estimated from triton’s astrometric observations. *Solar System Research*, (1), 2025. doi: 10.1134/S0038094624601440.
- Ye Yuan, Fan Li, Yanning Fu, and Jian Chen. The orbits of Triton and Nereid and the pole orientation of Neptune from Voyager, Hubble Space Telescope, and Earth-based astrometry in 1847-2020. , 654:A66, October 2021. doi: 10.1051/0004-6361/202140739.



Initial Research Questions and Mapping

During the initial phase of this study, a comprehensive set of research subtopics was formulated, organised into four categories: Data Evaluation and Processing, Estimation of Physical and Dynamical Parameters, Modelling Framework and Technical Implementation, and Analysis, Uncertainty, and Scientific Implications. The full list of original questions is reproduced below, followed by a mapping to the research questions addressed in this report.

Mapping to Final Research Questions

[Table A.1](#) maps the original questions to the final research questions of [section 1.3](#). Questions that are not addressed in this work are marked accordingly and are recommended as directions for future research in [section 8.3](#).

Table A.1: Mapping between the original research subtopics and the final research questions. Fully addressed questions are marked with ✓, partially addressed with ~, and unaddressed with —.

Original Question	Status	Maps to	Notes
1a (Data type benefits)	~	RQ4	Earth-based data only
1b (Weighting and biases)	✓	RQ3	Core of Weight Scheme Analysis
1c (Sensitivity to older data)	—	—	Not addressed
1d (Historic dataset biases)	—	—	Not addressed
1e (Gaia brightness limits)	—	—	Not addressed
1f (Time standard corrections)	—	—	Not addressed
2a (Initial state constraint)	✓	RQ1, RQ2	Both analyses
2b (Pole and harmonics sensitivity)	✓	RQ1	Simulated Obs. Analysis
2c (System GM constraint)	~	RQ1, RQ4	Negative result
2d (Tidal dissipation)	—	—	Not addressed
2e (Inner satellite masses)	—	—	Not addressed
2f (Observational biases)	~	RQ3	One bias identified
3a (Planetary ephemeris choice)	—	—	Not addressed
3b (Tudat pole precession)	✓	RQ1, RQ2	Implemented
3c (Old observation corrections)	—	—	Not addressed
3d (Earth rotation model)	—	—	Not addressed
3e (Justified model complexity)	✓	RQ1, RQ4	Simulated Obs. Analysis
4a (Data batch contributions)	~	RQ3	Per-file RMS analysis
4b (Consider parameters)	~	RQ4	Identified, not formalised
4c (Weighting robustness)	✓	RQ3	Two initialisations tested
4d (Covariance propagation)	—	—	Not addressed
4e (Forecast uncertainty)	—	—	Not addressed
4f (Mutual event prediction)	—	—	Not addressed
4g (Long-term evolution)	—	—	Not addressed

B

Integrator Step Size Validation

The numerical integrator used throughout this study is a fixed-step Runge–Kutta–Fehlberg 7(8) scheme. To select an appropriate step size, a convergence analysis is performed by propagating Triton’s orbit over the half the estimation arc with step sizes of 28 800 s (8 h), 14 400 s (4 h), 7 200 s (2 h), 3 600 s (1 h), 1 800 s (30 min), and 900 s (15 min), and comparing successive solutions.

Figure B.1 shows the absolute position difference between consecutive step size pairs as a function of time. For the larger step sizes, the differences grow monotonically over the arc, indicating that truncation error dominates and that the solution has not yet converged. As the step size decreases, the differences reduce by several orders of magnitude, consistent with the expected convergence behaviour of a high-order integrator. Between the 3 600 s and 1 800 s pair, the difference at the end of the arc is on the order of 100 m. However, the 1 800 s to 900 s pair shows a qualitatively different behaviour: the difference no longer decreases monotonically but instead exhibits an irregular, oscillatory pattern at the sub-metre level. This indicates that further reduction of the step size does not improve the solution, as floating-point precision loss begins to dominate the integration error. As discussed in subsection 3.1.3, this is a consequence of the integration being performed in the Solar System Barycentre frame, where Triton’s absolute coordinates are on the order of 30 AU.

Figure B.2 summarises the accumulated position difference at the end of the simulation as a function of step size. The curve shows a clear convergence trend down to a step size of approximately 1 800 s, below which the accumulated difference flattens, confirming the onset of numerical noise dominance.

Based on this analysis, a step size of 3 600 s (1 hour) is selected for all propagations in this study. This choice provides a margin of approximately one order of magnitude above the numerical noise floor while keeping the truncation error well below the accuracy of the astrometric observations (on the order of tens of milliarcseconds, corresponding to positional uncertainties of tens of kilometres at Neptune’s distance).




Figure B.1: Absolute position difference between propagations with consecutive step sizes as a function of time. The transition from truncation-error-dominated behaviour (upper curves) to numerical-noise-dominated behaviour (1 800 to 900 s pair) is clearly visible.





Figure B.2: Accumulated position difference at the end of the simulation arc as a function of integrator step size. The flattening below 1 800 s indicates the onset of floating-point precision loss.



Additional Figures for Weight Strategy Comparison Across Estimation Analyses

C.1. Hybrid Weighting Scheme Uncertainties

This appendix derives the observation uncertainties corresponding to the hybrid weighting schemes defined in [Equation 3.2.2](#).

C.1.1. Geometric Mean Hybrid

For the geometric mean hybrid, the weight for timeframe t in file f is $w_{t,f}^{(g)} = \sqrt{w_f \cdot w_t}$. The corresponding uncertainty is:

$$v_{t,f}^{(g)} = \frac{1}{\sqrt{w_{t,f}^{(g)}}} = \frac{1}{(w_f \cdot w_t)^{1/4}} \quad (\text{C.1})$$

which can be expressed as the geometric mean of the constituent uncertainties:

$$v_{t,f}^{(g)} = \sqrt{v_f \cdot v_t} \quad (\text{C.2})$$

Substituting the expressions for v_f and v_t in terms of residuals yields, for the scaled per-file + per-timeframe variant:

$$v_{t,f}^{(g)} = \sqrt{\sqrt{\frac{N_f}{T_f}} \sqrt{\frac{1}{N_f} \sum_{i=1}^{N_f} (r_{i,f})^2} \cdot \sqrt{\sum_{i=1}^{n_t} (r_{t,i})^2}} \quad (\text{C.3})$$

which simplifies to:

$$v_{t,f}^{(g)} = \left(\frac{N_f}{T_f}\right)^{1/4} \left(\frac{1}{N_f} \sum_{i=1}^{N_f} (r_{i,f})^2\right)^{1/4} \left(\sum_{i=1}^{n_t} (r_{t,i})^2\right)^{1/4} \quad (\text{C.4})$$

For the per-file + per-timeframe variant (without the scaling factor):

$$v_{t,f}^{(g)} = \sqrt{\sqrt{\frac{1}{N_f} \sum_{i=1}^{N_f} (r_{i,f})^2} \cdot \sqrt{\sum_{i=1}^{n_t} (r_{t,i})^2}} \quad (\text{C.5})$$

C.1.2. Arithmetic Mean Hybrid

For the arithmetic mean hybrid, $w_{t,f}^{(a)} = (w_f + w_t)/2$, and the uncertainty is:

$$v_{t,f}^{(a)} = \sqrt{\frac{2}{w_f + w_t}} \quad (\text{C.6})$$

Unlike the geometric mean case, this does not simplify to a closed-form combination of v_f and v_t , which is one reason the geometric mean variant is preferred for interpretability.

C.2. Additional Uncertainties figures for WA-FIT-PL

C.2.1. Per File Uncertainty

The uncertainty metric is a useful tool to assess the weight magnitude per observation file and the expected positional accuracy in mas for each file under each weighting scheme. The exact derivation is presented in [chapter 3](#). The per-file uncertainties for the three main weighting schemes, ID, ID $\nu 2$ and TF, are presented in [Figure C.1](#).

Looking at the figure, the ID weighting scheme assigns noticeably lower uncertainties (i.e. higher weights) compared to the other two schemes. This is because ID lacks a per-timeframe deweighting strategy, meaning any correlated noise or biases cannot be efficiently captured by the weights. This is reflected in the formal errors discussed in previous sections, where the ratio between the difference with NEP097 and the formal errors is large.

In contrast, ID $\nu 2$ and TF exhibit similar trends, with noticeable differences for some files. It is important to note that for the TF scheme, weights are assigned per timeframe rather than per file, so the values shown here represent a per-file average, while weights within the same file but across different timeframes may vary substantially.

However, this figure can be misleading in isolation, as both the number of observations and their temporal spread vary significantly from file to file. A useful companion figure is [Figure 6.5](#), which shows the ratio of observations to timeframes N_f/T_f for each file. A ratio of 1 indicates exactly one observation per timeframe. As derived in [chapter 3](#), when $N_f/T_f = 1$, the ID $\nu 2$ weights reduce exactly to the ID weights, since the scaling factor between the two schemes is $\sqrt{N_f/T_f}$. This ratio is equally relevant for the TF scheme: timeframes with very few observations are highly volatile, and a single observation lying close to the initial propagation would produce an extremely small RMS, resulting in an unrealistically high weight. For this reason, the TF scheme enforces a hard lower limit of 10 mas on the uncertainty. This behaviour is further explored with the TF `no cap` scheme in the next section.

From [Figure 6.5](#) it is apparent that many files exhibit a ratio close or equal to 1, particularly those from the 689 observatory. As a consequence, all weighting schemes assign low uncertainties (high weights) to these files, making it difficult to reliably model any correlated noise or systematic biases. This is not to say that these observations are inaccurate, but rather that their accuracy is hard to assess at the timeframe level.

In conclusion, files with a high N_f/T_f ratio benefit most from deweighting strategies, as the per-timeframe RMS is more representative of the true statistical noise and any correlated noise or biases can be more reliably detected. For files with a ratio close to 1, neither scheme can accurately capture whether systematic biases or correlated noise are present. Nevertheless, ID $\nu 2$ exhibits the most coherent formal errors overall, suggesting that these files most likely do not contain correlated noise or biases large enough to be detected at the precision of this study. Therefore, the most significant improvement from ID to ID $\nu 2$ stems from the more appropriate deweighting strategy applied to files with high N_f/T_f ratios, where such effects can be reliably identified and corrected.




Figure C.1: Per File weight uncertainty (defined in [chapter 3](#)) in RA (top) and Dec (bottom) for base weighting schemes: ID, ID v2, and TF

C.2.2. Per Timeframe Analysis

The difference between the three main weighting schemes can be further examined through the per-timeframe uncertainty of individual files. A good example is the 874 nm0004 observation file, which contains roughly 700 observations. The per-timeframe uncertainty for this file is presented in [Figure C.2](#), where the uncertainty of the ID, ID v2 and TF weighting schemes is shown as connected dots, and the number of observations per timeframe is indicated by semi-transparent gray bars on the right y-axis. Both ID and ID v2 appear as flat lines, as their weights are applied at the file level, while TF assigns weights per timeframe. There are significant differences between timeframes: for some, TF assigns much larger uncertainties (lower weights) than the other two schemes, while for most it assigns uncertainties closer to the more aggressive ID scheme, sometimes even lower.

This behaviour cannot be directly linked to the number of observations per timeframe alone. Timeframes with a relatively low number of observations exhibit highly varying TF uncertainties, as the RMS with respect to the initial propagation is not statistically representative for small samples, reflecting the volatility of the estimate. It is difficult to assess what sample size is sufficient to be statistically representative, as this depends on both the data gathering process and the peculiarity of the initial propagation. A possible test would be to run the estimation with slightly different initial conditions, from which the residuals are computed, and to assess whether the resulting weights, solutions and formal errors remain consistent. This has not been done in this study.




Figure C.2: Per Timeframe Uncertainty for observation data file 874 nm0004 in RA (top) and Dec (bottom) for base weighting schemes: ID, ID v2, and TF. Grey bars indicate the number of observations per file (right axis)

Per-file RMS update (initial – final) in RA (top) and Dec (bottom) for Group 1: base weighting schemes ID, ID v2, and TF. Grey bars indicate the number of observations per file (right axis).

C.2.3. RMS Update Per File

The RMS difference between the initial and final propagation for each estimation under the different weighting schemes is presented in [Figure C.3](#) and [Figure C.4](#) for the base weighting schemes (ID, ID v2, and TF) and the hybrid schemes (*hybrid G+v2* and *hybrid A+v2*), respectively.

The updates are overall negligible, amounting to at most a few milliarcseconds. The exception is file 83 nm0083, which contains a single observation near the end of the estimation arc in 1963 and therefore has little effect on the overall estimation performance.

The ID weighting scheme produces the largest per-file changes, which is expected given that its weights differ significantly from those of the other schemes. Notably, it is not only the magnitude of the changes that differs but also the sign for several files, compared to the other weighting schemes. This is most clearly observed for the problematic 83 nm0083 file.

The hybrid schemes *hybrid G+v2* and *hybrid A+v2* perform similarly to their constituent base schemes. Overall, the updates are on the order of a single milliarcsecond, and their magnitude falls between those of ID v2 and TF. For *hybrid A+v2*, the update is skewed more toward TF, as the arithmetic mean of the two weight sets is dominated by the TF weights, which are on average several orders of magnitude larger. An interesting outlier is file 188 nm0016, which exhibits a larger update under the hybrid schemes than under either TF or ID v2 individually.

In summary, all weighting schemes produce per-file RMS updates that are well below the observation noise level, confirming that the estimator converges consistently regardless of the weighting strategy. The ID scheme stands out as the most sensitive to individual files, while ID v2, TF, and the hybrid schemes all exhibit stable, small updates across the observation set. This further supports the selection of ID v2 as the preferred scheme.




Figure C.3: Per-file RMS update (initial - final) in RA (top) and Dec (bottom) for base weighting schemes ID, ID v2, and TF. Grey bars indicate the number of observations per file (right axis).




Figure C.4: Per-file RMS update (initial – final) in RA (top) and Dec (bottom) for hybrid weighting schemes hybrid G+v2 and hybrid A+v2, compared against ID v2 and TF. Layout as in [Figure C.3](#).

TODO: Maybe remove per file uncertainties. As the full weights (uncertainties) will be present.

TODO: Needs a proper discussion and conclusion about the results.

C.3. Analysis WA-IAU

This section presents the weighting scheme analysis for the WA-IAU initialisation, which uses Triton’s initial state from the NEP097 kernel and the IAU 2015 pole model, estimating only the initial state of Triton. The analysis follows the same structure as the WA-FIT-PL analysis in [chapter 6](#). The weights are derived from the residuals of the WA-IAU initial propagation and therefore differ from those used in WA-FIT-PL.

C.3.1. RSW Statistics

The RSW RMS difference, formal error RMS, and their ratio for all weighting schemes under the WA-IAU initialisation are presented in [Figure C.5](#). The overall RMS differences with NEP097 are substantially larger than those of WA-FIT-PL, which is expected given that the WA-IAU initialisation does not include the pole model correction that captures the dominant cross-track mismatch. The RMS differences in R range from approximately 80 to 185 km, in S from 360 to 460 km, and in W from 750 to 1 050 km.

In terms of relative ranking, TF no cap again achieves the lowest RMS differences, while ID exhibits the largest, consistent with the WA-FIT-PL findings. However, the ratios are considerably worse across all schemes compared to WA-FIT-PL. The ID v2 scheme produces the lowest ratios (approximately 4 to 8 depending on direction), while ID and TF no cap exhibit ratios exceeding 20 in some directions. This confirms that without the pole model correction,

no weighting scheme can produce formal errors that are consistent with the actual solution accuracy. Nevertheless, ID v2 remains the most balanced scheme in terms of the ratio metric, consistent with the conclusion from WA-FIT-PL.

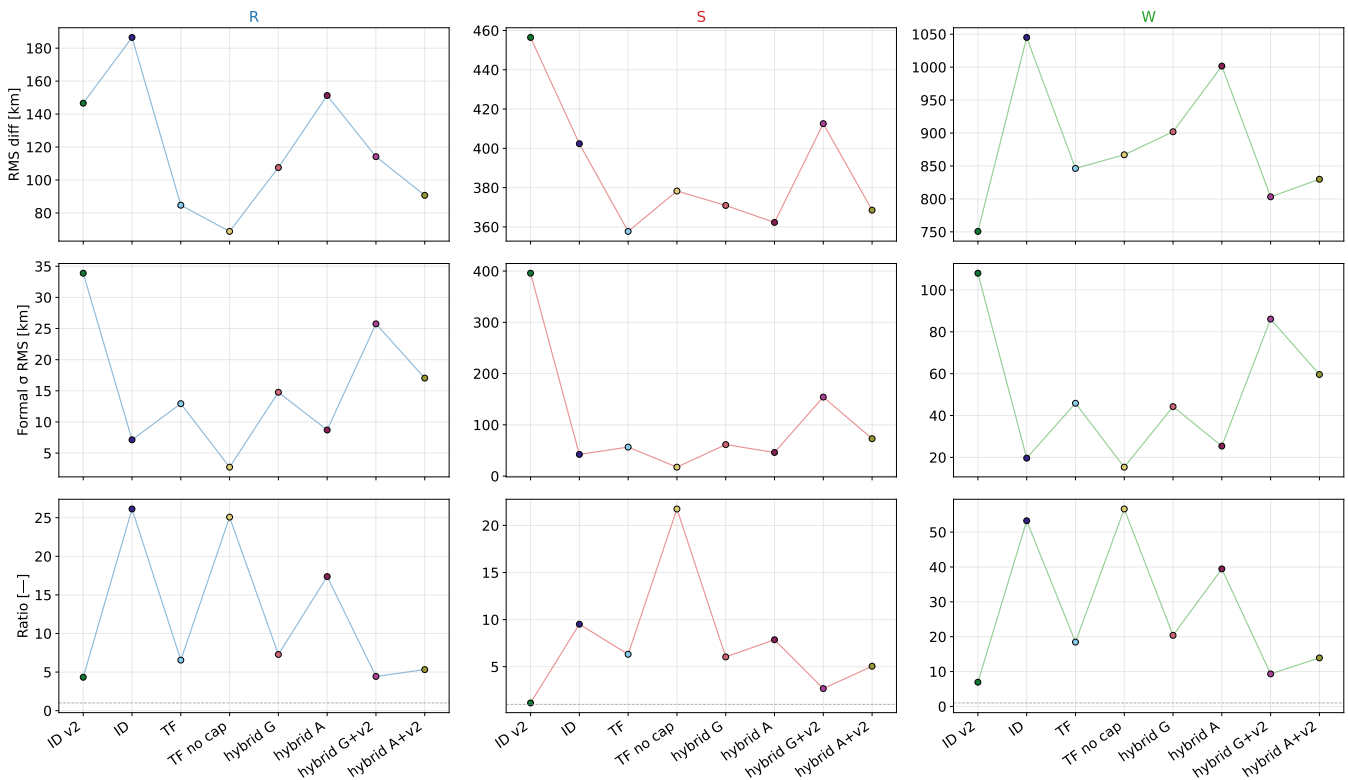


Figure C.5: RMS difference with NEP097, formal error RMS, and their ratio per RSW direction for all weighting schemes under the WA-IAU initialisation.

C.3.2. RSW Timeseries Analysis

Figure C.6 and Figure C.8 present the RSW position differences with respect to NEP097 for the base and hybrid weighting schemes respectively, while Figure C.7 and Figure C.9 show the corresponding formal errors.

Base Weighting Scheme Analysis

The RSW differences under WA-IAU are qualitatively similar to those of WA-FIT-PL but larger in magnitude, particularly in the along-track and cross-track directions. The ID scheme produces the largest RMS difference (655 km), while ID v2 (514 km) and TF (533 km) perform comparably. Unlike in WA-FIT-PL, where ID achieved the lowest cross-track differences, here all three base schemes produce similar cross-track envelopes, with peak values exceeding 2000 km toward 1963.

The formal errors show the same hierarchy as in WA-FIT-PL: ID v2 produces the largest formal errors (σ_R up to 37.7 km, σ_W up to 110 km), followed by TF (σ_R up to 15.4 km, σ_W up to 47 km), and ID (σ_R up to 8.1 km, σ_W up to 21 km). The gap between the formal errors and the actual differences is much larger than in WA-FIT-PL for all schemes, reflecting the absence of the pole model correction.

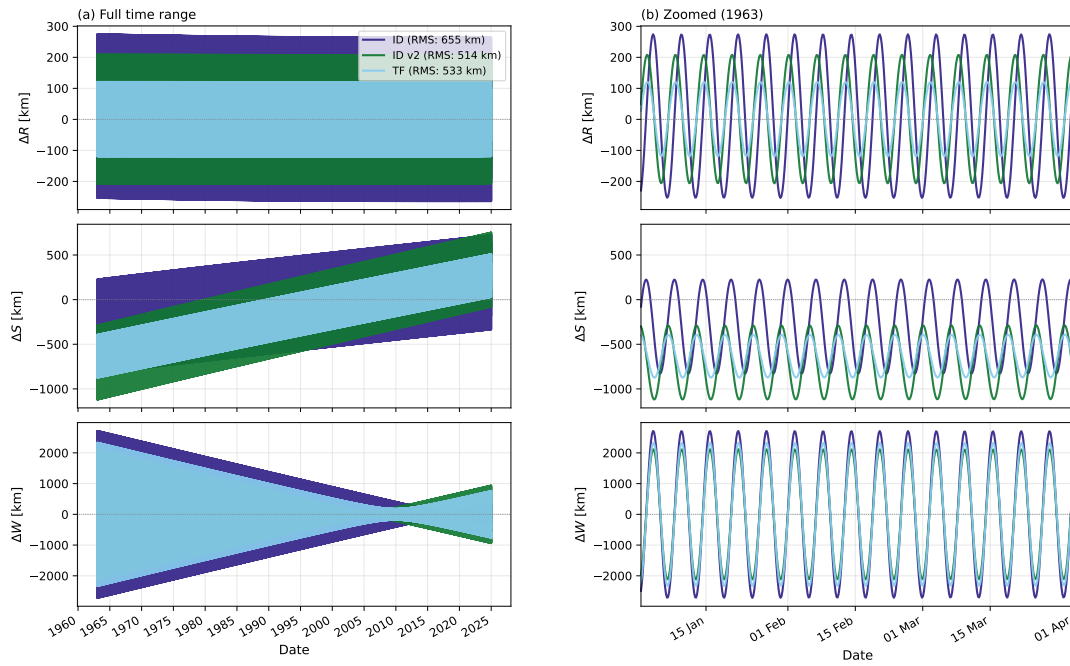


Figure C.6: RSW position difference with respect to NEP097 for base weighting schemes ID, ID v2, and TF under the WA-IAU initialisation. Left panels show the full time range (1960–2025); right panels show a zoomed view in early 1963. From top to bottom: radial (R), along-track (S), and cross-track (W) components.




Figure C.7: RSW formal errors for base weighting schemes ID, ID v2, and TF under the WA-IAU initialisation. Layout as in Figure C.6. From top to bottom: σ_R , σ_S , and σ_W .

Hybrid Weighting Scheme Analysis

The hybrid schemes `hybrid G+v2` and `hybrid A+v2` produce RMS differences (514 to 533 km) that are comparable to their constituent base schemes, with ID v2 achieving the lowest value. The spread between schemes is narrower than in WA-FIT-PL, suggesting that the weighting scheme has less influence on the solution when the dominant dynamical mismatch (the pole model) is not corrected.

The formal error hierarchy mirrors that of WA-FIT-PL: ID v_2 produces the largest formal errors, the hybrid schemes fall in between, and TF produces the smallest. The same conclusion applies: hybrid G+v2 produces marginally smaller formal errors than ID v_2 but inherits the instability of the TF scheme.

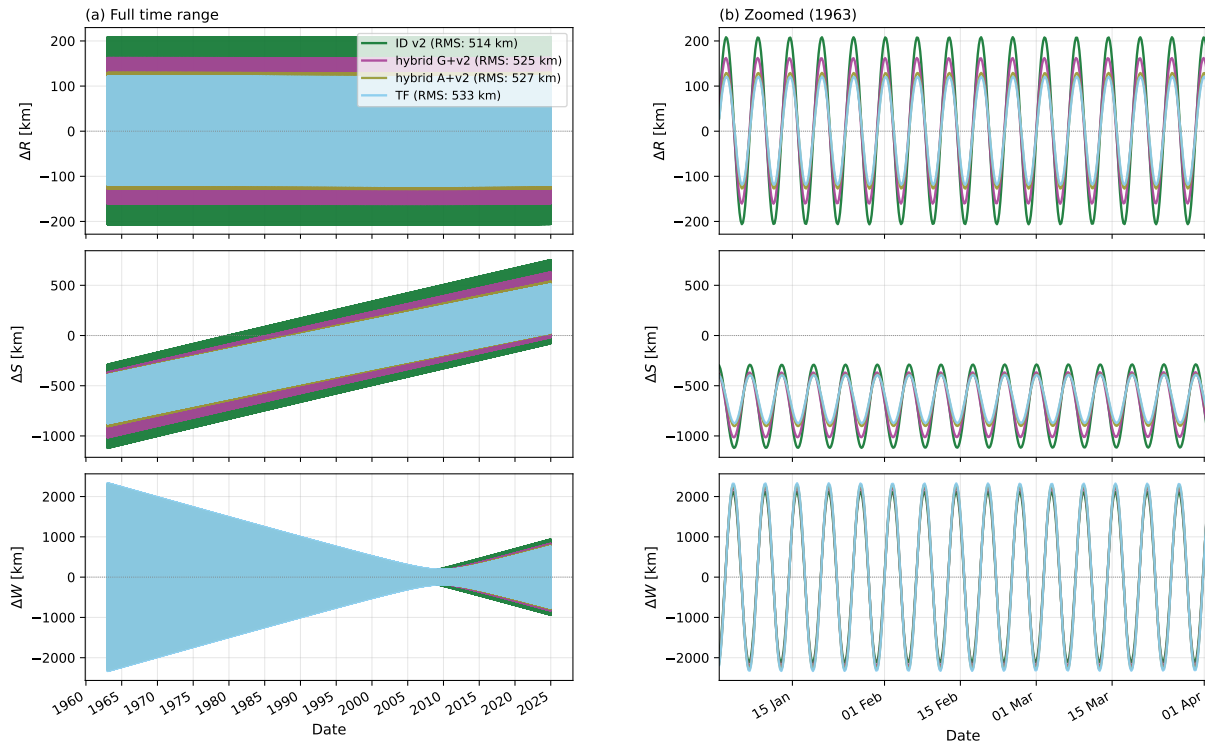


Figure C.8: RSW position difference with respect to NEP097 for hybrid weighting schemes hybrid G+v2, hybrid A+v2, and TF, compared against ID v_2 , under the WA-IAU initialisation. Layout as in [Figure C.6](#).




Figure C.9: RSW formal errors for hybrid weighting schemes hybrid G+v2, hybrid A+v2, and TF, compared against ID v2, under the WA-IAU initialisation. Layout as in Figure C.6.

RSW Timeseries Conclusion

The RSW timeseries analysis under WA-IAU confirms the findings of WA-FIT-PL: ID v2 consistently produces the most conservative formal errors and the most favourable ratio between formal errors and actual differences with NEP097. The absolute ratios are worse than in WA-FIT-PL across all schemes, which is expected given that the dominant dynamical mismatch is not corrected in the WA-IAU initialisation. Importantly, the relative ranking of the weighting schemes is preserved across both initialisations, reinforcing the robustness of ID v2 as the preferred weighting strategy.

C.3.3. Per-File Uncertainty

The per-file uncertainties for the three base weighting schemes under the WA-IAU initialisation are presented in Figure C.10. The overall pattern is consistent with the WA-FIT-PL analysis: ID assigns the lowest uncertainties (highest weights) across nearly all files, while ID v2 and TF assign higher uncertainties. The absolute uncertainty values differ from WA-FIT-PL because the weights are derived from the residuals of a different initial propagation, but the relative ordering between schemes is preserved.

A notable difference compared to WA-FIT-PL is that the ID v2 uncertainties are generally larger than in the WA-FIT-PL case, particularly for files with high N_f/T_f ratios. This is consistent with the larger initial residuals of the WA-IAU propagation, which lacks the pole model correction. The N_f/T_f ratios themselves (Figure C.11) are identical to those in WA-FIT-PL, as they depend only on the observational data structure and not on the initialisation.




Figure C.10: Per-file weight uncertainty in RA (top) and Dec (bottom) for base weighting schemes ID, ID v2, and TF under the WA-IAU initialisation.




Figure C.11: Number of observations over number of timeframes ratio (top) and number of observations (bottom) per file under the WA-IAU initialisation.

C.3.4. Per-Timeframe Analysis

The per-timeframe uncertainty for the 874 nm0004 file under the WA-IAU initialisation is presented in [Figure C.12](#). As in the WA-FIT-PL case, ID and ID v2 appear as flat lines (file-level weights), while TF varies substantially across timeframes. The TF uncertainties are notably more volatile than in WA-FIT-PL, with peaks exceeding 2000 mas in both RA and Dec for several timeframes (e.g. timeframes 12 to 15, 31 to 33). This increased volatility is attributed to the larger initial residuals of the WA-IAU propagation, which amplify the sensitivity of the per-timeframe RMS to individual observations. The ID v2 uncertainty is approximately 850 mas in both coordinates, substantially higher than the ID value of approximately 200 mas, reflecting the per-timeframe deweighting applied by ID v2.




Figure C.12: Per-timeframe uncertainty for observation file 874 nm0004 in RA (top) and Dec (bottom) for base weighting schemes ID, ID v2, and TF under the WA-IAU initialisation. Grey bars indicate the number of observations per timeframe (right axis).

C.3.5. RMS Update Per File

The per-file RMS updates for the WA-IAU initialisation are presented in [Figure C.13](#) and [Figure C.14](#). In contrast to the WA-FIT-PL analysis, where updates were on the order of a few milliarcseconds, the WA-IAU updates are significantly larger, reaching up to 120 mas in RA and 90 mas in Dec. This is expected, as the WA-IAU initialisation uses the uncorrected IAU pole model, resulting in larger initial residuals that the estimator must correct through the initial state adjustment alone.

All three base schemes produce broadly similar update patterns, with the largest updates concentrated on a subset of files: 188 nm0083, 286 nm0084, 327 nm0019, and the 689 observatory files (nm0007 through nm0010). These are files with relatively few observations per timeframe, where the initial residuals are large and the estimator applies the most significant corrections. The ID scheme produces slightly different updates for some files (e.g. 874 nm0003), but the overall pattern is consistent across all schemes.

The hybrid schemes ([Figure C.14](#)) produce updates that are virtually indistinguishable from their base scheme counterparts, with TF showing the largest updates for a few files. This is consistent with the WA-FIT-PL findings.

The large per-file updates under WA-IAU confirm that the estimator is working harder to compensate for the uncorrected pole model mismatch, absorbing the systematic signal through the initial state. This further motivates the use of the WA-FIT-PL initialisation, where the pole model correction has already been applied and the per-file updates are correspondingly smaller.




Figure C.13: Per-file RMS update (initial - final) in RA (top) and Dec (bottom) for base weighting schemes ID, ID v2, and TF under the WA-IAU initialisation. Grey bars indicate the number of observations per file (right axis).




Figure C.14: Per-file RMS update (initial – final) in RA (top) and Dec (bottom) for hybrid weighting schemes hybrid G+v2 and hybrid A+v2, compared against ID v2 and TF, under the WA-IAU initialisation. Layout as in Figure C.13.