



MULTIPLE TARGET TRACKING FOR SPACE OBJECT CATALOGUING

ON SEQUENTIAL ORBIT DETERMINATION AND DATA ASSOCIATION
USING RANDOM FINITE SETS

M.Sc. Thesis by
Simon Van Hulle

MULTIPLE TARGET TRACKING FOR SPACE OBJECT CATALOGUING

ON SEQUENTIAL ORBIT DETERMINATION AND DATA ASSOCIATION
USING RANDOM FINITE SETS

by

Simon Van Hulle

Thesis submitted in partial fulfillment
of the requirements for the degree

Master of Science

Student Number: 4998006

Project Duration: November 2024 - July 2025

Thesis Committee:	Dr. ir. Dominic Dirkx	(chair)
	Dr. ir. Bart Root	(independent examiner)
	Dr. Steve Gehly	(supervisor TU Delft)
	Dr. ir. David Gondelach	(supervisor Vyoma)

Version: 2025-07-10

To be defended: 2025-07-24 at 10:00



PREFACE

This work marks the end of my studies at the TU Delft faculty of Aerospace Engineering. Throughout these years, I was fortunate to be surrounded by people whose friendship, love, insight, and support made my time as a student and intern both educational and deeply enjoyable.

First of all, I could not be more grateful to Steve Gehly and David Gondelach – the outstanding mentors who introduced me to the subtleties of scientific research, and guided me through the maze of my often incoherent arguments and spurious ideas. For the hours of your time, the illuminating discussions, diligent feedback, patient explanations, and so much more: thank you.

I thank my colleagues at Vyoma, who made me feel at home in a new country, for their refreshing points of view, our absurdly entertaining lunch conversations, and the truly stimulating environment they all continue to create.

I am thankful to the friends I made in Antwerp, Delft, and Munich, to my wonderful parents, brother, and close-knit family. Thank you for being my examples and inspiration – I owe you everything.

Lastly, let me express my appreciation to you, dear reader, for engaging – whether in depth or in passing – with the outcome of my work.

Simon Van Hulle
Munich, June 2025

ABSTRACT

The growing population of satellites, rocket bodies, debris and other objects in orbit around the Earth calls for robust space cataloguing methods that can characterise this environment, to support continued and safe space operations.

Traditional approaches maintain such a catalogue by first combining consecutive measurements into *tracklets*, then associating these tracklets definitively to known objects or to other tracklets and finally updating the orbit estimates accordingly. In contrast, multi-object tracking (MOT) methods consider the weighted contributions of multiple measurement association hypotheses, but typically discard the pre-formed tracklets and do not draw deterministic conclusions as to which measurements belong to which objects.

This work explores the properties of the relatively new labelled multi-Bernoulli (LMB) filter and introduces a way to extract tracklet association probabilities directly from the filtering recursion, enabling more specific single-object post-processing. In addition, a modified filter is proposed that accepts complete tracklets instead of individual measurements, taking full advantage of all prior information.

In low-clutter environments, the tracklet filter is shown to be less sensitive to field-of-view edge effects, allows for more efficient propagation in periods without detection, and reduces the overall association complexity compared to the original LMB filter.

After integrating two types of initial orbit determination, it is demonstrated that this tracklet filter can discover and maintain state estimates for objects in low Earth orbit (LEO) and geosynchronous Earth orbit (GEO), using sparse optical measurements from both ground-based and space-based observers with various pointing strategies.

All methods are implemented in BASIL, a general-purpose Bayesian estimation library developed for this work to support flexible, modular, and scalable multi-object tracking in astrodynamic contexts.

CONTENTS

Preface	i
Abstract	ii
Nomenclature	xi
Acronyms	xv
1 Introduction	1
2 Fundamental background	4
2.1 Probability density	4
2.1.1 Relevant probability distributions	4
2.1.2 Gaussian mixtures	6
2.1.3 Distance and information divergence	9
2.2 Single-object Bayesian estimation	10
2.2.1 State space representation	10
2.2.2 Single-object Bayesian filter	10
2.2.3 Conjugate prior	11
2.2.4 Gaussian model simplification	12
3 Multi-object tracking with random finite sets	17
3.1 Multi-object representation	17
3.2 Random finite sets	17
3.2.1 Definitions and finite set statistics	17
3.2.2 Labelled random finite set	20
3.2.3 Relevant RFS models	20
3.3 Multi-object Bayesian filtering	22
3.3.1 Multi-object transition model	23
3.3.2 Multi-object measurement model	23
3.3.3 Optimal assignment	25
3.3.4 LMB filter implementation	28
3.3.5 Overview of the LMB parameters	38
3.3.6 Multi-object divergence metrics	40
3.4 Implementations in BASIL	41
4 Extracting associations	43
4.1 Time-marginal association probability	43
4.2 Accumulated (tracklet) association probability	45
5 Multi-object tracking in space	48
5.1 Orbital motion model	48
5.1.1 Reference frames	48
5.1.2 Equations of motion	50
5.1.3 Evolution of uncertainty	51
5.1.4 State noise compensation	52
5.1.5 Survival model	52
5.2 Angles-only optical observation model	53
5.2.1 Space-based observers	53
5.2.2 Probability of detection	55

5.2.3 Clutter model	56
5.2.4 Optical tracklets	56
5.3 Initial orbit determination - space birth models	56
5.3.1 Gooding and batch least squares	57
5.3.2 Attributable and constrained admissible region	58
5.3.3 Adaptive birth with tracklets	62
6 The tracklet LMB filter	64
6.1 Tracklets as measurements	64
6.1.1 Tracklet detection probability	66
6.1.2 Tracklet log likelihood	67
6.1.3 Tracklet clutter model	67
6.2 Tracklet grouping	69
7 Simulations	71
7.1 Tracklet to tracklet association in GEO	71
7.1.1 Setup	71
7.1.2 Results	72
7.1.3 Summary	75
7.2 Closely spaced GEO satellites with space-based observer	75
7.2.1 Setup	75
7.2.2 Single-measurement LMB filter with prior knowledge	77
7.2.3 Tracklet LMB filter with prior knowledge	81
7.2.4 Tracklet LMB filter without prior knowledge	83
7.2.5 Evaluating association performance	87
7.2.6 Summary	88
7.3 Discovering GEO population from LEO	89
7.3.1 Setup	89
7.3.2 Birth model	90
7.3.3 Filter results	91
7.3.4 Summary	95
7.4 Polar payload separation	96
7.4.1 Setup	96
7.4.2 CAR birth model	96
7.4.3 Gaussian birth from rocket upper stage	101
7.4.4 Summary	103
8 Conclusion	104
8.1 Revisiting the research questions	105
8.2 Limitations and future work	106
References	109
List of Appendices	116
A Literature review	117
A.1 Single-object Bayesian filtering	117
A.2 Multiple object tracking	118
A.2.1 Tracking after association	119
A.2.2 Tracking while associating	119
A.3 MOT for SSA	122

B Project plan	124
C Testing the LMB in two dimensions	126
C.1 Reuter and Vo LMB EKF	126
D Constrained admissible region IOD	129
D.1 Constraining the admissible region	130
D.2 Attributable prediction	131
D.3 CAR Examples	132
D.3.1 Ground-based observer	132
D.3.2 Space-based observer	134
E Auxiliary computations	135
E.1 Covariance ellipsoids and the χ^2 distribution	135
E.2 LogSumExp normalisation	135
F Association matrices	136
G Simulated object data	140

LIST OF FIGURES

Figure 2.1	Simple 2D illustration of Gaussian mixture reduction methods	7
Figure 2.2	1D linear example – scenario	11
Figure 2.3	1D linear example – Gaussian state distribution	11
Figure 2.4	1D linear example – Kalman filter	13
Figure 3.1	Why multi-object densities are hard to compare	19
Figure 3.2	Example assignment score matrix	27
Figure 3.3	Example of grouping and gating	28
Figure 3.4	LMB filter schematic overview	29
Figure 3.5	Conversion from LMB to GLMB; example with 3 Bernoullis	32
Figure 3.6	GLMB update schematic representation	34
Figure 3.7	Computing existence probabilities for LMB approximation	34
Figure 3.8	1D linear example – initial LMB	35
Figure 3.9	1D linear example – predicted LMB	35
Figure 3.10	1D linear example – prior GLMB and predicted measurements	36
Figure 3.11	1D linear example – prior GLMB and predicted measurements	36
Figure 3.12	1D linear example – posterior LMB	37
Figure 3.13	1D linear example – estimated multi-object state	37
Figure 3.14	Overview of core MOT components in BASIL	42
Figure 4.1	Example label association result	44
Figure 5.1	Earth-centred reference frames	49
Figure 5.2	RTN reference frame	50
Figure 5.3	Uncertainty evolution in LEO under Keplerian dynamics.	52
Figure 5.4	Light-time effect for space-based observer	54
Figure 5.5	Aberration effect for space-based observer	54
Figure 5.6	Detection probability based on field of view	56
Figure 5.7	Example of initial orbit determination using Gooding’s method	58
Figure 5.8	Identical attributable for very different orbits	59
Figure 5.9	Example CAR with ground-based observer	60
Figure 5.10	Example ground-based CAR – accepted orbits	60
Figure 5.11	Example ground-based CAR – distribution of semi-major axis and eccentricity	61
Figure 5.12	Example ground-based CAR – range-marginal probability density function	61
Figure 5.13	Example ground-based CAR – Gaussian mixture components	62
Figure 5.14	Example ground-based CAR – Gaussian mixture density	62
Figure 6.1	Multi-scan detection probability with field of view	66
Figure 6.2	Tracklet grouping	69
Figure 7.1	Observation geometry for GEO objects from ground-based observer	72
Figure 7.2	OSPA errors for ground-based GEO	73
Figure 7.3	Position errors for ground-based GEO	73
Figure 7.4	RTN filter estimation errors and uncertainty bounds for object 26038	74
Figure 7.5	Confusion matrix for GEO objects with ground-based observer	75
Figure 7.6	Overview of components in the LMB filter	75

Figure 7.7	Distribution of tracklet duration	76
Figure 7.8	Observation geometry of GEO satellites observed from LEO	76
Figure 7.9	Predicted uncertainty	78
Figure 7.10	Euclidean distance between simulated objects	78
Figure 7.11	OSPA for classical LMB filter with 7 closely spaced GEO objects	78
Figure 7.12	Confusion matrix for classical LMB filter with time-marginal probabilities	79
Figure 7.13	Confusion matrix for classical LMB filter with accumulated probabilities	79
Figure 7.14	Detection probabilities for Arabsat 6B	80
Figure 7.15	single-measurement detection probabilities for Arabsat 7B	81
Figure 7.16	OSPA for tracklet LMB filter with 7 closely spaced GEO objects	81
Figure 7.17	Confusion matrix for closely spaced GEO objects with tracklet LMB filter	82
Figure 7.18	Post-fit estimation errors with tracklet filter	82
Figure 7.19	Post-fit Mahalanobis distance with tracklet filter	82
Figure 7.20	Confusion matrix for tracklet LMB filter without prior knowledge	84
Figure 7.21	Confusion matrix for tracklet LMBfilter, rearranged by final association	84
Figure 7.22	OSPA for tracklet LMB filter without any initial information	85
Figure 7.23	Tracklet detection probabilities with tracklet LMB filter	85
Figure 7.24	Predicted densities for Arabsat 4B and 6B	85
Figure 7.25	CAR birth GM approximation with space-based observer	87
Figure 7.26	CAR birth range-marginal PDF with space-based observer	87
Figure 7.27	Confusion matrix for tracklet LMB filter with more GM components	87
Figure 7.28	Distribution of a and e of considered GEO objects	89
Figure 7.29	Distribution of longitude for considered GEO objects	89
Figure 7.30	Observation strategy to discover GEO objects with polar observer	90
Figure 7.31	Shape of GEO tracklets with polar space-based observer	91
Figure 7.32	Filter component counts, observing 171 objects in GEO	92
Figure 7.33	OSPA errors for GEO discovery case	92
Figure 7.34	Cardinality, Mahalanobis distance and errors over time for GEO discovery	93
Figure 7.35	True distribution of inclination and RAAN	95
Figure 7.36	Estimated distribution of inclination and RAAN	95
Figure 7.37	Final RTN errors for 200 objects in GEO	95
Figure 7.38	Observer attitude profile for polar payload separation	97
Figure 7.39	Tracklet distribution for polar targets from polar observer	97
Figure 7.40	Constrained admissible region for polar observer and target	98
Figure 7.41	GM approximation of space-based CAR in LEO	98
Figure 7.42	Distribution of RAAN and inclination resulting from a CAR in LEO	98
Figure 7.43	Association summary for polar targets from polar observer	99
Figure 7.44	Rearranged association summary for polar targets from polar observer	99
Figure 7.45	Association counts per object over time	100
Figure 7.46	Assignment summary for payload separation case without birth model	101
Figure 7.47	Assignment summary for payload separation case without birth model (reordered)	101
Figure 7.48	Cumulative association counts per object without birth model	102
Figure 7.49	Detection and existence probabilities with insufficient GM components	103
Figure B.1	Work Breakdown Structure	124
Figure B.2	Preliminary activity planning	125
Figure C.1	Vo test case ground truth and measurements	127
Figure C.2	2D test case LMB estimates	127

Figure C.3	2D test case LMB cardinality	128
Figure C.4	2D test case LMB OSPA metric	128
Figure D.1	Constrained admissible region for Gehly case (BASIL)	133
Figure D.2	Constrained dmissible region for Gehly test case (METIS)	133
Figure D.3	Range-marginal PDF for Gehly test case (BASIL)	133
Figure D.4	GM approximation for Gehly test case (BASIL)	133
Figure D.5	Range-marginal PDF for Gehly test case (METIS)	134
Figure D.6	GM approximation for Gehly test case (METIS)	134
Figure D.7	Example CAR with space-based observer	134
Figure D.8	Example space-based CAR – Gaussian mixture approximation components	134
Figure D.9	Example space-based CAR – range-marginal probability density function	134
Figure F.1	Association probabilities for GEO with ground-based observer	136
Figure F.2	Association probabilities for GEO with space-based observer – low GM cap	137
Figure F.3	Association probabilities for GEO with space-based observer – no GM cap	138
Figure F.4	Association probabilities for payload separation with space-based observer	139

LIST OF TABLES

Table 1.1	Overview of contributions	3
Table 3.1	Summary of LMB parameters and their meaning	38
Table 7.1	Overview of simulation cases	71
Table 7.2	Common parameters for the case with closely spaced GEO objects	77
Table 7.3	Admissible region settings for closely spaced GEO objects	83
Table 7.4	Overview of association performance for GEO objects from LEO observer	88
Table 7.5	Common parameters for case with 200 GEO objects	90
Table 7.6	OSPA distance for discovering GEO population	93
Table 7.7	GEO object discovery results from polar LEO observer	94
Table 7.8	Association quality for GEO discovery case	95
Table 7.9	Final OSPA for payload separation case with CAR birth	99
Table 7.10	Association count for all objects	100
Table 7.11	Association counts summary for payload separation without birth model	102
Table C.1	LMB parameters for 2D test case	126
Table C.2	Constant birth model parameters for 2D test case	126
Table E.1	Cumulative density of the χ distribution	135
Table G.1	Orbital elements for Pirovano test case objects	140
Table G.2	TLE data for closely spaced GEO objects, obtained from SpaceTrack	140
Table G.3	Orbital elements for payload separation test case	140

LIST OF ALGORITHMS

Algorithm 2.1	Gaussian mixture merging	8
Algorithm 2.2	Extended Kalman Filter	13
Algorithm 2.3	Unscented Kalman Filter	15
Algorithm 2.4	Gaussian mixture prediction	15
Algorithm 2.5	Gaussian mixture single-measurement Bayes update	16
Algorithm 3.1	Sampling from a Poisson RFS	21
Algorithm 3.2	Sampling measurements from the standard measurement model	25
Algorithm 3.3	Converting LMB to δ GLMB using shortest path formulation	30
Algorithm 4.1	Marginal association probability mapping at time t_k	45
Algorithm 4.2	Accumulated association probability mapping up to time t_k	47
Algorithm 6.1	Simple tracklet grouping	70

NOMENCLATURE

SYMBOL	DESCRIPTION
General	
\mathfrak{Y}	Any finite-dimensional vector space
$Y \subset \mathfrak{Y}$	Subset of \mathfrak{Y}
$\mathcal{F}(\mathfrak{Y})$	Class of all finite subsets of \mathfrak{Y}
$\mathbf{y} \in \mathfrak{Y}$	Vector
$\boldsymbol{\mu}$	Mean vector
\mathbf{P}	Covariance matrix
$\mathcal{D}(f)$	Domain of function f
Sets and spaces	
\mathbb{N}	Natural numbers $\{1, 2, \dots\}$
\mathbb{X}	Finite-dimension state space
\mathbb{Z}	Finite-dimension measurement space
\mathbb{T}	Discrete time window
\mathbb{L}	State label space
\mathbb{M}	Measurement label space
$\mathbb{B}_k \subseteq \mathbb{L}$	Set of labels for tracks initialised (born) at t_k
Π_k	Set of permutations on $\{1, \dots, k\}$ for any $k \in \mathbb{N}$
Distributions	
\mathcal{N}	Multivariate normal (Gaussian) distribution
\mathcal{U}	Unifrom distribution
\mathcal{G}	Gaussian mixture distribution
χ_n^2	Chi-squared distribution with n degrees of freedom
d_M	Mahalanobis distance
Single-object	
$\mathbf{x} \in \mathbb{X}$	Single-object state vector
$\ell \in \mathbb{L}$	Target label
$\hat{\mathbf{x}} \triangleq (\mathbf{x}, \ell)$	Labelled single-object state
$\varphi_k(\mathbf{x}_{k-1}, \boldsymbol{\nu}_k)$	(Non-linear) evolution of \mathbf{x}_{k-1} to t_k , under the noise process $\boldsymbol{\nu}_k$
$h_k(\mathbf{x}_k, \boldsymbol{\varepsilon}_k)$	(non-linear) measurement mapping for \mathbf{x}_{k-1} , under the noise process $\boldsymbol{\varepsilon}_k$
$f_k(\cdot \mathbf{x}_{k-1})$	Markov transition density to time t_k given \mathbf{x}_{k-1}
$g_k(\mathbf{z}_k \mathbf{x}_k)$	Likelihood of observing \mathbf{z}_k given \mathbf{x}_k
$p_k(\mathbf{x}_k)$	$\triangleq p_{k k}(\mathbf{x}_k)$. Filtering density at \mathbf{x}_k
\mathbf{S}_k	Innovation covariance matrix at time t_k
\mathbf{H}_k	Measurement Jacobian $\frac{\partial h_k}{\partial \mathbf{x}_k} _{\mathbf{x}_k}$
$\boldsymbol{\Phi}_{k,k-1}$	Error state transition matrix $\frac{\partial \varphi_k}{\partial \mathbf{x}_k} _{\mathbf{x}_k}$
\mathbf{K}_k	Kalman gain matrix

SYMBOL	DESCRIPTION
Multi-object	
$X \subset \mathbb{X}$	Multi-object state
$\dot{X} \subset \mathbb{L} \times \mathcal{F}(\mathbb{X})$	Labelled Multi-object state
$X_{j:k}$	Multi-object state history on time sequence $\{t_j, \dots, t_k\}$
$\dot{X}_{j:k}$	Multi-object trajectory on the time sequence $\{t_j, \dots, t_k\}$
$T(\ell)$	Set of instants where multi-object trajectory $\dot{X}_{j:k}$ contains label ℓ
$\dot{\mathbf{x}}_{T(\ell)}^{(\ell)}$	Trajectory of label ℓ in multi-object trajectory $\dot{X}_{j:k}$
$\mathbf{f}_k(\cdot \dot{X}_{k-1})$	multi-object transition density given \dot{X}_{k-1}
$\mathbf{g}_k(Z_k \dot{X}_k)$	Likelihood of observing Z_k given \dot{X}_k
$P_{s,k}(\mathbf{x}_{k-1}, \ell)$	Probability of survival to time t_k for state (\mathbf{x}_{k-1}, ℓ)
$P_{D,k}(\mathbf{x})$	Probability of detecting state \mathbf{x} at time t_k
$\kappa_c(\mathbf{z})$	Clutter intensity at \mathbf{z}
λ_c	Clutter rate
Measurements	
$\mathbf{z} \in \mathbb{Z}$	Single measurement vector
$m \in \mathbb{M}$	Measurement label
$\dot{\mathbf{z}} \triangleq (\mathbf{x}, m)$	Labelled measurement
$Z \subset \mathbb{Z}$	Measurement set $Z = C \uplus O_X$
$C \subseteq Z$	Set of clutter measurements
$O_X \subseteq Z$	Set of observations originating from the multi-object state X
$\dot{Z} \subset \mathbb{M} \times \mathcal{F}(\mathbb{Z})$	Labelled measurement set
Random finite sets	
Ψ	Random finite set (RFS) on some hyperspace $\mathcal{F}(\mathfrak{Y})$
$\dot{\Psi}$	Labelled random finite set (LRFS) on some hyperspace $\mathbb{L} \times \mathcal{F}(\mathfrak{Y})$
ρ_Ψ	Cardinality distribution of RFS Ψ
π_Ψ	Multi-object (belief/FISST) density of the RFS Ψ
$\dot{\pi}_{\dot{\Psi}}$	Multi-object density of the LRFS $\dot{\Psi}$
$\dot{\pi}_k(\dot{X}_k)$	Multi-object filtering density at X_k
$w_{\dot{\Psi}}(L)$	Joint existence probability. $\Pr(L \subseteq \mathcal{L}(\dot{\Psi}))$
v_Ψ	PHD or intensity of RFS Ψ
$\dot{v}_{\dot{\Psi}}$	PHD or intensity of LRFS $\dot{\Psi}$
$\dot{p}(\cdot, \ell)$	Attribute density for label ℓ
$r^{(\ell)}$	Probability of existence for label ℓ in an LMB density
$q^{(\ell)}$	$\triangleq 1 - r^{(\ell)}$. Probability of non-existence for label ℓ in an LMB density
$\mathcal{H}_+ \triangleq \mathcal{H}_{k k-1}$	$\triangleq \{(\xi, L)\}$. Set of hypotheses in the predicted GLMB density
\mathcal{H}_k	$\triangleq \{(\xi_\theta, L)\}$. Set of hypotheses in the corrected GLMB density
(ξ, L)	GLMB hypothesis with label set L and association history ξ
$w^{(\xi)}(L)$	Joint existence probability of all labels in L under association history ξ – also referred to the <i>weight</i> of the hypothesis (ξ, L) in a GLMB density

SYMBOL	DESCRIPTION
Functions and operators	
$ Y $	Cardinality of a set Y
$\langle f, g \rangle$	$\triangleq \int f(\mathbf{y})g(\mathbf{y})d\mathbf{y}$. Inner product of functions f and g
$\langle \mathbf{a}, \mathbf{b} \rangle$	$\triangleq \mathbf{a}^T \mathbf{b}$. Inner product of vectors \mathbf{a} and \mathbf{b}
$\sqrt{\mathbf{M}}$	Lower-triangular Cholesky factorisation matrix, such that $\sqrt{\mathbf{M}}\sqrt{\mathbf{M}}^T = \mathbf{M}$
$\mathbf{1}_S(\mathbf{y})$	Indicator function for a set S . $\mathbf{1}_S(\mathbf{y}) = 1$ if $\mathbf{y} \in S$, else 0
$\delta_a[b]$	$\triangleq 1$ if $a = b$ else 0. Kronecker-delta
h^Y	Multi-object exponential. $h^Y \triangleq \prod_{\mathbf{y} \in Y} h(\mathbf{y})$, with $h^\emptyset = 1$
$\mathcal{L}(\overset{\circ}{Y})$	Labels of set $\overset{\circ}{Y}$ with projection $\mathcal{L} : (\mathbf{y}, \ell) \mapsto \ell$
$\mathcal{A}(\overset{\circ}{Y})$	Attributes of set $\overset{\circ}{Y}$ with projection $\mathcal{A} : (\mathbf{y}, \ell) \mapsto x$
$\dot{\mathbf{y}}$	$\triangleq \frac{\partial}{\partial t} \mathbf{y}$. First time-derivative of variable \mathbf{y}
Association probability	
$\alpha_k^{(\ell, \bar{m})}$	Time-marginal association probability between target ℓ and measurement \bar{m}
$\bar{\alpha}_k^{(m)}$	Time-marginal probability that \bar{m} is clutter or unknown a time t_k
$\alpha_{1:k}^{(m)}(\ell)$	Probability that measurement \bar{m} is associated to ℓ , accumulated up to t_k
$\bar{\alpha}_{1:k}^{(m)}$	Probability that \bar{m} is clutter or unknown a time t_k
$\mathcal{M}(\ell)$	set of measurement labels assigned to target ℓ .
Fundamental astrodynamics	
\mathcal{E}	Specific orbital energy
\mathbf{h}	Angular momentum vector
R_E	Equatorial radius of the Earth
a_{GEO}	Semi-major axis for equatorial GEO orbit around a point mass Earth
Orbital state parameters	
\mathbf{r}	$\triangleq (x, y, z)$. Cartesian position
\mathbf{v}	$\triangleq \dot{\mathbf{r}} \triangleq (\dot{x}, \dot{y}, \dot{z})$. Cartesian velocity
a	Semi-major axis
e	Eccentricity
i	Orbital inclination w.r.t. the equatorial plane
Ω	right ascension of the ascending node (RAAN)
ω	Argument of pericentre/perigee
θ	True anomaly
Optical measurements and attributable IOD	
α	right ascension in the topocentric frame parallel to ECI
δ	declination in the topocentric frame parallel to ECI
\mathbf{a}	$\triangleq (\alpha, \delta, \dot{\alpha}, \dot{\delta})$. Attributable measurement
ρ	Slant range vector
Tracklets as measurements	
$\mathfrak{T}_{a,b}^{(m)}$	Tracklet with label \bar{m} and measurements at times between $t_a^{(m)}$ and $t_b^{(m)}$

SYMBOL	DESCRIPTION
\mathfrak{Z}_k	$\triangleq \{\mathfrak{Z}\}$ Tracklet group for filter update at time t_k
$\lambda_c^{\mathfrak{Z}}$	Tracklet clutter rate
$\kappa_c^{\mathfrak{Z}}$	Tracklet clutter intensity
Δt_{valid}	Tracklet validity padding
Sub/superscripts	
$(\cdot)_k$	at time t_k
$(\cdot)_{k k-1}$	filtering prediction at t_k , based on measurements up to t_{k-1}
$(\cdot)_{k k}$	filtering correction at time t_k , based on measurements up to t_k

ACRONYMS

AEGIS	Adaptive Entropy-Based Gaussian-Mixture Information Synthesis
BASIL	Bayesian Association and State Inference Library
BVP	Boundary Value Problem
CAR	Constrained Admissible Region
COE	Classical Orbital Elements
CPHD	Cardinalised Probability Hypothesis Density
CS	Cauchy-Schwarz
ECEF	Earth-Centred Earth-Fixed
ECI	Earth-Centred Inertial
EKF	Extended Kalman Filter
EME2000	Earth Mean Equinox Reference Frame at Julian Epoch J2000
FISST	Finite Set Statistics
FN	False Negative
FOV	Field of View
FP	False Positive
GEO	Geosynchronous Earth Orbit
GLMB	Generalised Labelled Multi-Bernoulli
GM	Gaussian Mixture
GNN	Global Nearest Neighbour
GNSS	Global Navigation Satellite System
GOSPA	Generalised Optimal Sub-Pattern Assignment Metric
HMM	Hidden Markov Model
i.i.d.	Independent and Identically Distributed
IERS2010	International Earth Rotation Service 2010
IOD	Initial Orbit Determination
ITRF	International Terrestrial Reference Frame
JoM	Joint Multi-Object Estimate
JPDA	Joint Probabilistic Data Association
KL	Kullback–Leibler
LEO	Low Earth Orbit
LMB	Labelled Multi-Bernoulli
LOS	Line of Sight
LRFS	Labelled Random Finite Set
LS	Least Squares
M2TA	Measurement-to-Track Association
MaM	Marginal Multi-Object Estimate
MAP	Maximum a-Posteriori
MB	Multi-Bernoulli Random Finite Set
MBM	Multi-Bernoulli Mixture
MEE	Modified Equinoctial Elements
MEO	Medium Earth Orbit
METIS	Multitarget Estimation, Tracking and Information Synthesis
MHT	Multiple Hypothesis Tracking
MO	Multi-Object
MOT	Multi-Object Tracking
OD	Orbit Determination
OMAT	Optimal Mass Transfer Metric
Orekit	ORbits Extrapolation KIT
OSPA	Optimal Sub-Pattern Assignment Metric
PAR	Probabilistic Admissible Region

PDA	Probabilistic Data Association
PDF	Probability Density Function
PF	Particle Filter
PGF	Probability Generating Function
PGFl	Probability Generating Functional
PHD	Probability Hypothesis Density
PLMB	Poisson Labelled Multi-Bernoulli
PMBM	Poisson Multi-Bernoulli Mixture
PMF	Probability Mass Function
PPP	Poisson Point Process
PPV	Positive Predictive Value
RAAN	Right Ascension of the Ascending Node
RFS	Random Finite Set
RMS	Root Mean Square
RSO	Resident Space Object
RTN	Radial, Transverse, Normal
RV	Random Variable
SBO	Space-Based Observer
SMC	Sequential Monte Carlo
SNR	Signal-to-Noise Ratio
SO	Single-Object
SRP	Solar Radiation Pressure
SSA	Space Situational Awareness
SSM	State Space Model
SSN	Space Surveillance Network
TFRM	Telescope Fabra ROA Montsec
TLE	Two-Line Element
TO	Track-Oriented
TOSPA	Temporal Generalisation of Optimal Sub-Pattern Assignment Metric
TP	True Positive
TPR	True Positive Rate
UCT	Uncorrelated Tracklet
UKF	Unscented Kalman Filter
δGLMB	δ -generalised Labelled Multi-Bernoulli

Definitions

object	An existing item with unknown state, which is partially observable by measurements and can be estimated in a tracking process
target	A hypothesised object, as part of a multi-object filter
track	The history of estimated states for a particular target. This is synonymous to the estimated trajectory
tracking	The process of estimating the state trajectory over time for one or more objects of interest, based on measurements
tracklet	A short group of measurements, certainly originating from a single object

INTRODUCTION

Over the last several decades, the number of artificial satellites in Earth’s orbit has increased drastically. Certain orbital regimes are of particular interest for specific applications, leading to rapidly increasing satellite densities in those regions. Low Earth orbit (LEO) is often attractive because of the relatively low energy to orbit. The low latency is useful for communication purposes, and the short orbital period and broad surface coverage make low orbits ideal for imaging and other Earth observation missions. Meanwhile, the geosynchronous Earth orbit (GEO) regime is very suitable for applications like telecommunication, weather monitoring and global navigation satellite system (GNSS) augmentation.

This increased local concentration of objects in space inevitably increases collision risks. In 2009, Iridium 33 and Cosmos 2251 were involved in the first impact between intact satellites [1]. This generated two clouds with thousands of debris fragments, which quickly distributed across their respective orbits and later migrated even further away under the influence of dynamic perturbations. Before that, there had already been several collisions with mission-related debris, in-orbit break-ups, anti-satellite tests and other events that spiked the number of debris objects in space. New debris is still being generated on a daily basis because of normal launch operations, material ageing, leaks in thermal control systems and further fragmentation effects [76]. In 2016, a debris particle with a characteristic length of 1 cm hit a solar panel of Sentinel-1A, causing a change in its orbit and partial power loss [71]. Incidents like this one demonstrate that even tiny pieces of debris can have destructive consequences for operational satellites, underlining the need for improved, up-to-date and accurate understanding of the space environment. This objective is often referred to as space situational awareness (SSA).

According to [ESA Space Debris Office](#) [34], there are almost 1 million resident space objects (RSOs) in Earth’s orbit larger than 1 cm, but only 29 568 of those are actively tracked and catalogued by the space surveillance network (SSN) as of March 2025 [94]. To improve the safety of the space environment, this portion of catalogued objects should increase significantly, which brings about all sorts of challenges. Observing and tracking this many objects in space, usually with very short measurement arcs, dense target clusters, long intervals between detections and imperfect dynamical models leads to very challenging measurement association problems. Since most SSA tracking measurements are angles-only optical streaks or radar observations [20][65] with no direct way to infer the source object, increasingly sophisticated methods are required to deduce their originators.

Traditional approaches start by collecting observations into *tracklets* – short sequences of closely spaced measurements from the same tracked object. Since single passes through the observer field of view (FOV) are usually orders of magnitude shorter than a single orbital period, this tracklet formation process is nearly linear in 2 or 3 dimensions (for optical and radar, respectively) and does not require full orbital dynamics. After this, various track-to-orbit association methods are used to determine to which object in the catalogue each tracklet belongs. If this fails, track-to-track methods (often based on initial and boundary value formulations [87][98][116][117]) can also be employed. Every time an object is assigned a tracklet, its estimate is updated in the catalogue and the process continues. However, in case of densely packed target orbits, very short

tracklets or outdated catalogue states with large uncertainties, these algorithms can result in ambiguous associations, where a single tracklet might belong to multiple objects and a final decision cannot be obtained.

Instead of aiming for a deterministic association before updating the catalogue, *multi-object tracking* (MOT) methods use a statistical approach and consider the contributions of multiple association hypotheses, weighted by some measure of likelihood. MOT algorithms that were successfully applied in SSA include joint probabilistic data association (JPDA) [120] and multiple hypothesis tracking (MHT) [3][10][118], but the more recent suite of methods based on finite set statistics (FISST) [81][84] has sparked particular interest. Modelling the multi-object state and measurements using random finite sets (RFSs), the framework extends the concepts of Bayesian estimation and rigorously integrates data association in the filtering process [123]. A thorough review of the literature on this topic is provided in [Appendix A](#).

Although RFS-based methods are becoming more prevalent in the SSA literature, they are not yet used in an operational catalogue and still leave much room for research. Three elements further motivate the objectives of this work.

- (i) First of all, most MOT filters consider the measurements from a single image or scan at each filtering step. Individual observations are thus treated as independent and correlations obtained from reliable tracklet formation methods are discarded. Although some efforts exist in the literature to introduce tracklets in the FISST framework, they typically reconstruct the tracklets from scratch and use fixed time intervals to do so. In reality, however, the length of a tracklet depends on how long an object is inside the observer’s FOV, and objects rarely enter and exit the FOV simultaneously. As a result, fixed-interval approaches fragment tracklets into multiple segments, limiting the effectiveness of measurement association. An approach that preserves complete pre-formed tracklets and dynamically adapts the time interval for each step of the multi-object (MO) recursion could reduce sensitivity to outliers and improve filtering performance.
- (ii) The second shortcoming of MOT approaches is that their probabilistic nature prevents definitive conclusions on which measurements belong to the same object. In MOT filters, measurements contribute to the state updates of multiple objects, and there is often no single “final decision” on the most probable measurement association. In the cataloguing context, however, it is often necessary to collect all measurements that confidently belong to a specific object for tasks such as precise orbit determination, conjunction assessment, and further post-processing. A consistent way to extract such definitive measurement-to-object associations from MO filters is not readily available in the literature.
- (iii) Finally, space cataloguing research usually focuses on measurements from ground-based observers. Now that the German company Vyoma is launching the first European surveillance and tracking satellite [31], in-situ observations are more relevant than ever, and their behaviour in MOT algorithms warrants further investigation.

These three points lead to the main research question for this work, which is further divided into three sub-questions below.

Main research question

How can multi-object tracking methods be employed for orbit cataloguing and tracklet association using space-based optical observations?

Sub-questions

- (i) Which available MOT methods are suitable for tracking and data association with space-based optical observations?
- (ii) How can the optimal measurement association with catalogued RSOs be extracted from MOT methods, independent of the state estimation?
- (iii) How can the MOT methods be adapted to leverage complete pre-formed optical tracklets?

In the process of addressing these research questions, this work focuses on the relatively new *generalised labelled multi-Bernoulli (GLMB)* filter [133][135] and its LMB approximation [104]. The main contributions are summarised in Table 1.1.

Table 1.1: Overview of contributions

Contribution	Section
A flexible Java library named BASIL is created. It provides model-agnostic MO filtering functionality which can be implemented for specific tracking scenarios. Orbital tracking capabilities are further developed based on the widely used Orekit [85] toolbox.	2, 3, 5
A robust method is proposed to extract measurement-to-object assignment probabilities from the GLMB filter, using both unlabelled and labelled observations.	4
The model for detection probability with an optical sensor is improved to be less sensitive to discontinuous boundaries in the telescope FOV.	5.2.2
The labelled multi-Bernoulli (LMB) filter is reformulated to directly use complete tracklets instead of individual observations.	6

The thesis is structured as follows. Chapter 2 first introduces fundamental background on probability density and single-object Bayesian estimation, and Chapter 3 reviews the theory of multi-object tracking using the RFS framework with a focus on the implementation of the LMB filter. These background chapters also serve as technical specification for most core functionality of BASIL, the Java library developed for this thesis. A method to robustly extract association results from GLMB/LMB filters is developed in Chapter 4. After that, Chapter 5 describes the relevant models for MOT with angles-only observations of resident space objects and Chapter 6 proposes an extension of the LMB filter that accepts multi-scan tracklets as measurements. The methods are tested and discussed based on numerical simulations in Chapter 7 and the main conclusions are summarised in Chapter 8, along with limitations and suggestions for future work. A thorough review of the literature on single-object and multi-object Bayesian estimation is provided in Appendix A.

FUNDAMENTAL BACKGROUND

This section introduces several core concepts and algorithms that are heavily used in multi-object tracking and form the basis for the methods in this work. [Section 2.1](#) focuses on representing and approximating probability distributions, as well as quantifying the “difference” between them. Then, [Section 2.2](#) provides an overview of single-object Bayesian estimation, forming a consistent background to extend to multiple objects in [Chapter 3](#).

Additionally, these theory chapters serve as technical specification for the core algorithms in the Bayesian Association and State Inference Library (BASIL), an extensible Java toolbox that was implemented from scratch in the context of this thesis.

2.1 Probability density

For a *continuous* random variable in a Euclidean space \mathbb{X} , the probability density function (PDF) $p(\mathbf{x}; \boldsymbol{\theta})$ describes the relative likelihood that the variable has value \mathbf{x} , conditioned on the distribution parameters $\boldsymbol{\theta}$ and on the assumption that it exists. As a result:

$$\forall S \subset \mathbb{X} : \Pr(\mathbf{x} \in S) = \int_{\xi \in S} p(\xi; \boldsymbol{\theta}) d\xi, \quad (2.1a)$$

$$\text{and } \Pr(\mathbf{x} \in \mathbb{X}) = 1 \quad (2.1b)$$

Similarly, if a random variable can only take values in a countable sample space Ω , it is said to be *discrete* and the probability of sampling each value $x \in \Omega$ is defined by its probability mass function (PMF) $\Pr(x; \boldsymbol{\theta})$.

$$\forall S \subset \Omega : \Pr(\mathbf{x} \in S) = \sum_{\xi \in S} \Pr(\xi; \boldsymbol{\theta}) \quad (2.2a)$$

$$\Pr(\mathbf{x} \in \Omega) = 1 \quad (2.2b)$$

Note that these two are similar, but not equivalent. The probability density function (PDF) of a continuous random variable has units $[\prod_{x \in \mathbf{x}} x]^{-1}$ and can take any value ≥ 0 as long as the integral over some region of the domain represents a probability $\in [0, 1]$. Conversely, a discrete PMF directly provides probabilities for every $x \in \Omega$.

2.1.1 Relevant probability distributions

A selection of the probability distributions with relevance in MO tracking algorithms are discussed below. Note that a distribution is completely described by its PDF. For brevity, the phrases *distribution f* and *distribution with probability density function f* are therefore used interchangeably in the context of this work.

A. Binomial distribution (discrete)

If a scalar random variable $k \sim \text{Bin}(n, r)$ follows a binomial distribution with parameters $r \in [0, 1]$ and $n \in \mathbb{N}$, its PMF is described as

$$\forall k \in \{0, \dots, n\} : \text{Bin}(k; n, r) = \binom{n}{k} r^k (1 - r)^{n-k}, \quad (2.3a)$$

$$\text{where } \binom{n}{k} = \frac{n!}{k!(n-k)!}. \quad (2.3b)$$

This can be interpreted as the probability of observing a *success* result k times in a series of n independent and identically distributed (i.i.d.) *trials*, each with exactly two discrete possible outcomes: $\Pr(\text{success}) = r$ and $\Pr(\text{fail}) = 1 - r$.

B. Bernoulli distribution (discrete)

A special case of the binomial distribution occurs when only a *single trial* is evaluated ($n = 1$),

$$\forall k \in \{0, 1\} : \quad \text{Ber}(k; r) = r^k(1 - r)^{1-k}. \quad (2.4)$$

In the context of MOT, the Bernoulli distribution is often used to represent the probability of existence for a certain object, where $k = 1$ represents existence and $k = 0$ represents non-existence.

C. Poisson distribution (discrete)

If $k \sim \text{Po}(\lambda)$ follows a Poisson distribution, this can be interpreted intuitively as representing the number of events, given an average of λ and the assumption that the events are independent. This results in the PMF

$$\text{Po}(k; \lambda) = \frac{\lambda^k e^{-\lambda}}{k!}. \quad (2.5)$$

This distribution is often used to model clutter measurements in tracking scenarios, where the number of clutter observations in a given scan is Poisson-distributed.

D. Uniform distribution (continuous)

A vector random variable $\mathbf{y} \sim \mathcal{U}(S)$ is uniformly distributed over some region S if its probability density is equal at all points in S and zero elsewhere:

$$\mathcal{U}(\mathbf{y}; S) = \mathbf{1}_S(\mathbf{y}) \frac{1}{\int_S d\mathbf{y}}, \quad (2.6)$$

where $\int_S d\mathbf{y}$ is the *hypervolume* of S and $\mathbf{1}_S$ is its *indicator function*

$$\mathbf{1}_S(\mathbf{y}) = \begin{cases} 1 & \text{if } \mathbf{y} \in S \\ 0 & \text{otherwise} \end{cases}. \quad (2.7)$$

E. Normal/Gaussian distribution (continuous)

If a continuous vector random variable $\mathbf{y} \in \mathbb{R}^n$ follows a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \mathbf{P})$ with mean $\boldsymbol{\mu} \in \mathbb{R}^n$ and covariance matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$, then the PDF is described as

$$\mathcal{N}(\mathbf{y}; \boldsymbol{\mu}, \mathbf{P}) = \frac{1}{\sqrt{\det(2\pi\mathbf{P})}} \cdot \exp\left[-\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu})^T \mathbf{P}^{-1}(\mathbf{y} - \boldsymbol{\mu})\right], \quad (2.8)$$

where the term in the exponent is negative one half of the squared *Mahalanobis distance* (see [Eq. \(2.19\)](#)).

One particularly useful property of this distribution is referred to by [Mahler \[82, App. D\]](#) as the *fundamental Gaussian identity* and is the basis for linearised Kalman filter formulations in [Section 2.2.4](#). For any $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{z} \in \mathbb{R}^m$, $\mathbf{P} \in \mathbb{R}^{n \times n}$, $\mathbf{H} \in \mathbb{R}^{m \times n}$ and $\mathbf{R} \in \mathbb{R}^{m \times m}$,

$$\mathcal{N}(\mathbf{z}; \mathbf{H}\mathbf{x}, \mathbf{R}) \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{P}) = \mathcal{N}(\mathbf{z}; \mathbf{H}\boldsymbol{\mu}, \mathbf{R} + \mathbf{H}\mathbf{P}\mathbf{H}^T) \mathcal{N}(\mathbf{x}; \mathbf{c}, \mathbf{C}), \quad (2.9a)$$

$$\text{where } \begin{cases} \mathbf{c} = \boldsymbol{\mu} + \mathbf{K}(\mathbf{z} - \mathbf{H}\boldsymbol{\mu}) \\ \mathbf{C} = (\mathbf{I}_{m \times m} - \mathbf{K}\mathbf{H}\mathbf{P}) \\ \mathbf{K} = \mathbf{P}\mathbf{H}^T(\mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R})^{-1} \end{cases}. \quad (2.9b)$$

F. Chi-squared distribution (continuous)

If a continuous vector random variable $\mathbf{x} \in \mathbb{R}^n$ follows a standard Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ as described above, then its norm is *chi-squared* (χ^2) distributed¹ with n degrees of freedom [\[140\]](#). The PDF of this scalar quantity is then

$$\chi^2(x; n) = \frac{x^{\frac{n}{2}-1} \exp(-\frac{x}{2})}{2^{\frac{n}{2}} \Gamma(\frac{n}{2})}, \quad (2.10)$$

where Γ is the gamma function. This distribution is very commonly used in statistical tests and particularly useful in the context of measurement association techniques like gating (see [Section 3.3.3](#)) or to represent confidence intervals.

2.1.2 Gaussian mixtures

Given the properties of Gaussian distributions, it is often useful to formulate more complex distributions as a weighted mixture of Gaussian components.

$$\mathcal{G}(w_{1:n}, \boldsymbol{\mu}_{1:n}, \mathbf{P}_{1:n}) \triangleq \{(w_i, \boldsymbol{\mu}_i, \mathbf{P}_i)\}_{i=1}^n \quad (2.11)$$

The probability intensity² of a variable following a Gaussian mixture (GM) distribution is simply the weighted sum of the components in the mixture

$$p(\mathbf{x}) = \sum_{i=1}^n w_i \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_i, \mathbf{P}_i). \quad (2.12)$$

In the context of MO filters, Gaussian mixture (GM) representations often end up with many components as a consequence of multiple association hypotheses. To limit the computational complexity, it is usually desirable to remove components with very low weights (*pruning*), combine components that are close enough together (*merging*) or even set a hard limit on the number of components in the GM (*capping*). These three techniques are described below and illustrated for a two-dimensional case in [Figure 2.1](#).

¹This is not the only occurrence of the Chi-squared distribution, but the most relevant in MOT.

²In the general GM formulation [Eq. \(2.11\)](#), there are no restrictions on the weights $w_{1:n}$ of the components, so \mathcal{G} does not necessarily meet the requirements of a PDF ([Eq. \(2.1\)](#)). This broader kind of probability function is referred to as *intensity* and will be used in the multi-object context (see [Section 3.2.1–E](#)).

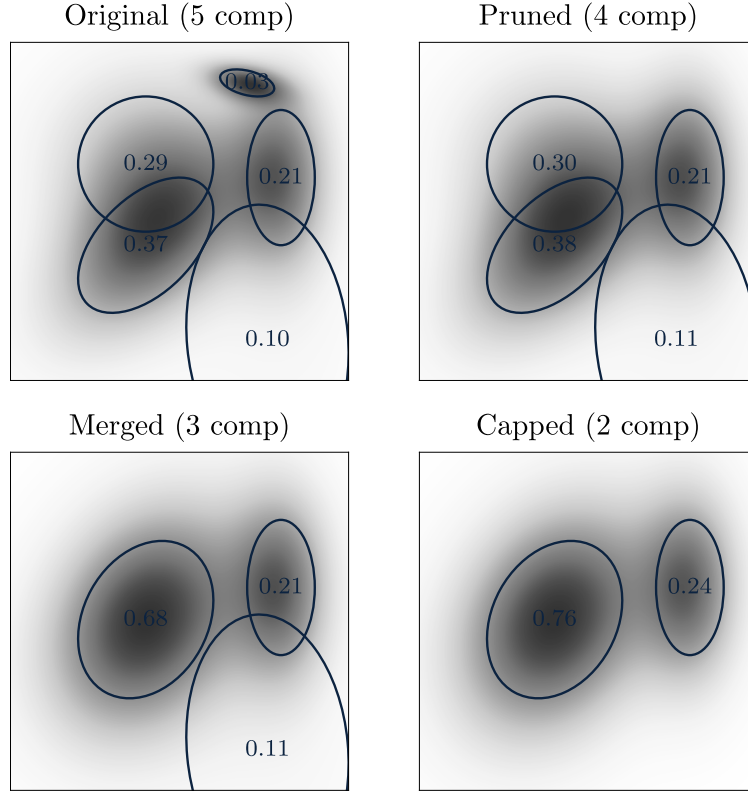


Figure 2.1: Simple 2D illustration of GM reduction methods. Every component is represented by its weight and $d_M^2 = 1$ covariance ellipsoid (see Eq. (2.19)). The original mixture has 5 components and $\rho_{\mathcal{G}} = 1$. Pruning with $\vartheta_P = 0.05$ removes one component and its weight is distributed over the others to maintain the cardinality. Merging with $\vartheta_M = 3$ combines two components and capping with $n_{\max} = 2$ limits the size of the mixture.

Note that a Gaussian mixture in this definition does not enforce that all the weights sum to 1, which is why Eq. (2.12) was referred to as the GM *intensity* instead of *density* or PDF. For MOT, the sum of weights in GM often represents the (expected) number of objects $\rho_{\mathcal{G}}$ that the mixture contains:

$$\rho_{\mathcal{G}} = \sum_{i=1}^{|\mathcal{G}|} w_i. \quad (2.13)$$

It follows that \mathcal{G} is only a PDF if its weights are *normalised* such that $\rho_{\mathcal{G}} = 1$.

A. Normalisation

Normalising a GM refers to adjusting the weights so that they sum to the desired value. This is often required after applying the GM reduction methods *pruning* and *capping*, described below. If an original GM \mathcal{G}_o is reduced to \mathcal{G}_r , it should be normalised to ensure that $\rho_{\mathcal{G}_r} = \rho_{\mathcal{G}_o}$:

$$\forall i \in \{1, \dots, |\mathcal{G}_r|\} : \quad \tilde{w}_i = w_i \frac{\rho_{\mathcal{G}_o}}{\rho_{\mathcal{G}_r}}. \quad (2.14)$$

B. Pruning

Given a pruning threshold ϑ_P , the reduced GM \mathcal{G}_r contains all components of the original mixture \mathcal{G}_o with a weight above the threshold.

$$\mathcal{G}_r = \text{prune}(\mathcal{G}_o; \vartheta_P) = \{(w_i, \boldsymbol{\mu}_i, \mathbf{P}_i) \in \mathcal{G}_o | w_i \geq \vartheta_P\} \quad (2.15)$$

C. Merging

Merging components in a GM model is usually done through a processes called *moment matching*. Merging collapses multiple components of a mixture to a single Gaussian distribution with an associated weight.

Let \mathcal{G}_o be a GM with sum of weights $\rho_{\mathcal{G}_o}$. The merged mixture has a single component with weight $\rho_{\mathcal{G}_o}$ and the first and second moments are matched so as to minimise the Kullback–Leibler (KL) divergence [Eq. \(2.20\)](#) [73] between \mathcal{G}_o and \mathcal{G}_r [123]:

$$\mathcal{G}_r = \text{merge}(\mathcal{G}_o) = \{(\rho_{\mathcal{G}_o}, \boldsymbol{\mu}_r, \mathbf{P}_r)\}, \quad (2.16a)$$

$$\text{with } \boldsymbol{\mu}_r = \frac{1}{\rho_{\mathcal{G}_o}} \sum_{i=1}^{|\mathcal{G}_o|} w_i \boldsymbol{\mu}_i \quad (2.16b)$$

$$\mathbf{P}_r = \frac{1}{\rho_{\mathcal{G}_o}} \sum_{i=1}^{|\mathcal{G}_o|} \underbrace{w_i \mathbf{P}_i}_{\text{average cov.}} + \underbrace{w_i (\boldsymbol{\mu}_i - \boldsymbol{\mu}_r)(\boldsymbol{\mu}_i - \boldsymbol{\mu}_r)^T}_{\text{spread of the mean}}. \quad (2.16c)$$

Note that, if $\rho_{\mathcal{G}_o} = 1$, $\text{merge}(\mathcal{G}_o) = \mathcal{N}(\boldsymbol{\mu}_r, \mathbf{P}_r)$.

In many cases, however, it is useful to reduce a GM by merging only components that are relatively close together. Given a merging threshold ϑ_M , [Algorithm 2.1](#) can be used to combine components where the means are close enough together. This procedure will iteratively select the component with the highest weight as reference and merge it with all components in the mixture for which the squared Mahalanobis distance from the reference to the mean is smaller than ϑ_M .

Algorithm 2.1: Gaussian mixture merging with threshold ϑ_M (adapted from [128])

```

MERGE( $\mathcal{G}_o, \vartheta_M$ ):
1  Given  $\mathcal{G}_o = \{\mathcal{E}_i \triangleq (w_i, \boldsymbol{\mu}_i, \mathbf{P}_i)\}_{i=1}^{|\mathcal{G}_o|}$ 
2  Create index set  $I = \{1, \dots, |\mathcal{G}_o|\}$ 
3  Initialise  $\mathcal{G}_r = \emptyset$ 
4  while  $I \neq \emptyset$ 
5       $j := \arg \max_{i \in I} (w_i)$ 
6       $L := \{i \in I | (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j) \mathbf{P}_j^{-1} (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)^T \leq \vartheta_M\}$ 
7       $\mathcal{G}_r := \mathcal{G}_r \cup \text{merge}(\{\mathcal{E}_i \in \mathcal{G}_o | i \in L\})$  ▷ Eq. \(2.16\)
8       $I := I \setminus L$ 
9  return  $\mathcal{G}_r$ 

```

D. Capping

Finally, the most straightforward GM reduction method is capping. Given a maximum number of components n_{\max} , the updated mixture will keep the $\min(|\mathcal{G}_i|, n_{\max})$ components with the highest weight values. All other components are deleted from the set and the remaining weights are normalised per [Eq. \(2.14\)](#). The term *capping* will be applied to other sets with similar meaning.

2.1.3 Distance and information divergence

It is often desirable to evaluate the similarity between elements of tracking algorithms. This could be to evaluate performance compared to a ground truth or other solutions, but also within the algorithm to gauge convergence, measurement association likelihood or the influence between parameters.

When comparing vectors to each other or to some probability distribution, the (dis)similarity is usually referred to as *distance*. Alternatively, one can also express the agreement between distributions, in which case the term (*information*) *divergence* is preferred as these quantities do not satisfy the formal requirements of a *metric* [24]. Several such measures for single objects are discussed below, and Section 3.3.6 extends the notion to the multi-object context.

A. Euclidean distance

The Euclidean *miss distance* between two vectors is one of the most straightforward similarity metrics:

$$d_E(\mathbf{y}_1, \mathbf{y}_2) = \|\mathbf{y}_1 - \mathbf{y}_2\|. \quad (2.17)$$

It can be used to measure the *absolute performance* of a tracking result $\mathbf{x}_{1:k}$ as compared to a ground truth $\mathbf{x}_{1:k}^{(g)}$ using an appropriate accumulation over all the time steps. One example is the root mean square (RMS) Euclidean distance

$$d_E^{\text{RMS}}(\mathbf{x}_{1:k}, \mathbf{x}_{1:k}^{(g)}) = \sqrt{\frac{1}{k} \sum_{i=1}^k \|\mathbf{x}_i - \mathbf{x}_i^{(g)}\|^2}. \quad (2.18)$$

Note that this can only be used if all state dimensions have the same units, or some form of normalisation is applied. This can cause RMS miss distances to be hard to interpret. In many cases, *sub-distances* can be useful, where only parts of the vectors are compared. A good example is comparing position and velocity errors separately.

B. Mahalanobis distance

Given an arbitrary reference distribution with first and second moments $\boldsymbol{\mu}$ and \mathbf{P} , the *Mahalanobis distance* [80] from that distribution to a vector \mathbf{y} is

$$d_M(\mathbf{y}; \boldsymbol{\mu}, \mathbf{P}) = \sqrt{(\mathbf{y} - \boldsymbol{\mu})^T \mathbf{P}^{-1} (\mathbf{y} - \boldsymbol{\mu})}. \quad (2.19)$$

This can be interpreted as a multivariate generalisation of the *standard deviation*, so that the region $d_M^2(\mathbf{y}) < 1$ defines the 1σ uncertainty of the distribution, accounting for the (co)variances of all variables [82]. For MOT, this concept is often applied in the context of measurement association and gating (see Section 3.3.3).

C. Kullback-Leibler divergence

The Kullback–Leibler (KL) divergence [73] between a distribution f_1 and a reference f_0 is

$$D_{\text{KL}}(f_1; f_0) = \int f_1(\mathbf{x}) \cdot \log\left(\frac{f_1(\mathbf{x})}{f_0(\mathbf{x})}\right) d\mathbf{x}. \quad (2.20)$$

Note that D_{KL} is always larger than 0, unless $f_0(\mathbf{x}) = f_1(\mathbf{x})$ for all values of \mathbf{x} . A divergence of 0 thus indicates equality between the distributions. Recall that the moment matching approach used in GM merging minimises D_{KL} between the original and reduced GM. As discussed later, similar methods are common in the theoretical derivation of tractable MO filters.

Other similarity measures between distributions exist, such as the *chi-squared* (χ^2) *divergence*, *Rényi divergence* [112] and *Cauchy-Schwarz (CS) divergence* [56]. Although these are often used in the literature, they are not implemented for this thesis and therefore not discussed here.

2.2 Single-object Bayesian estimation

This section is included to provide the relevant background in single object Bayesian estimation. The purpose is not to give a complete and rigorous overview, but rather to introduce the relevant concepts, terminology and notation, and detail the algorithms as implemented in the BASIL. The reader is referred to e.g. [82, Ch. 2][123] for more details.

2.2.1 State space representation

Let $\mathbf{x}_k \in \mathbb{X}$ describe the n -dimensional state of an *object* or *target* at time t_k , where $\mathbb{X} = \mathbb{R}^n$ is the set of all possible states. Similarly, let $\mathbf{z}_k \in \mathbb{Z} = \mathbb{R}^m$ be a *measurement* or *observation* of that object, provided that there exists a mapping $\mathbf{x}_k \mapsto \mathbf{z}_k$. In other words, the state space \mathbb{X} must be (partially) observable for the available sensor.

In order to track the object over time, a notion of state dynamics is needed. The dynamics of the state \mathbf{x} are represented by a motion model

$$\dot{\mathbf{x}} = F_k(\mathbf{x}_k, \mathbf{u}_k, \boldsymbol{\nu}_k), \quad (2.21)$$

where \mathbf{u}_k is a control vector, $\boldsymbol{\nu}_k$ represents some noise process and the subscript $(\cdot)_k$ indicates dependency on time, evaluated at t_k . For the purpose of this work, the control vector \mathbf{u} is ignored, assuming that any control forces can be incorporated in the dynamics of F . For sequential estimation methods, it is often useful to formulate the motion model in its integrated form, describing how a state transitions from t_{k-1} to t_k , under the influence of noise:

$$\mathbf{x}_k = \varphi_{k,k-1}(\mathbf{x}_{k-1}, \boldsymbol{\nu}_k). \quad (2.22)$$

Similarly, the observation model describes how a measurement vector \mathbf{z}_k is generated from \mathbf{x}_k , under the influence of another noise process $\boldsymbol{\varepsilon}$

$$\mathbf{z}_k = h_k(\mathbf{x}_k, \boldsymbol{\varepsilon}_k), \quad (2.23)$$

2.2.2 Single-object Bayesian filter

Consider an object with some prior state density $p_0(\mathbf{x}_0)$ at time t_0 and a single-sensor measurement sequence $Z_{1:k} = \{\mathbf{z}_0, \dots, \mathbf{z}_k\}$, originating from that object at times t_1, \dots, t_k . The single-object Bayesian filter is an estimation method that chronologically evaluates the measurements to update the state estimate and its uncertainty accordingly. Since each update depends only on the previous state density and a single measurement, the filter can be formulated recursively in two steps, often referred to as the *predictor* and *corrector* [5][81][125].

Example

In this chapter and the next, a very simple one-dimensional (1D) example will be used to visualise the steps in Bayesian filtering. Consider a train on a straight track. It's state $\mathbf{x} = (s, v)^T$ consists of a position s along that track as well as a velocity v . Assume the train is not accelerating and its dynamics are deterministic, such that $\dot{\mathbf{x}} = (v, 0)^T$. Figure 2.3 shows an arbitrary state distribution for \mathbf{x} . The observations $\mathbf{z} = (s + \varepsilon_s)$ are position-only to reduce the measurement space to a single dimension for simplicity.

(cont. on p. 13)




Figure 2.2: 1D linear example – scenario




Figure 2.3: 1D linear example – Gaussian state distribution. The red line indicates the $d_M = 1$ uncertainty bound.

A. Predictor

The motion model from Eq. (2.21) and Eq. (2.22) gives rise to the *Markov transition density* $f_k(\mathbf{x}_k|\mathbf{x}_{k-1})$, which describes the (probability) density that an initial state \mathbf{x}_{k-1} will transition to \mathbf{x}_k between t_{k-1} and t_k . If this is multiplied by the distribution of that initial state and integrated over all possible values of \mathbf{x}_{k-1} , one obtains the predicted density at time t_k based on all the measurements up to time t_{k-1} :

$$p_+(\mathbf{x}_k) \triangleq p_{k|k-1}(\mathbf{x}_k|Z_{1:k-1}) = \int f_k(\mathbf{x}_k|\mathbf{x}'_{k-1})p_{k-1}(\mathbf{x}'_{k-1})d\mathbf{x}'_{k-1}, \quad (2.24)$$

which is also called the *Chapman Kolmogorov* equation [122][134].

B. Corrector

Similarly, the *measurement likelihood function*³ $g_k(\mathbf{z}_k|\mathbf{x}_k)$ can be derived from the measurement model Eq. (2.23) and describes the likelihood of observing the measurement \mathbf{z}_k , assuming the object state is \mathbf{x}_k . The *Bayes posterior* filtering density, updated by the observation \mathbf{z}_k , is then formulated as

$$p_k(\mathbf{x}_k) \triangleq p_{k|k}(\mathbf{x}_k|Z_{1:k}) = \frac{g_k(\mathbf{z}_k|\mathbf{x}_k)p_+(\mathbf{x}_k)}{\int g_k(\mathbf{z}_k|\mathbf{x}'_k)p_+(\mathbf{x}'_k)d\mathbf{x}'_k}, \quad (2.25)$$

where the denominator is commonly referred to as the *Bayes normalisation factor* and represents the complete probability density of observing the measurement \mathbf{z}_k , for any value of the state \mathbf{x}_k .

2.2.3 Conjugate prior

Since the Bayes' filter described by Eq. (2.24) and Eq. (2.25) is intractable for general distributions, it is necessary to approximate the filtering density. One way to do that without losing the mathematical properties of Bayes optimality is to choose a state distribution for which the predictor and corrector steps result in a distribution of the same family with the relevant measurement model. A distribution that is *closed* under the Bayes update is called *conjugate prior*

³Note that, in the case of single-object, single-sensor filtering where the object always generates exactly one measurement, the measurement likelihood function is a PDF. In the general case, however, this is not necessarily the case.

[121]. Similarly, one can informally call the prediction “conjugate”, if the Chapman-Kolmogorov equation Eq. (2.24) conserves the family of the state density.

Conjugate Prior [121]

Let \mathfrak{G} be a class of likelihood functions $g(z|x)$, and \mathfrak{F} a class of prior distributions for \mathbf{x} , then the class \mathfrak{F} is conjugate for \mathfrak{G} if

$$\forall g(z|x) \in \mathfrak{G}, p_+(x) \in \mathfrak{F} : p(x|z) \in \mathfrak{F}. \quad (2.26)$$

2.2.4 Gaussian model simplification

One particularly useful conjugate prior is the Gaussian distribution (see Section 2.1.1), which is closed under both prediction and correction if $f_k(\mathbf{x}_k|\mathbf{x}_{k-1})$ and $g_k(z_k|\mathbf{x}_k)$ are Gaussian too.

First of all, assume the previous state density is Gaussian with mean $\mathbf{x}_{k-1|k-1}$ and covariance $\mathbf{P}_{k-1|k-1}$. In addition, the noise processes in the motion model Eq. (2.22) and observation model Eq. (2.23) are assumed to be additive white noise [124], such that the models can be simplified as

$$\mathbf{x}_k = \varphi_k(\mathbf{x}_{k-1}) + \boldsymbol{\nu}_k, \quad \text{with } \boldsymbol{\nu}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k), \quad (2.27a)$$

$$z_k = h_k(\mathbf{x}_k) + \epsilon_k, \quad \text{with } \epsilon_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k). \quad (2.27b)$$

Here, $\varphi_k \triangleq \varphi_{k,k-1}(\cdot, \mathbf{0})$ and $h_k \triangleq h_k(\cdot, \mathbf{0})$ represent the noiseless versions of Eq. (2.21) and Eq. (2.23).

A. Extended Kalman filter

If it is further assumed that the (non-linear) mappings φ_k and h_k can be reasonably linearised around the filter estimate $\mathbf{x}_{k-1|k-1}$, then the transition density and measurement likelihood simplify to

$$f_k(\mathbf{x}_k|\mathbf{x}_{k-1}) = \mathcal{N}(\mathbf{x}_k; \varphi_k(\mathbf{x}_{k-1}), \mathbf{Q}_k), \quad (2.28a)$$

$$g_k(z_k|\mathbf{x}_k) = \mathcal{N}(z_k; h_k(\mathbf{x}_k), \mathbf{R}_k), \quad (2.28b)$$

along with the respective Jacobian matrices (first-order Taylor approximation)

$$\boldsymbol{\Phi}_k \triangleq \boldsymbol{\Phi}_{k,k-1} = \left. \frac{\partial \varphi_k}{\partial \mathbf{x}_{k-1}} \right|_{\mathbf{x}_{k-1|k-1}}, \quad (2.29a)$$

$$\text{and } \mathbf{H}_k \triangleq \mathbf{H}_{k|k-1} = \left. \frac{\partial h_k}{\partial \mathbf{x}_k} \right|_{\mathbf{x}_{k|k-1}}. \quad (2.29b)$$

This set of assumptions leads to the well-known and widely used *extended Kalman filter (EKF)* formulation [5][21][82]. Derivations for Algorithm 2.2 are omitted here but rely mostly on the fundamental identity in Eq. (2.9).

Algorithm 2.2: Extended Kalman Filter

```

EKF( $\mathbf{x}_{0|0}$ ,  $\mathbf{P}_{0|0}$ ,  $Z_{1:K}$ ):
1  for  $k$  in  $1, \dots, K$ 
2       $\mathbf{x}_{k|k-1} = \varphi_k(\mathbf{x}_{k-1|k-1})$   $\triangleright$  noiseless prediction
3       $\mathbf{P}_{k|k-1} = \Phi_k \mathbf{P}_{k-1|k-1} \Phi_k^T + \mathbf{Q}_k$ 
4       $p_+(\mathbf{x}_k) = \mathcal{N}(\mathbf{x}_k; \mathbf{x}_{k|k-1}, \mathbf{P}_{k|k-1})$   $\triangleright$  predicted density
5
6       $\Delta \mathbf{z}_k = \mathbf{z}_k - h_k(\mathbf{x}_{k|k-1})$   $\triangleright$  mean innovation
7       $\mathbf{P}_{\Delta \mathbf{z}_k} = \mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^T + \mathbf{R}_k$   $\triangleright$  innovation covariance
8       $\mathbf{K}_k = \mathbf{P}_{k|k-1} \mathbf{H}_k^T \mathbf{P}_{\Delta \mathbf{z}_k}^{-1}$   $\triangleright$  Kalman gain
9       $\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + \mathbf{K}_k \Delta \mathbf{z}_k$ 
10      $\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_{k|k-1}$ 
11      $p_k(\mathbf{x}_k) = \mathcal{N}(\mathbf{x}_k; \mathbf{x}_{k|k}, \mathbf{P}_{k|k})$   $\triangleright$  corrected density
12 end for

```

Example (cont. from p. 11)

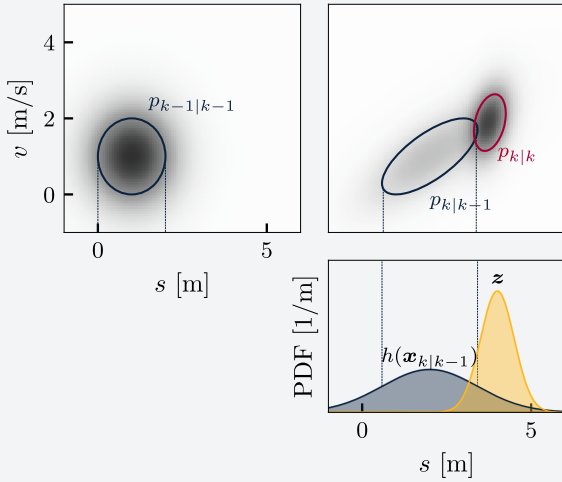


Figure 2.4: Kalman filter step for 1D linear example. *upper left:* prior PDF, *upper right:* predicted and corrected PDF, *lower right:* predicted measurement $h(\mathbf{x})$ with variance $\mathbf{H}\mathbf{P}\mathbf{H}^T$ and observed measurement $\mathcal{N}(\mathbf{z}, \mathbf{R})$. The ellipses indicate the $d_M = 1$ Mahalanobis distance bounds.

Figure 2.4 shows a single EKF step for the 1D linear example, introduced before. Recalling that the motion model assumes constant velocity and the measurements are position-only, the Jacobian matrices (recall Eq. (2.29)) for this problem are

$$\Phi_{k,k-1} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{H}_k = \begin{pmatrix} 1 & 0 \end{pmatrix}.$$

At time t_{k-1} , there is a prior density $p_{k-1|k-1}$ for the state of the train. The upper left image shows it is expected at $s = 1$ with a velocity of $v = 1$ and the state uncertainty is uncorrelated Gaussian with $\sigma_s = \sigma_v = 1$. Predicting this state PDF 1 second later increases the uncertainty and also correlates s and v . This is apparent from the non-circular uncertainty ellipse for $p_{k|k-1}$ in the upper-right image. Then mapping $p_{k|k-1}$ to measurement space (lower right) and comparing to the observed measurement \mathbf{z} allows computing the innovation. The Kalman gain incorporates the uncertainties in measurement and predicted state to compute the corrected PDF (red in upper right image), which is always smaller than the predicted density.

(cont. on p. 35)

B. Unscented Kalman filter

However, if the assumptions leading to [Eq. \(2.28\)](#) do not hold, the linearisation around $\mathbf{x}_{k-1|k-1}$ may not be valid and the Jacobians \mathbf{H}_k and $\Phi_{k,k-1}$ [Eq. \(2.29\)](#) are not constant in a sufficiently large region. In such case, the EKF loses information on the higher-order terms of the uncertainty evolution.

To address this issue, [Julier and Uhlmann \[68\]](#) introduced an improved method which guarantees at least second-order accuracy on the mean and covariance for any type of non-linear mappings [\[139\]](#): the *unscented Kalman filter (UKF)*.

Consider a Gaussian distribution $\mathcal{N}(\mathbf{x}, \mathbf{P})$ with $\mathbf{x} \in \mathbb{R}^n$. It is possible to describe this distribution using $(2n + 1)$ well chosen *sigma points* $\mathbf{X} = (\mathbf{x}_1 \dots \mathbf{x}_{2n+1})^4$, along with appropriate weights to reconstruct a Gaussian distribution. This allows performing the true non-linear mappings on each sigma point, without the need for linearisation. The (inverse) unscented transform is formulated as

$$\text{UT} : \mathcal{N}(\mathbf{x}, \mathbf{P}_x) \mapsto \mathbf{X} = \begin{pmatrix} \mathbf{x} & \mathbf{x} + \sqrt{\eta \mathbf{P}_x} & \mathbf{x} - \sqrt{\eta \mathbf{P}_x} \end{pmatrix} \quad (2.30a)$$

$$\text{UT}^{-1} : (\mathbf{X}, \mathbf{Q}) \mapsto \mathcal{N}(\mathbf{x}, \mathbf{P}_x) \quad (2.30b)$$

$$\text{with } \begin{cases} \mathbf{x} = \sum_{i=0}^{2n} w_i^{\text{mean}} \mathbf{X}_i \\ \mathbf{P}_x = \sum_{i=0}^{2n} w_i^{\text{cov}} (\mathbf{X}_i - \mathbf{x})(\mathbf{X}_i - \mathbf{x})^T + \mathbf{Q} \end{cases} \quad (2.30c)$$

$$\begin{cases} w_0^{\text{mean}} = \frac{\lambda}{\eta} \\ w_0^{\text{cov}} = \frac{\lambda}{\eta} + 1 - \alpha^2 + \beta \\ w_i^{\text{mean}} = w_i^{\text{cov}} = \frac{1}{2\eta} \quad \text{if } i > 0 \end{cases}, \quad (2.30d)$$

$$\eta = n + \lambda, \quad \text{and} \quad \lambda = \alpha^2(n + \kappa) - n. \quad (2.30e)$$

Here, λ is a scaling parameter and the notation $\sqrt{\mathbf{M}}$ refers to the lower-triangular Cholesky factorisation of a matrix \mathbf{M} . The meaning of α, β and κ is addressed by [Wan and Merwe \[139, p. 229\]](#).

Additionally, the *cross-covariance* between two distributions can be computed from their sigma point representation as

$$\mathbf{P}_{xz} = \sum_{i=0}^{2n} w_i^{\text{cov}} (\mathbf{X}_i - \mathbf{x})(\mathbf{Z}_i - \mathbf{z})^T, \quad (2.31)$$

provided that \mathbf{X} and \mathbf{Z} both have $2n + 1$ sigma points (i.e. equal number of columns). Based on [Eq. \(2.30\)](#) and [Eq. \(2.31\)](#), the unscented Kalman filter (UKF) is formulated in [Algorithm 2.3](#).

⁴Note that $\mathbf{X} \in \mathbb{R}^{n \times (2n+1)}$ where every column is one of the sigma point vectors

Algorithm 2.3: Unscented Kalman filter [139, Tab. 7.3]

```

UKF( $\mathbf{x}_{0|0}, \mathbf{P}_{0|0}, Z_{1:K}$ ):
1  for  $k$  in  $1, \dots, K$ 
2       $\mathbf{X}_{k-1|k-1} \xleftarrow{\text{UT (2.30)}} \mathcal{N}(\mathbf{x}_{k-1|k-1}, \mathbf{P}_{k-1|k-1})$        $\triangleright$  draw sigma points
3       $\mathbf{X}_{k|k-1}^* = \varphi_k(\mathbf{X}_{k-1|k-1})$        $\triangleright$  noiseless predictions
4      Compute  $\mathbf{Q}_k$        $\triangleright$  process noise
5       $\mathcal{N}(\mathbf{x}_{k|k-1}, \mathbf{P}_{k|k-1}) \xleftarrow{\text{UT}^{-1} (2.30)} (\mathbf{X}_{k|k-1}^*, \mathbf{Q}_k)$        $\triangleright$  collapse sigma points
6       $p_+(\mathbf{x}_k) = \mathcal{N}(\mathbf{x}_k; \mathbf{x}_{k|k-1}, \mathbf{P}_{k|k-1})$        $\triangleright$  predicted density
7
8       $\mathbf{X}_{k|k-1} \xleftarrow{\text{UT (2.30)}} \mathcal{N}(\mathbf{x}_{k|k-1}, \mathbf{P}_{k|k-1})$        $\triangleright$  redraw sigma-points5
9       $\mathbf{Z}_k^* = h_k(\mathbf{X}_{k|k-1})$        $\triangleright$  noiseless meas. prediction
10     Compute  $\mathbf{R}_k$        $\triangleright$  measurement noise
11      $\mathcal{N}(z_{k|k-1}, \mathbf{P}_{\Delta z_k}) \xleftarrow{\text{UT}^{-1} (2.30)} (\mathbf{Z}_k^*, \mathbf{R}_k)$        $\triangleright$  predicted meas. distribution
12      $\mathbf{P}_{xz_k^*} \xleftarrow{\text{Eq. (2.31)}} (\mathbf{X}_{k|k-1}, \mathbf{Z}_k^*)$        $\triangleright$  cross-covariance
13      $\mathbf{K}_k = \mathbf{P}_{xz_k^*} \mathbf{P}_{z_k}^{-1}$        $\triangleright$  Kalman gain
14
15      $\Delta z_k = z_k - z_{k|k-1}$        $\triangleright$  innovation
16      $\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + \mathbf{K}_k(\Delta z_k)$ 
17      $\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} + \mathbf{K}_k \mathbf{P}_{\Delta z_k} \mathbf{K}_k^T$ 
18      $p_k(\mathbf{x}_k) = \mathcal{N}(\mathbf{x}_k; \mathbf{x}_{k|k}, \mathbf{P}_{k|k})$        $\triangleright$  corrected density
19  end for

```

C. Gaussian mixture as filtering density

As will become clear, it is often convenient in MOT algorithms to represent complex distributions as Gaussian mixtures. When faced with a new measurement, the mixture must be predicted and updated accordingly. Algorithm 2.4 outlines the prediction step from t_{k-1} to t_k .

Algorithm 2.4: Gaussian mixture prediction

```

GM-PRED( $\mathcal{G}_{k-1|k-1}, t_k$ ):
1  define  $\mathcal{G}_{k|k-1} \triangleq \left\{ \mathcal{E}_{k|k-1}^{(i)} \right\}_{i=1}^{|\mathcal{G}_{k-1|k-1}|}$ 
2  for  $\left( w_{k-1|k-1}^{(i)}, \mathbf{x}_{k-1|k-1}^{(i)}, \mathbf{P}_{k-1|k-1}^{(i)} \right)$  in  $\mathcal{G}_{k-1|k-1}$ 
3       $\left( \mathbf{x}_{k|k-1}^{(i)}, \mathbf{P}_{k|k-1}^{(i)} \right) \xleftarrow{\text{EKF/UKF/...}} \left( \mathbf{x}_{k-1|k-1}^{(i)}, \mathbf{P}_{k-1|k-1}^{(i)} \right)$        $\triangleright$  prediction
4       $\mathcal{E}_{k|k-1}^{(i)} = \left( w_{k-1|k-1}^{(i)}, \mathbf{x}_{k|k-1}^{(i)}, \mathbf{P}_{k|k-1}^{(i)} \right)$ 
5  end for
6  return  $\mathcal{G}_{k|k-1}$ 

```

In turn, Algorithm 2.5 details the update step for a GM density based on a single measurement z_k . Note that the sum of weights before normalisation q_z is the measurement likelihood for the measurement on the entire GM. Normalisation is required to ensure the mixture is still a PDF,

⁵Wan and Merwe [139, p. 233] note that *redrawing* sigma points discards odd-moment uncertainty information and suggest an alternative *augmenting* approach. However, under the additive noise assumptions of Eq. (2.27), this gives the same result.

so it might be omitted in case the GM represents something else (such as a multi-object intensity function; see [Section 3.2.1–E](#)).

Algorithm 2.5: Gaussian mixture single-measurement Bayes update

```

GM-CORR( $\mathcal{G}_{k|k-1}, \mathbf{z}_k$ ):
1  define  $\mathcal{G}_{k|k} \triangleq \left\{ \mathcal{E}_{k|k}^{(i)} \right\}_{i=1}^{|\mathcal{G}_{k|k-1}|}$ 
2  for  $\left( w_{k-1|k-1}^{(i)}, \mathbf{x}_{k-1|k-1}^{(i)}, \mathbf{P}_{k-1|k-1}^{(i)} \right)$  in  $\mathcal{G}_{k-1|k-1}$ 
3       $\left( \mathbf{x}_{k|k}^{(i)}, \mathbf{P}_{k|k}^{(i)} \right) \xleftarrow{\text{EKF/UKF/...}} \left( \mathbf{x}_{k|k-1}^{(i)}, \mathbf{P}_{k|k-1}^{(i)} \right)$   $\triangleright$  component update
4
5      Extract  $\Delta \mathbf{z}^{(i)}, P_{\Delta \mathbf{z}_k}^{(i)}$  from EKF/UKF/...  $\triangleright$  innovation and cov
6       $w_{k|k}^{(i)} = w_{k|k-1}^{(i)} \cdot \mathcal{N}\left(\Delta \mathbf{z}_k^{(i)}, \mathbf{0}, P_{\Delta \mathbf{z}_k}^{(i)}\right)$   $\triangleright$  component weight
7       $\mathcal{E}_{k|k}^{(i)} = \left( w_{k|k}^{(i)}, \mathbf{x}_{k|k}^{(i)}, \mathbf{P}_{k|k}^{(i)} \right)$ 
8  end for
9   $q_{z_k} = \rho_{\mathcal{G}_{k|k}}$   $\triangleright$  combined “likelihood”
10 normalise( $\mathcal{G}_{k|k-1}$ )  $\triangleright$  Eq. \(2.14\)
11 return  $\mathcal{G}_{k|k-1}, q_{z_k}$ 

```

MULTI-OBJECT TRACKING WITH RANDOM FINITE SETS

This section addresses the key aspects of the multi-object tracking (MOT) problem, with a focus on its formulation and solutions based on finite set statistics (FISST), here referred to as random finite sets (RFSs). As in [Chapter 2](#), the theory, methods and algorithms described here serve as technical specification for the implementations in BASIL. The Java library is briefly introduced in [Section 3.4](#).

3.1 Multi-object representation

At a given time t_k , a multi-object state $X_k \subset \mathbb{X}$ is a discrete set of target states \mathbf{x}_k , belonging to different objects. The purpose of a multi-object tracking algorithm is to estimate a multi-object trajectory $X_{j:k}$, consisting of an estimated multi-object state at every instant in the discrete time window $\mathbb{T}_{j:k} = \{t_j, \dots, t_k\}$. However, for any practical application, this is not enough and one is interested in extracting individual trajectories $\tau : \mathbb{T} \rightarrow \mathbb{X}$ [\[8\]](#) for each distinct object. [Vo et al. \[134\]](#) claim that “pragmatically, *labelling* is unavoidable in real multi-object systems” and the notion of a target label ℓ is required to distinguish between objects.

Given an *attribute state* $\mathbf{x} \in \mathbb{X}$ and *label* $\ell \in \mathbb{L}$, the single-object (SO) labelled state is $\hat{\mathbf{x}} \triangleq (\mathbf{x}, \ell)$. In principle, any type of label space \mathbb{L} is possible as long as it ensures a unique label for all objects, but [Vo et al. \[133\]\[135\]](#) suggest a common choice is $\ell \triangleq (t_B, \iota)$, where t_B is the time where the label first occurs –*birth*– and ι is an index, allowing multiple labels to be created simultaneously. It follows that $\hat{X}_k \subset \mathbb{X} \times \mathbb{L}$ is a labelled MO state and $\hat{X}_{j:k}$ represents a labelled MO trajectory. Note that the labelled MO trajectory satisfies the need to extract the individual state history per object over time. For every label $\ell \in \mathcal{L}(\hat{X}_{j:k})$, one can construct a single-object trajectory:

$$\tau_\ell : t_i \mapsto \mathbf{x}_i, \quad \forall t_i \in \mathcal{D}(\tau_\ell), \quad (3.1)$$

Here, $\mathcal{L} : (\mathbf{x}, \ell) \mapsto \ell$ maps a state to its label and the domain $\mathcal{D}(\tau_\ell) \subseteq \mathbb{T}$ is the set of instants where the label ℓ exists [\[11\]\[134\]](#). The remainder of this chapter addresses how these trajectories and their uncertainties can accurately be extracted from measurements.

3.2 Random finite sets

One way to solve the MOT problem is based on RFSs, a branch of point process theory [\[19\]\[27\]\[63\]](#) that was specifically tailored for the context of tracking and information fusion. It provides a generalisation of the Bayesian framework discussed in [Section 2.2](#) by defining states as sets of object states, along with probability distributions on those sets. This section describes some of the fundamentals [\[84\]](#).

3.2.1 Definitions and finite set statistics

Integer random variables (e.g. on \mathbb{N}_0) and vector random variables (on \mathbb{R}^n) were already briefly addressed in [Section 2.1](#). RFSs take this concept one step further and transform sets into random variables. This means that the number of elements in the set – its *cardinality* – is also subject to randomness.

Let \mathfrak{Y} be some underlying space, such as the state space \mathbb{X} or measurement space \mathbb{Z} . The *hyperspace of all finite subsets* of \mathfrak{Y} is then referred to as $\mathcal{F}(\mathfrak{Y})$. The random finite set Ψ is a random variable on such a hyperspace. This means that a realisation Y of Ψ is a set containing a random number of elements, each randomly selected from \mathfrak{Y} , according to some governing distributions [82].

For example, $X \in \mathcal{F}(\mathbb{X})$ is a random target set, where each of the targets $\mathbf{x} \in X$ is drawn from a distribution on \mathbb{X} . Similarly, $Z \in \mathcal{F}(\mathbb{Z})$, is a random measurement set, where each of the measurements $\mathbf{z} \in Z$ is a realisation of some random variable on \mathbb{Z} .

The value of set distributions for MOT algorithms is that they allow extending the rigorous principles of single-object Bayesian estimation (recall Section 2.2) to multi-object scenarios with an unknown number of targets.

A. Set integral

For a function $f : \mathcal{F}(\mathfrak{Y}) \rightarrow \mathbb{R}$ that maps any set of values $\mathbf{y}_j \in \mathfrak{Y}$ to a real number, the set integral over a region $S \subseteq \mathfrak{Y}$ is defined by Goodman et al. [54] as

$$\int_S f(Y) \delta Y = f(\emptyset) + \sum_{i=1}^{\infty} \frac{1}{i!} \int_{\underbrace{S \times \dots \times S}_i} f(\{\mathbf{y}_1, \dots, \mathbf{y}_i\}) d\mathbf{y}_1 \dots d\mathbf{y}_i, \quad (3.2)$$

where the component $i!$ can be understood intuitively as a normalisation factor to account for all the permutations to order the elements in the set $\{\mathbf{y}_1, \dots, \mathbf{y}_i\}$. By definition, they should all give the same result [82, p. 361].

B. Multi-object probability density

Very similar to a single-object density function, FISST offers the notion of a *belief* density. This is a non-negative real-valued function $\pi_\Psi : \mathcal{F}(\mathfrak{Y}) \rightarrow \mathbb{R}_0^+$, for which the set integral over a region S is the probability that a realisation of Ψ is a subset of S . This so-called (*FISST*) *belief mass* becomes unity if $S = \mathfrak{Y}$:

$$\forall S \subseteq \mathfrak{Y} : \int_S \pi_\Psi(Y) \delta Y = \Pr(\Psi \subseteq S) \quad \text{and} \quad \int_{\mathfrak{Y}} \pi_\Psi(Y) \delta Y = 1. \quad (3.3)$$

For the purpose of MOT, this is often referred to as the multi-object probability density function, even though various other notions of set probability density exist [63][134]. Note that, if $\mathbf{y} \in \mathfrak{Y}$ has units u , then the units of $\pi_\Psi(Y)$ must be $u^{-|Y|}$ to cancel out the dimensions of $d\mathbf{y}_1 \dots d\mathbf{y}_{|Y|}$ in Eq. (3.2) [82]. For this reason, it is hard to meaningfully compare values of multi-object density. Vo et al. [130, Sec. 1][134, Sec. II-C] illustrate this intuitively with a clear example, reproduced here for reference.

Example [130]

Imagine collecting apples in a 1D forest. Every day, the number of apples is uniformly distributed between 0 and 9, and each apple falls i.i.d. following the PDF in Figure 3.1. On day one, the set of apples is $X_1 = \{a, b\}$ and on day two, $X_2 = \{c\}$. Which of these has a higher likelihood? Noting that there is a $\frac{1}{10}$ chance of observing any number of apples $m < 10$, the joint density $\pi(\{x_1, \dots, x_m\}) = \frac{1}{10} \prod_{i=1}^m p(x_i)$. So, $\pi(X_1) = 3.6 \times 10^{-2} \text{ m}^{-2}$ and $\pi(X_2) = 2 \times 10^{-2} \text{ m}^{-1}$. At first glance X_1 might appear more likely, but since the densities have different units, they cannot be meaningfully compared. If we would measure in mm instead, the PDF would be scaled accordingly and the joint density values would become $\pi(X_1) = 3.6 \times 10^{-8} \text{ mm}^{-2}$ and $\pi(X_2) = 2 \times 10^{-4} \text{ mm}^{-1}$, respectively. [134]

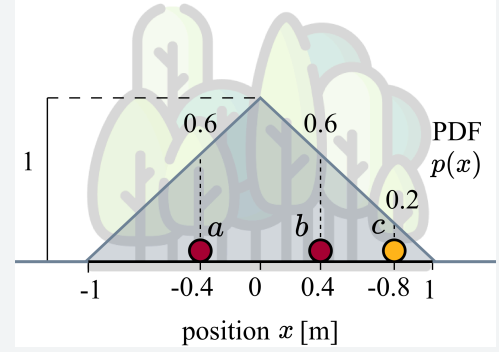


Figure 3.1: Why multi-object densities are hard to compare. The MO densities of events $\{a, b\}$ and $\{c\}$ have different units⁶

C. Convolution of RFS densities

If $\Psi = \biguplus_{i=1}^n \Psi_i$ is the union of n disjoint and independent RFSs, then its multi-object density follows from the *RFS convolution* [84, p. 86][134]

$$\pi_{\Psi}(Y) = \sum_{\biguplus_{i=1}^n Y_i = Y} \left[\prod_{i=1}^n \pi_{\Psi_i}(Y_i) \right]. \quad (3.4)$$

D. Cardinality distribution

The cardinality distribution of a random finite set is a discrete distribution defined by [19][82][134]

$$\rho_{\Psi}(n) \triangleq \Pr(|\Psi| = n) = \frac{1}{n!} \int_{\underbrace{\mathfrak{Y} \times \dots \times \mathfrak{Y}}_n} \pi_{\Psi}(\{\mathbf{y}_{1:n}\}) d\mathbf{y}_{1:n}. \quad (3.5)$$

So, $\rho_{\Psi}(n)$ is the probability that a realisation of the RFS Ψ contains exactly n elements.

E. Probability hypothesis density – intensity

A commonly used RFS statistic is the *probability hypothesis density (PHD)* or *probability intensity* v_{Ψ} . It is the first moment approximation of an RFS [81], and can be seen intuitively as the equivalent of “expected value” in the single-object case. More specifically, the integral of the probability hypothesis density (PHD) over some region $S \subseteq \mathfrak{Y}$ results in the expected cardinality of the RFS inside that region,

$$\int_S v_{\Psi}(\mathbf{y}) d\mathbf{y} = \mathbb{E}[|\Psi \cap S|]. \quad (3.6)$$

Notice that, if the cardinality is certainly 1 (i.e. $\rho_{\Psi}(n) = 1$ if $n = 1$ and 0 otherwise), then Eq. (3.6) reduces to Eq. (2.1) and v_{Ψ} becomes equivalent to a single-object PDF.

⁶Pictogram from Flaticon.com

3.2.2 Labelled random finite set

As discussed in [Section 3.1](#), labels are essential to model, estimate and extract MO states and trajectories. [Vo et al. \[133\]\[135\]](#) introduced this notion into the RFS framework, by defining the labelled random finite set (LRFS) $\dot{\Psi}$ as a RFS on $\mathbb{X} \times \mathbb{L}$ such that its realisations (labelled MO states) are of the form

$$\dot{X} \triangleq \{(\mathbf{x}_i, \ell_i)\}_{i=1}^{|\dot{X}|}. \quad (3.7)$$

For convenience, we define $\mathcal{L}(\dot{X})$ as the set of *labels* and *attributes* in \dot{X} , respectively. The MO density of $\dot{\Psi}$ is denoted $\dot{\pi}_{\Psi}$.

In addition, LRFS theory uses the so-called *joint existence probability* $w_{\dot{\Psi}}(L)$ to indicate the probability that a realisation of $\dot{\Psi}$ contains each of the labels in $L \subseteq \mathbb{L}$ [\[134\]](#). Clearly, this joint existence probability can be used to determine the cardinality distribution by summing all possible $w_{\dot{\Psi}}(L)$ with the same number of existing labels:

$$\rho_{\dot{\Psi}}(n) = \sum_{L \subseteq \mathbb{L}} \delta_n[|L|] w_{\dot{\Psi}}(L), \quad (3.8)$$

where $\delta_a[b]$ is the *Kronecker delta* and evaluates to 1 if $a = b$ and 0 otherwise. Read [Eq. \(3.8\)](#) as “take all possible label sets with exactly n elements, compute the probability that all labels in these subsets exist at the same time and then sum those probabilities”. Based on the labelled PHD/intensity $\dot{v}(\mathbf{x}, \ell)$, one can also infer the existence probability of a single label and the attribute distribution for a given label provided that it exists.

$$r_{\dot{\Psi}}(\ell) = \int \dot{v}_{\dot{\Psi}}(\mathbf{x}, \ell) d\mathbf{x}, \quad p_{\dot{\Psi}}(\cdot, \ell) = \frac{\dot{v}_{\dot{\Psi}}(\cdot, \ell)}{r_{\dot{\Psi}}(\ell)} \quad (3.9)$$

3.2.3 Relevant RFS models

A. Independent and identically distributed Cluster RFS

An i.i.d. cluster is a type of RFS where the cardinality follows some discrete distribution ρ_{Ψ} , and each of the components has the same independent “spatial” distribution p_{Ψ} .

B. Poisson RFS

A special case of i.i.d. cluster appears when the cardinality distribution is Poisson. This is referred to interchangeably as a Poisson point process (PPP) or Poisson RFS and has many applications in MOT, specifically when modelling clutter (see [Section 3.3.2-B](#)).

Notice that a Poisson point process (PPP) is completely characterised by its PHD v_{Ψ} and the cardinality is Poisson distributed with mean $\lambda_{\Psi} = \langle v_{\Psi}, 1 \rangle$, such that

$$\pi_{\Psi}(Y) = e^{-\lambda_{\Psi}} \prod_{\mathbf{y} \in Y} v_{\Psi}(\mathbf{y}), \quad \rho_{\Psi}(n) = \frac{e^{-\lambda_{\Psi}} \lambda_{\Psi}^n}{n!}. \quad (3.10)$$

Intuition is easily obtained from [Algorithm 3.1](#), which describes that one can sample a set from a PPP by sampling the number of targets n from a discrete Poisson distribution and then sampling n targets from the single-object PDF $p_{\Psi} = v_{\Psi}/\lambda_{\Psi}$.

Algorithm 3.1: Sampling from a Poisson RFS

```

SAMPLE( $v_\Psi$ ):
1   $\lambda_\Psi = \langle v_\Psi, 1 \rangle$ 
2   $p_\Psi = v_\Psi / \lambda_\Psi$ 
3  sample  $n \sim \text{Po}(\lambda_\Psi)$  ▷ cardinality Eq. (3.10)
4   $Y = \{\text{sample } \mathbf{y}_i \sim p_\Psi\}_{i=1}^n$  ▷ states
5  return  $Y$ 
    
```

Note that the term $n!$ in Eq. (3.10) emerges from the integrated MO density because of the fact that there are $n! = |Y|!$ permutations to order the elements in Y (see p. 18).

C. Bernoulli RFS

A Bernoulli (RFS) is parametrised by a single-object PDF p and an *existence probability* $r \in [0, 1]$. It is particularly appealing because of its intuitive tracking interpretation as a *hypothesised target* [84, p. 100], i.e. if the target exists, $\Psi = \{\mathbf{y}\}$ with $\mathbf{y} \sim p$ and if not, the set is empty $\Psi = \emptyset$. The corresponding MO density, cardinality distribution and PHD are

$$\pi_\Psi(Y) = (1 - r)\delta_\emptyset[Y] + rp(\mathbf{y})\delta_{\{\mathbf{y}\}}[Y], \quad (3.11a)$$

$$\rho_\Psi(n) = (1 - r)\delta_0[n] + r\delta_1[n], \quad (3.11b)$$

$$v_\Psi(\mathbf{y}) = rp(\mathbf{y}). \quad (3.11c)$$

D. Multi-Bernoulli RFS

A union of independent Bernoullis with distinct components is called a multi-Bernoulli (RFS). Giving the identifiers $\mathcal{B} = \{\beta_1, \dots, \beta_{|\mathcal{B}|}\}$ to those Bernoullis, the multi-Bernoulli is completely parametrised by $\pi_\Psi = \{(r^{(\beta)}, p^{(\beta)})\}_{\beta \in \mathcal{B}}$. The probability that none of the components exist is $\pi_\Psi(\emptyset) = \prod_{\beta \in \mathcal{B}} (1 - r^{(\beta)})$ and the characteristic distributions are

$$\pi_\Psi(Y) = \sum_{\substack{\mathcal{B}_i \subseteq \mathcal{B} \\ |\mathcal{B}_i| = |Y|}} \left[\overbrace{\prod_{\beta_j \in \mathcal{B} \setminus \mathcal{B}_i} (1 - r^{(\beta_j)})}^{\text{all excluded don't exist}} \overbrace{\prod_{\beta_j \in \mathcal{B}_i} r^{(\beta_j)} p^{(\beta_j)}(\mathbf{y}_j)}^{\text{all included exist and SO probability density}} \right] \quad (3.12a)$$

$$= \pi_\Psi(\emptyset) \cdot \sum_{\substack{\mathcal{B}_i \subseteq \mathcal{B} \\ |\mathcal{B}_i| = |Y|}} \left[\prod_{\beta_j \in \mathcal{B}_i} \frac{r^{(\beta_j)} p^{(\beta_j)}(\mathbf{y}_j)}{1 - r^{(\beta_j)}} \right], \quad (3.12b)$$

$$v_\Psi(\mathbf{y}) = \sum_{\beta \in \mathcal{B}} r^{(\beta)} p^{(\beta)}(\mathbf{y}). \quad (3.12c)$$

The factorisation of Eq. (3.12b) is essential in many MOT implementations and particularly to allow *truncation before computation*, discussed in several parts of Section 3.3.4. Notice that $\pi_\Psi(\emptyset)$ cancels out all the denominators and leaves the relevant components in the numerator, thereby modelling the non-existence probabilities *implicitly*.

E. Labelled multi-Bernoulli RFS

A labelled multi-Bernoulli (LMB) is an LRFS that consists of a set of Bernoullis components, where each component has a unique label $\ell \in \mathbb{L}$. Since the LMB is completely characterised by those components, the labelled MO density can be written as the set

$$\dot{\pi}_{\dot{\Psi}} \triangleq \{(r^{(\ell)}, \dot{p}_{\dot{\Psi}}(\cdot, \ell))\}_{\ell \in L}. \quad (3.13)$$

From an MOT perspective, this density has a very clear interpretation: The distribution has labels for $|L|$ different objects. For each of those objects, there is a probability of $r^{(\ell)}$ that it exists and if it does, its state follows the attribute PDF $\dot{p}_{\dot{\Psi}}(\cdot, \ell)$ [84, p. 453].

F. Generalised labelled multi-Bernoulli RFS

A final step to incorporate even more variability is made with the generalised labelled multi-Bernoulli (GLMB) RFS [133][135], which is parametrised by a set of hypotheses

$$\dot{\pi}_{\dot{\Psi}} = \{(w^{(\xi)}(L), \dot{p}^{(\xi)})\}_{(\xi, L) \in \mathcal{H}}, \quad (3.14)$$

that represent the MO density. The GLMB probability density of a labelled state \dot{X} can be written in various equivalent ways, but for numerical implementations, the δ GLMB form is most convenient. The total density is then a weighted sum over all hypotheses:

$$\dot{\pi}_{\dot{\Psi}}(\dot{X}) = \underbrace{\Delta_{\dot{X}}}_{\text{label uniqueness hypotheses}} \sum_{(\xi, L) \in \mathcal{H}} \left[\underbrace{w^{(\xi)}(L)}_{\text{joint existence probability}} \underbrace{\prod_{(\mathbf{x}, \ell) \in \dot{X}} \dot{p}^{(\xi)}(\mathbf{x}, \ell)}_{\text{combined density}} \right], \quad (3.15)$$

where $\Delta_{\dot{X}} \triangleq \delta_{|\dot{X}|}[\mathcal{L}(X)]$ is 1 if the labels in \dot{X} are unique and 0 otherwise. Again, reading this from a tracking point of view, the MO density takes into account a set of *hypotheses* $(\xi, L) \in \mathcal{H}$, where L contains all the targets (labels) that exist in that hypothesis and ξ reflects the history of measurement associations (see Eq. (3.18)). The state of these targets then follows the distribution $\dot{p}^{(\xi)}(\cdot, \ell)$ for each $\ell \in L$. Finally, the probability that all these targets exist together with this particular distribution is $w^{(\xi)}(L)$, often called the *weight* of the hypothesis.

Vo and Vo [135] showed that this RFS is *conjugate prior* (recall Section 2.2.3) with respect to the multi-object measurement likelihood from Eq. (3.22) and also *closed* under the prediction step. As a result, the GLMB is a valid MO state density representation for Bayesian filtering.

3.3 Multi-object Bayesian filtering

As mentioned before, the entire purpose of introducing RFSs is to extend principles of single-object Bayesian estimation (Section 2.2) to the multi-object context. The focus here is on the elements needed to implement a Bayesian filter based on the LMB/GLMB formulation (Section 3.2.3-E and 3.2.3-F). Given some prior LRFS distribution $\dot{\pi}_{k-1}$ and scan of measurements Z_k , the filter must do three things:

- (i) Use the system dynamics to transition the multi-object distribution to $\dot{\pi}_{k|k-1}$.
- (ii) Use the measurement model to update the distribution to $\dot{\pi}_{k|k}$.
- (iii) Estimate the multi-object state $\dot{X}_{k|k} = \{(\dot{\mathbf{x}}_{k|k}, \ell)\}$ and update the trajectory $\tau^{(\ell)}$ for each of the targets ℓ .

These steps are conceptually equivalent to the single-object filtering logic. However, there are some key differences when dealing with set states, caused by the additional uncertainty as to the number of distinct objects that are being tracked.

First of all, the number of hypothetical objects (*targets*) represented by the LRFS state can vary over time. The process of introducing new labels into the MO density is commonly referred

to as *target birth*. Conversely, the existence of a target that is no longer supported by any observations can become so improbable that it is removed from the state (*target death*). From an SSA perspective, recall that only a very small fraction of the RSO population is tracked by active catalogues. In any measurement scan, it is thus possible that some of the observations belong to previously unknown objects. A *birth process* inside the transition model (the first step of the filtering recursion) will account for this possibility. This is already touched upon in [Section 3.3.1](#), but [Section 5.3](#) later develops this for the orbital tracking problem.

Furthermore, there can now be multiple measurements at every observation epoch (a single image can capture light reflected from multiple satellites) and the MO filter must deal with the consistent assignment of these measurements to the available tracklets. This means that the *joint* measurement likelihood not only depends on how well the observations fit the assigned objects, but also how this likelihood compares to all other possible assignments. This leads to the MO measurement model introduced in [Section 3.3.2](#) and requires carefully ranking likely measurement assignments, discussed in [Section 3.3.3](#).

3.3.1 Multi-object transition model

Compared to the single-object case, the MO transition model has the additional “responsibility” of dealing with *target birth* and *death* (see above). That is, the MO Bayesian filter must allow for additional objects to be initialised as new measurements come in, or to conclude that an object no longer exists.

In the *standard multi-object model* [135], the birth density is represented by a GLMB ([Eq. \(3.14\)](#)), but an LMB ([Eq. \(3.13\)](#)) is usually sufficient and offers the more intuitive interpretation, where each component represents a potential new object with distribution $p_{B,k}$ and some birth probability P_B

$$\dot{\pi}_{B,k} = \left\{ \left(P_{B,k}^{(\ell)}, p_{B,k}(\cdot, \ell) \right) \right\}_{\ell \in \mathcal{L}(B_k)}. \quad (3.16)$$

Here, $\mathbb{B}_k = \mathcal{L}(B_k)$ is a set of labels that did not previously exist, i.e. $\mathbb{B}_k \cap \mathbb{L}_{k-1} = \emptyset$. The *surviving* GLMB is defined by a survival probability $P_{S,k}(\dot{\mathbf{x}}_{k-1})$ and a single object transition density $f_{S,k}(\cdot | \dot{\mathbf{x}}_{k-1})$. Both survival and birth are incorporated in the MO transition density \dot{f}_k , which allows for a direct generalisation of the Chapman-Kolmogorov equation in [Eq. \(2.24\)](#) to RFSs [135]:

$$\dot{\pi}_+(\dot{X}_k) = \int \dot{f}(\dot{X}_k | \dot{X}'_{k-1}) \dot{\pi}_{k-1}(\dot{X}'_{k-1}) \delta \dot{X}'_{k-1}. \quad (3.17)$$

If both the survival and birth densities are in δ -generalised labelled multi-Bernoulli (δ GLMB) form [Eq. \(3.14\)](#), then the predicted density can be written conveniently as $\dot{\pi}_+^{\text{GLMB}} = \dot{\pi}_{B,k}^{\text{GLMB}} \cup \dot{\pi}_{S,k}^{\text{GLMB}}$, which is useful in implementation, as discussed in [Section 3.3.4](#).

3.3.2 Multi-object measurement model

Another elementary requirement to transform the Bayesian framework to MOT is a multi-object version of the measurement likelihood. In the discussion on single-object estimation ([Section 2.2](#)), it was implicitly assumed that the measurement sequence $Z_{1:k}$ consisted of exactly one measurement per time step, which always originated from the object of interest. However, an object might not be detected, or clutter measurements might exist that contain no information about the object’s state. In the general case, a measurement scan $Z_k = O_k \uplus C_k$, where O_k is the

set of observations originating from some object states in \dot{X}_k (others could be undetected) and C_k is the set of clutter measurements.

Apart from the actual state estimations, a tracker must solve the problem of *association*, i.e. determining which objects in \dot{X}_k caused which measurements in the scan Z_k . An answer to this problem is referred to as an *association hypothesis* $\theta_k : \mathcal{L}(\dot{X}_k) \rightarrow \{0, \dots, |Z_k|\}$, which maps every label in \dot{X}_k either to a distinct measurement index or to 0 if the label is undetected.

$$\theta_k(\ell) = \begin{cases} i & \text{if object } \ell \text{ caused measurement } i : \mathbf{o}_k^{(\ell)} \triangleq \mathbf{z}_k^{(i)} \\ 0 & \text{if object } \ell \text{ is undetected} \end{cases} \quad (3.18)$$

Consequently, one can define the “inverse” mapping $\theta_k^{-1} : \{1, \dots, |Z_k|\} \rightarrow \mathcal{L}(\dot{X}_k) \cup \{\mathbf{c}\}$, where \mathbf{c} represents the *clutter/unknown label*:

$$\theta_k^{-1}(i) = \begin{cases} \ell & \text{if object } \ell \text{ caused measurement } \mathbf{z}_k^{(i)} \\ \mathbf{c} & \text{if no known object caused } \mathbf{c}_k^{(i)} \triangleq \mathbf{z}_k^{(i)} \end{cases} \quad (3.19)$$

For a given set of prior states \dot{X}_k and a measurement scan Z_k , $\Theta_k = \{\theta_k^i, \dots\}$ is the set of all possible association hypotheses.

A. Object observations

Assume an object with state $\hat{\mathbf{x}} \in \dot{X}$ has a detection probability $P_D(\hat{\mathbf{x}})$ and if it is detected, the resulting measurement $\mathbf{o}_k^{(\ell)}$ follows the SO density $g_k(\mathbf{o}|\hat{\mathbf{x}})$. This means that the SO observation set $O_k^{(\ell)}$ follows a Bernoulli RFS [Eq. \(3.11\)](#) and as a result, the MO observation set $O_k = \biguplus_{(\mathbf{x}, \ell) \in \dot{X}} O_k^{(\ell)}$ is multi-Bernoulli [\[122\]](#) and its MO likelihood follows the convolution [Eq. \(3.4\)](#)

$$g_k^o(O|\dot{X}_k) = \sum_{\biguplus_{\ell \in \mathbb{L}_k} O^{(\ell)} = O} \left[\prod_{(\mathbf{x}, \ell) \in \dot{X}_k} g_k(O^{(\ell)}|\{\hat{\mathbf{x}}\}) \right], \quad (3.20a)$$

$$\text{where } \mathbb{L}_k = \mathcal{L}(\dot{X}), \quad (3.20b)$$

$$\text{and } g_k(O|\{\hat{\mathbf{x}}\}) \stackrel{\text{Eq. (3.11)}}{=} \begin{cases} P_D(\ell)g_k(\mathbf{o}|\hat{\mathbf{x}}) & \text{if } O = \{\mathbf{o}\} \\ 1 - P_D(\hat{\mathbf{x}}) & \text{if } O = \emptyset \end{cases}. \quad (3.20c)$$

Intuitively, this means that one should consider all possible ways that the states could have generated these observations, evaluate the probability that each of those actually happened and sum these probabilities.

B. Clutter measurements

The other source of measurements –the clutter set C_k – is commonly modelled as a Poisson point process [Eq. \(3.10\)](#) with time-dependent clutter intensity function $\kappa_c \triangleq \kappa_c(\cdot, t_k)$:

$$\pi_c(C) = e^{-\lambda_c} \prod_{\mathbf{z} \in C} \kappa_c(\mathbf{z}), \quad (3.21)$$

such that $\lambda_c \triangleq \lambda_c(t_k) \triangleq \langle \kappa_c, 1 \rangle$ represents the expected number of clutter measurements at time t_k .

C. Complete multi-object measurement likelihood

Combining the two measurement sources, and considering the likelihood of all possible association scenarios in a single MO measurement model leads to the measurement scan likelihood

$$g_k(Z_k | \dot{X}_k) = \underbrace{\pi_c(Z_k)}_{\text{all clutter}} \cdot \underbrace{g_k^o(\emptyset | \dot{X}_k)}_{\text{all undetected}} \cdot \sum_{\theta_k \in \Theta_k} \overbrace{\left[\prod_{\substack{\dot{x} \triangleq (\mathbf{x}, \ell): \\ \theta_k(\ell) > 0}} \frac{P_D(\dot{x}) g_k(z_k^{\theta_k(\ell)} | \mathbf{x})}{\kappa_c(z_k^{\theta_k(\ell)}) (1 - P_D(\dot{x}))} \right]}_{\text{detected (so not clutter or undetected)}}, \quad (3.22)$$

where $\pi_c(Z_k)$ is the probability that all measurements are false positives – computed by [Eq. \(3.21\)](#) – and $g_k^o(\emptyset | \dot{X}_k) = \prod_{\dot{x} \in \dot{X}} (1 - P_D(\dot{x}))$ is the probability that none of the objects are detected. All other options are present in the sum over the possible association hypotheses Θ_k , where the denominator is used to cancel out all unnecessary contributions to the first two terms. As such, any combination of detected, misdetections and clutter cases are accounted for in this model. This formulation of the likelihood is very useful from an implementation point of view (see [Section 3.3.3](#)).

D. Sampling measurements from the standard measurement model

For the purpose of testing filter performance on theoretical cases, it is often useful to generate simulated measurements. [Algorithm 3.2](#) describes how this can be done while keeping consistency with the standard measurement model in [Eq. \(3.22\)](#).

Algorithm 3.2: Sampling measurements from the standard measurement model

```

MEAS SAMPLE(likelihood  $g_k(Z | \dot{X}_k; p_D, \kappa_c)$ ):
1   $O_k = C_k = \emptyset$ 
2  for  $\dot{x}_k$  in  $\dot{X}_k$  ▷ sample observations
3      sample  $u \sim \mathcal{U}([0, 1])$ 
4      if  $u < p_D(\dot{x}_k, t_k)$ 
5          sample  $\mathbf{o}_k^{(\ell)} \sim g_k(\cdot, \mathbf{x}_k)$ 
6           $O_k = O_k \cup \{\mathbf{o}_k\}$ 
7      end if
8  end for
9   $\lambda_c = \int_{\mathbb{Z}} \kappa_c(z) dz$  ▷ clutter rate
10 sample  $|C_k| \sim \text{Po}(\lambda_c)$  ▷ clutter cardinality
11 for  $i$  in  $\{1, \dots, |C_k|\}$  ▷ sample clutter
12     sample  $\mathbf{c}_k \sim \kappa_c(\cdot, t_k)$ 
13      $C_k = C_k \cup \{\mathbf{c}_k\}$ 
14 end for
15  $Z_k = O_k \cup C_k$ 
16 return  $Z_k$ 
    
```

3.3.3 Optimal assignment

One of the key bottlenecks in MO trackers occurs in the step of measurement association. Since the complexity of the last term in [Eq. \(3.22\)](#) scales combinatorially with the number of measurements and labels, the likelihood function is intractable for large-scale tracking problems

such as SSA. The Bayesian update must therefore be truncated in some meaningful manner. The practice of finding the hypotheses in Θ_k that contribute the most to the MO likelihood can be formulated as an optimal assignment problem.

To achieve this, let $L_k^{\theta_k} \triangleq \log(g_k^{\theta_k}(Z_k | \hat{X}_k))$ be the log likelihood of observing the measurements Z_k based on the prior \hat{X}_k and *conditioned* on specific association mapping $\theta_k \in \Theta_k$ [Eq. \(3.18\)](#):

$$L_k^{\theta_k} = \overbrace{\log[\pi_c(Z_k)]}^{\text{all clutter}} + \overbrace{\log[g_k^o(\emptyset | \hat{X}_k)]}^{\text{all undetected}} + \overbrace{\sum_{\substack{\hat{x} \triangleq (\mathbf{x}, \ell) \\ \theta_k(\ell) > 0}} \psi_{\ell}^{\theta_k(\ell)}}^{\text{assignment score}}, \quad (3.23a)$$

$$\text{with } \psi_{\ell}^{\theta_k(\ell)} = \begin{cases} \log \left[\frac{P_D(\hat{\mathbf{x}}) P_G g_k(\mathbf{z}_k^{\theta_k(\ell)} | \mathbf{x})}{\kappa_c(\mathbf{z}_k^{\theta_k(\ell)}) (1 - P_D(\hat{\mathbf{x}}))} \right] & \text{if } \theta_k(\ell) > 0 \\ 0 & \text{if } \theta_k(\ell) = 0 \text{ (undetected)} \end{cases}, \quad (3.23b)$$

where P_G is the gating probability (discussed further in [Section 3.3.3-B](#)). If gating is not applied, $P_G = 1$ and this factor disappears. Since θ_k implies every label causes at most 1 measurement, the association is an assignment problem where the total score is $L_k^{\theta_k}$. Note that taking the logarithm does not change the order of the scores, but is commonly used to avoid numerical problems since the score values are typically very low.

Example

The assignment problem can be visualised in a 2D table like [Figure 3.2](#), where 4 track labels form the rows and the first 5 columns correspond to the measurements. The association score from [Eq. \(3.23b\)](#) is then computed for each of the target-measurement associations. To represent the possibility of missed detections, an additional column is added for each target with one value set to the score $\psi_{\ell_i}^0 = 0$ and all others set to $-\infty$, since they are not allowable associations.

After assigning each target (row) to a different column, the complete multi-object measurement log-likelihood is obtained by summing all the relevant scores and adding the constant terms from [Eq. \(3.23a\)](#). The benefit of formulating the problem like this is that existing algorithms from other fields can be used to find the assignments with the highest scores.

Again, notice that false-positive observations are modelled implicitly. Any measurement to which no target is assigned, must be associated to the clutter process. Since the assignment scores have the clutter intensity in the denominator, the clutter contributions from all assigned measurements cancel out and what remains of the first constant term in [Eq. \(3.23a\)](#) correctly represents likelihood that all unassigned observations are false-positives.

	z_1	z_2	z_3	z_4	z_5	mis_1	mis_2	mis_3	mis_4
ℓ_1	$\psi_{\ell_1}^1$	$\psi_{\ell_1}^2$	$\psi_{\ell_1}^3$	$\psi_{\ell_1}^4$	$\psi_{\ell_1}^5$	$\psi_{\ell_1}^0$	$-\infty$	$-\infty$	$-\infty$
ℓ_2	$\psi_{\ell_2}^1$	$\psi_{\ell_2}^2$	$\psi_{\ell_2}^3$	$\psi_{\ell_2}^4$	$\psi_{\ell_2}^5$	$-\infty$	$\psi_{\ell_2}^0$	$-\infty$	$-\infty$
ℓ_3	$\psi_{\ell_3}^1$	$\psi_{\ell_3}^2$	$\psi_{\ell_3}^3$	$\psi_{\ell_3}^4$	$\psi_{\ell_3}^5$	$-\infty$	$-\infty$	$\psi_{\ell_3}^0$	$-\infty$
ℓ_4	$\psi_{\ell_4}^1$	$\psi_{\ell_4}^2$	$\psi_{\ell_4}^3$	$\psi_{\ell_4}^4$	$\psi_{\ell_4}^5$	$-\infty$	$-\infty$	$-\infty$	$\psi_{\ell_4}^0$

Figure 3.2: Example assignment score matrix with potential assignment mapping $\theta_k : \{\ell_1 \mapsto 5, \ell_2 \mapsto 1, \ell_3 \mapsto 0, \ell_4 \mapsto 3\}$ indicated in grey. This implies that the object with label ℓ_3 is undetected at time t_k and z_2 and z_4 are clutter measurements.

A. Auction and Murty for n -best optimal assignment

A variety of algorithms exist that can find the optimal assignment mapping for a score matrix like in Figure 6 with varying computational complexity [9][11][66][72].

However, finding the single best assignment (hypothesis) is not sufficient. As mentioned before, the goal is to consider the weighted contributions from multiple hypotheses. Although the theoretical formulation just considers *all* hypotheses, this is intractable for realistic numbers of targets and measurements. Moreover, the large majority of assignments usually has a negligible likelihood and therefore almost no influence on the final result of the filter update. It is therefore desirable to select the n best assignments for the weighted update and ignore all the others. The result will be an approximation of the corrected MO density before actually computing it. Vo and Vo [135, Sec. IV-C] showed that this type of *truncation* minimises the L_1 distance for the δ GLMB corrected density, which is used in this work (see Section 3.3.4–C).

Solving this *ranked assignment* problem is often done by Murty’s algorithm [91], which iteratively uses an optimal assignment solver (here: Auction) and then removes that solution from the search space until the desired number of solutions is extracted [11]. For this thesis, a version of the Auction algorithm [9] is implemented, following the description by Blackman and Popoli [11, Sec. 6.5] and the implementation of Murty’s algorithm was based on a paper by Cox and Hingorani [25, p. 143].

B. Measurement gating

Since the association step is notoriously one of the computational bottlenecks in MOT – Murty’s algorithm has about $\mathcal{O}[n \cdot |L| \log |L| \cdot (|Z| + |L|)]$ complexity when using Auction assignment algorithm –, it is unfeasible to use it as-is for larger target and measurement sets. In such cases it is better to eliminate certain associations a-priori, based on the statistical distance between expected and observed measurements. This is often called *gating*.

Given a Gaussian measurement model $h : \mathbb{X} \rightarrow \mathbb{Z}$ where $\mathbb{Z} \subseteq \mathbb{R}^m$, every prior state estimate (\mathbf{x}, ℓ) results in an estimated measurement $h(\mathbf{x})$, and an innovation covariance $\mathbf{P}_{\Delta z}$ (recall Section 2.2.4). Given this probability distribution, the squared *Mahalanobis distance* Eq. (2.19) between $\mathcal{N}(h(\mathbf{x}), \mathbf{P}_z(\mathbf{x}))$ and any measurement \mathbf{z} originating from object ℓ will be χ^2 distributed with m degrees of freedom. One can then define a gate with *gating probability* P_G , that accepts measurements \mathbf{z} for which

$$d_M^2(\mathbf{z}; h(\mathbf{x}), \mathbf{P}_z(\mathbf{x})) < Q_{\chi^2}(P_G; m), \quad (3.24)$$

where $Q_{\chi^2}(\cdot; m)$ is the *inverse cumulative density function* or *quantile function* of the χ^2 distribution for m degrees of freedom. This means that there is a $1 - P_G$ chance that a correct measurement falls outside the gate.

The result of gating can be implemented by setting the assignment score in Figure 3.2 to $-\infty$ for the measurements that do not pass a target's gate. This immediately discards assignment mappings where any rejected association is included.


C. Assignment grouping

After gating, it is possible to reduce the dimensionality of the measurement association problem by *grouping* those targets whose gates admitted the same measurements. Intuitively, it follows that if two disjoint measurement sets pass the gates of two disjoint target sets, the ranked assignment problem may be split into two separate assignment problems that can be solved in parallel. This is illustrated below for the example from Figure 3.2. Reuter et al. [104] show that grouping before the filter update introduces negligible errors provided that the gating threshold is sufficiently large.

Example (cont. from p. 26)

Figure 3.3 visualises how the process of gating and grouping splits the dense assignment matrix from Figure 3.2 into two smaller problems of reduced complexity. Completely disjoint gates form separate groups (i.e. assignment matrices) and within groups, various assignments might be disallowed, resulting in sparse assignment matrices.

Gating in measurement space:



Assignment matrix 1:

	z_1	z_3	mis_1	mis_3	mis_4
ℓ_1	$\psi_{\ell_1}^1$	$-\infty$	$\psi_{\ell_1}^0$	$-\infty$	$-\infty$
ℓ_3	$-\infty$	$\psi_{\ell_3}^3$	$-\infty$	$\psi_{\ell_3}^0$	$-\infty$
ℓ_4	$\psi_{\ell_4}^1$	$\psi_{\ell_4}^3$	$-\infty$	$-\infty$	$\psi_{\ell_4}^0$

Assignment matrix 2:

	z_2	z_5	mis_2
ℓ_2	$\psi_{\ell_2}^2$	$\psi_{\ell_2}^5$	$\psi_{\ell_2}^0$

Figure 3.3: Example of grouping and gating with the resulting reduced assignment score matrices compared to Figure 3.2.

3.3.4 LMB filter implementation

The main part of this work uses the LMB filter for MOT, so this section briefly summarises the key elements in that algorithm as proposed by Reuter et al. [104] with a focus on how they are implemented in BASIL for this thesis. The reader is referred to the original paper for details and derivations. A schematic overview of the LMB recursion is provided in Figure 3.4, referencing the relevant subsections where each step is further addressed.

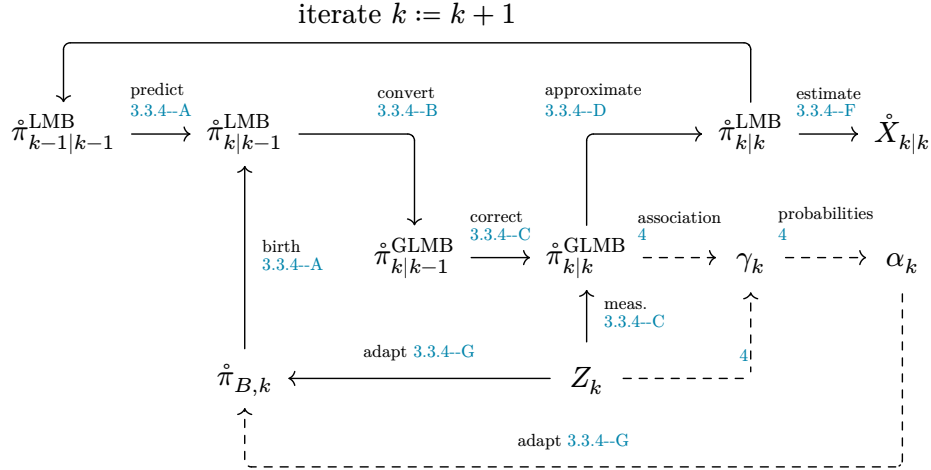


Figure 3.4: LMB filter schematic overview. The links direct to the section where each step is addressed. Note that everything related to association extraction (here shown with dashed lines) is not part of the original LMB filter.

To provide a more intuitive insight into the procedure, all steps in a single iteration are illustrated for a very simple example starting on p. 35.

Let $\hat{\pi}_{k-1}$ be a prior LMB multi-object density, where the single-object state densities are restricted to be labelled Gaussian mixtures $\mathcal{G}_{k-1}^{(\ell)}$ (recall [Section 2.1.2](#)) for the purpose of this work⁷:

$$\hat{\pi}_{k-1} = \left\{ \left(r_{k-1}^{(\ell)}, \mathcal{G}_{k-1}^{(\ell)} \right) \right\}_{\ell \in \mathbb{L}_{k-1}}, \quad (3.25a)$$

$$\text{with } \mathcal{G}_k^{(\ell)} \triangleq \left\{ \mathcal{E}_k^{(\ell,i)} \triangleq \left(w_k^{(\ell,i)}, \mathbf{x}_k^{(\ell,i)}, \mathbf{P}_k^{(\ell,i)} \right) \right\}_{i=1}^{|\mathcal{G}_k^{(\ell)}|}. \quad (3.25b)$$

A. Prediction

Since the LMB is closed under the Chapman-Kolmogorov equation [Eq. \(3.17\)](#), the prediction step consists of (1) predicting the components of the previously existing Gaussian mixtures, (2) adjusting the existence probabilities to account for potential target disappearance (*death*) and (3) adding the components of an LMB birth density $\hat{\pi}_{B,k}$:

$$\hat{\pi}_{k|k-1} = \left\{ \left(r_{k|k-1}^{(\ell)}, \mathcal{G}_{k|k-1}^{(\ell)} \right) \right\} = \hat{\pi}_{B,k} \cup \hat{\pi}_{S,k} \quad (3.26a)$$

$$\text{with } \hat{\pi}_{S,k} = \left\{ \left(r_{k|k-1}^{(\ell)}, \mathcal{G}_{k|k-1}^{(\ell)} \right) \right\}_{\ell \in \mathbb{L}_{k-1}}, \quad (3.26b)$$

where $r_{k|k-1}^{(\ell)} = P_{S,k}(\hat{\mathbf{x}})r_{k-1}^{(\ell)}$ and the mixtures $\mathcal{G}_{k|k-1}^{(\ell)}$ are predicted using [Algorithm 2.4](#). Note that the birth process causes the label space to be extended to $\mathbb{L}_k = \mathbb{L}_{k-1} \uplus \mathbb{B}_k$. To limit the computational complexity, the implementation sets configurable capping and pruning thresholds that removes the components with lowest existence probability.

⁷This GM limitation is not necessary for the LMB filter, but it is chosen for simplicity and to reflect the present implementation. The notation here does assume that the GM is closed under both the prediction and Bayes update, which is not true in general, but holds for the Gaussian simplifications from [Section 2.2.4](#).

B. Conversion from LMB to GLMB

As Reuter et al. [104][135] pointed out, the LMB is not conjugate prior, meaning that the result of the Bayes update is not an LMB, but rather a GLMB. For this reason, an implementation of the LMB filter first needs to reformulate the predicted density as δ GLMB. Recall from Eq. (3.14) that the δ GLMB contains weighted hypotheses where each hypothesis (ξ, L) is defined by a set of existing labels L and a measurement association history ξ . The PDF of every target is conditioned on ξ and the joint existence probability of the labels in L determines the weight. Since there is no notion of association history in the LMB, ξ is omitted here and the hypotheses are formed by computing the *joint existence probability* $w(L)$ for various label sets $L \subseteq \mathbb{L}_k$.

The naive approach is to create one hypothesis for every possible subset of the labels represented by the LMB density:

$$\hat{\pi}_{k|k-1}^{\text{GLMB}} = \left\{ \left(w(L), p_{k|k-1}^{\mathcal{G}} \right) \right\}_{L \subseteq \mathbb{L}_k}, \text{ where } p_{k|k-1}^{\mathcal{G}}(\cdot, \ell) = \mathcal{G}_{k|k-1}^{(\ell)}. \quad (3.27)$$

However, this leads to $2^{|\mathbb{L}_k|}$ prior hypotheses, which is infeasible for most realistic scenarios. Since checking all tracks and keeping only the ones with highest weight is still a task of exponential complexity, a more efficient solution is needed. One option that is used by Vo and Vo [136] is to translate this into a graph problem, such that it can be solved using *n-shortest path* algorithms like Yen [145] or Eppstein [32], which are significantly less demanding than the exhaustive search approach. The procedure to transform the problem is shown in Algorithm 3.3 and illustrated for a small LMB of 3 components in Figure 3.5.

Algorithm 3.3: Converting LMB to δ GLMB using graph shortest path formulation

```

MAKE GLMB( $\hat{\pi}_{k|k-1}^{\text{LMB}}, n$ ):
1   $\mathbb{L}_k$  label set in  $\hat{\pi}_{k|k-1}^{\text{LMB}}$ 
2   $\ell_1, \dots, \ell_{|\mathbb{L}_k|}$  s.t.  $\forall i < j : r_k^{(\ell_i)} \geq r_k^{(\ell_j)}$  ▷ sort labels by descending  $r_k^{(\ell)}$ 
3   $N = \{\text{entry}, \text{exit}\}$  ▷ initialise nodes
4   $E = \{(\text{entry} \rightarrow \text{exit}, 0)\}$  ▷ initialise edges
5  for  $1 \leq i \leq |\mathbb{L}_k|$ 
6       $N = N \cup \{\ell_i\}$  ▷ add node
7       $E = E \cup \{(\ell_i \rightarrow \text{exit}, 0)\}$  ▷ add 0-weight edge
8       $c^{(\ell_i)} = -\log\left(\frac{r_k^{(\ell_i)}}{1-r_k^{(\ell_i)}}\right)$  ▷ compute weight
9      for  $i < j < |\mathbb{L}_k|$ 
10          $E = E \cup \{(\ell_i \rightarrow \ell_j, c^{(\ell_i)})\}$  ▷ add edge
11     end for
12 end for
13  $G = (N, E)$  ▷ Create graph
14  $\mathcal{P}_{\text{entry} \rightarrow \text{exit}} \xleftarrow{\text{Eppstein}} n$  shortest paths given  $G$  ▷ from entry to exit node
15  $\log[w(\emptyset)] = \sum_{\ell \in \mathbb{L}_k} (1 - r_k^{(\ell)})$  ▷ weight that none exist
16  $\hat{\pi}_{k|k-1}^{\text{GLMB}} = \emptyset$ 
17 for  $\mathcal{P} \in \mathcal{P}_{\text{entry} \rightarrow \text{exit}}$ 
18      $L^{(\mathcal{P})} = \{\ell \in \mathcal{P} \mid \ell \in \mathbb{L}_k\}$  ▷ label set included in the path
19      $w^{(\mathcal{P})} = \exp\left[-\left(\sum_{\ell \in L^{(\mathcal{P})}} c^{(\ell)} + \log[w(\emptyset)]\right)\right]$  ▷ correct weight

```

20	$\hat{\pi}_{k k-1}^{\text{GLMB}} = \hat{\pi}_{k k-1}^{\text{GLMB}} \cup \{((w^{\mathcal{I}})p^{\mathcal{G}})\}$	▷ Add to GLMB
21	end for	
22	return $\hat{\pi}_{k k-1}^{\text{GLMB}}$	

Example

This example visualises how the GLMB-to-LMB conversion is translated to a shortest path problem for a small case with 3 labels. In Figure 3.5, the symbols $r_i \triangleq r^{(\ell_i)}$ and $q_i \triangleq 1 - r^{(\ell_i)}$ are used for simplicity, representing the probability of (non-)existence for each of the target labels ℓ_i . On the left, the exhaustive list of $2^3 = 8$ hypotheses is shown, where a circle is blue if the label is included in the label set L_j and red if it is not. The hypothesis weight $w(L_j)$ is the product of the corresponding probabilities. For example, the fourth row represents $L_4 = \{\ell_2, \ell_3\}$ and the weight is $w(\{\ell_2, \ell_3\}) = (1 - r_1) \cdot r_2 \cdot r_3$, i.e. the probability that both ℓ_2 and ℓ_3 are included, and that ℓ_1 is excluded.

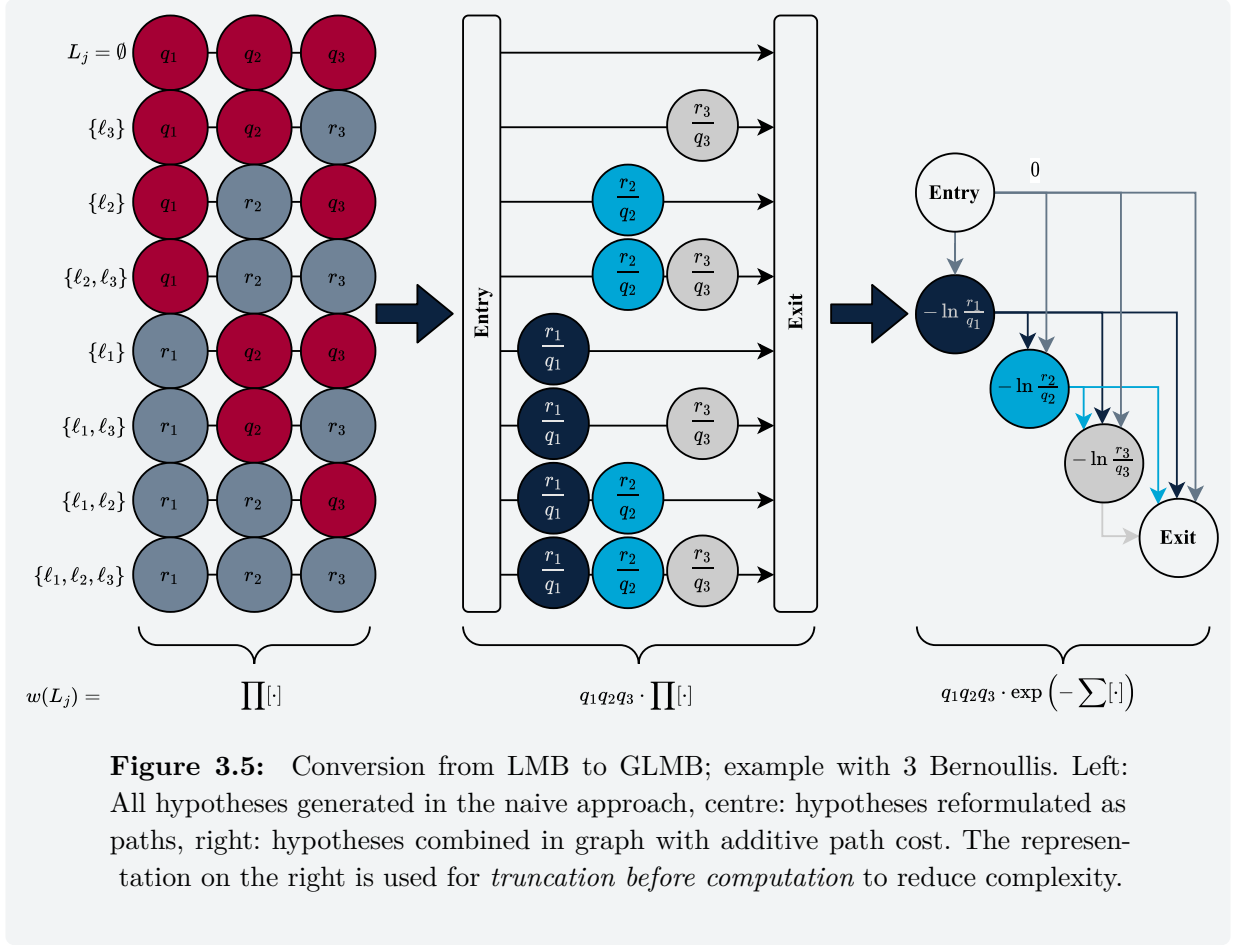
In the second stage (middle of figure Figure 3.5), the exclusion of labels is modelled implicitly, by factoring out q_i for all labels. Further, the *entry* and *exit* nodes are added so that every hypothesis can be interpreted as a *path*. Notice that the nodes for each label are always $\frac{r_i}{q_i}$ (indicated by a different colour for every label), so this middle step is really an unfolded representation of a *directed graph*, where every path from *entry* to *exit* represents a possible label set. The corresponding hypothesis weight is then the product of all included nodes and the constant $\prod_i q_i$. Considering again the fourth row, the weight has not changed but is now computed as $w(\{\ell_2, \ell_3\}) = q_1 q_2 q_3 \frac{r_2}{q_2} \frac{r_3}{q_3}$.

In the last stage, the directed graph is collapsed and the logarithm is taken on all the node contributions to transform the product to a sum. By assigning the value in the starting node to each of the edges, one obtains a *weighted, directed graph* and the objective of finding the most likely label sets has become a “longest path” problem. Since most available algorithms instead minimise path length, the edge weights are negated to obtain a *shortest path problem*. Once more addressing the 4th label hypothesis, it is represented in the graph by a 0-weight jump from *entry* to the ℓ_2 node and then passing through ℓ_3 to arrive at *exit*. The corresponding weight is again unchanged but computed by

$$\log[w(\{\ell_2, \ell_3\})] = \overbrace{\log(q_1) + \log(q_2) + \log(q_3)}^{\text{constant}} \quad (3.28a)$$

$$- \overbrace{(-\log(r_2) - \log(r_3) + \log(q_1) + \log(q_2) + \log(q_3))}^{\text{path length (minimise)}}. \quad (3.28b)$$

In this form, the most probable label sets can be found efficiently using established graph algorithms. It should be noted that the approach results in improved scalability and therefore reduced run times for large numbers of labels, but does not provide any speed-up in the 3-component case.



In this thesis, the algorithm from Eppstein [32] is used, resulting in a complexity of $\mathcal{O}(l \log(l) + m + n)$ where $l = (|L_k| + 2)$, $m = \frac{l(l+1)}{2}$ and n is the number of shortest paths to find. For illustration, a label set of 20 objects would already require over 10^6 operations with the naive approach, whereas the graph implementation is still $\mathcal{O}(10^3)$.

C. GLMB Bayes update

In the δ GLMB form, one can perform the Bayes update for the MO density, i.e. all prior hypotheses in $\mathcal{H}_+ \triangleq \mathcal{H}_{k|k-1}$ are evaluated independently, generating new posterior hypotheses in $\mathcal{H}_k \triangleq \mathcal{H}_{k|k}$ for every measurement assignment option. Since this would lead to hypothesis explosion, the *truncation before computation* properties of the δ GLMB [135] are used to limit the complexity. This correction step is represented schematically in Figure 3.6 and can be summarised as follows:

- (i) If the measurement set is empty, this means that all existing objects were undetected. The posterior hypotheses are the same as the prior hypotheses, but their weights are scaled by the probability of all missed detections and no clutter.
- (ii) Otherwise, the update for every prior state and measurement combination is precomputed⁸.
- (iii) For every prior hypothesis (each with a different label set), an n -best assignment problem is solved using Murty's algorithm (see Section 3.3.3) and the appropriate weights are computed to reflect the posterior joint existence probabilities conditioned on the prior hypotheses.

⁸This can only be done because in this case, the GLMB is derived from the LMB, so there is only one SO distribution per label. For the general GLMB, this is intractable.

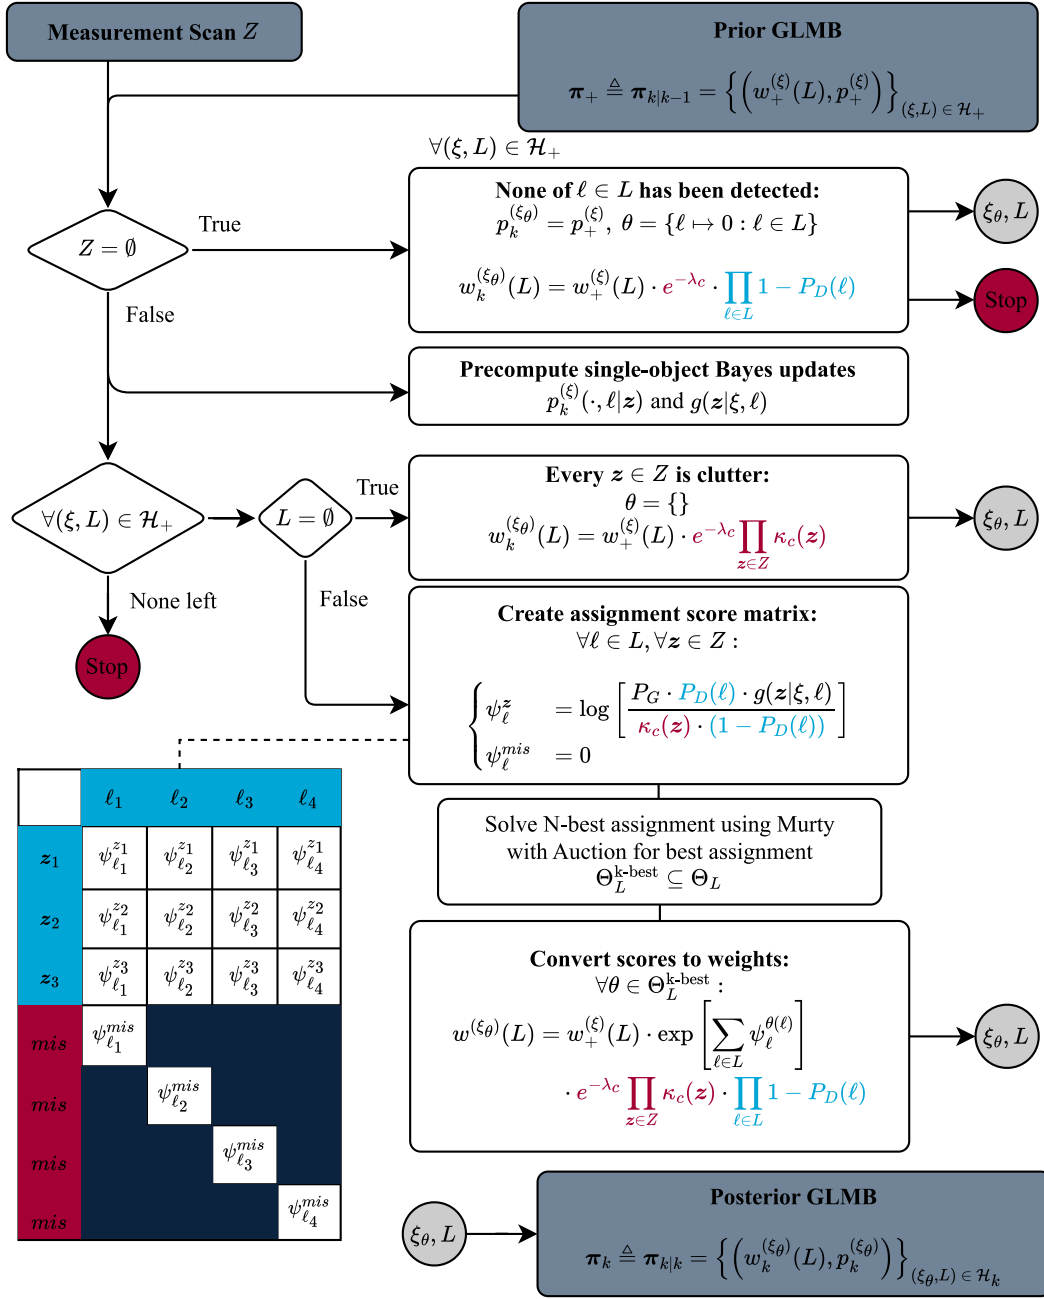


Figure 3.6: Schematic representation of the GLMB update. The parts highlighted in red represent the contributions of clutter Eq. (3.21) and blue is related to detection probability. The formulation of the joint existence probabilities corresponds to the MO measurement likelihood from Eq. (3.22). The yellow circles indicate all hypotheses forming the posterior GLMB.

Note that this step lends itself very well to parallelisation. Each of the prior hypotheses can be updated entirely independently, after which all (mutually disjoint) posterior hypothesis sets merge to form the posterior δ GLMB.

D. Approximating the posterior GLMB

After the GLMB update, the correction step of the LMB filter is completed by approximating the posterior δ GLMB with an LMB density. Reuter et al. [104] showed that this is essentially a moment matching process, resulting in an identical unlabelled PHD (first moment) and the exact same decomposition into individual tracks.

Figure 3.7 illustrates how the existence probability for each labelled Bernoulli can be computed from the posterior δ GLMB hypotheses. This results in the LMB formulation

$$\hat{\pi}_k(\cdot | Z_k) = \left\{ \left(r_k^{(\ell)}, \mathcal{G}_k^{(\ell)}(\cdot | Z_k) \right) \right\}, \quad (3.29a)$$

with

$$\begin{cases} r_k^{(\ell)} &= \sum_{\substack{(\xi, L) \in \mathcal{H}_k \\ \ell \in L}} w_k^{(\xi)}(L) \\ \mathcal{G}_k^{(\ell)}(\cdot | Z_k) &= \frac{1}{r_k^{(\ell)}} \left[\sum_{(\xi, L) \in \mathcal{H}_k} w_k^{(\xi)}(L) p_k^{(\xi)}(\cdot, \ell) \right] \end{cases}$$

Notice that the labelled densities are now a weighted superposition of all the association hypotheses, so this step erases the information on which measurement contributed most to each of the target state updates. A solution to this problem is proposed in Chapter 4.

E. LMB reduction

It was already discussed in Section 3.3.1 and Section 3.3.4–A that new components – representing new objects – can be added to the LMB filtering density by means of a *birth process*. However, depending on how this is modelled, many new targets might be created in every step with a relatively low existence probability. If there are no later measurements that *confirm* a new label represents a real object, then the component can be removed to reduce unnecessary computations. This can be done by LMB *pruning* and *capping*, similar to the GM reduction methods with the same names (Section 2.1.2). Given a pruning threshold $\vartheta_P^{\text{LMB}} \in [0, 1]$ and capping parameter $\vartheta_C^{\text{LMB}} \in \mathbb{N}^+$, the original LMB $\hat{\pi}_o^{\text{LMB}}$ can be reduced to

$$\hat{\pi}_r^{\text{LMB}} \subseteq \hat{\pi}_o^{\text{LMB}} \quad (3.30a)$$

$$\text{s.t.} \quad |\hat{\pi}_r^{\text{LMB}}| \leq \vartheta_C \quad (3.30b)$$

$$\text{and} \quad \forall \ell_1 \in \mathcal{L}(\hat{\pi}_r^{\text{LMB}}), \ell_2 \in \mathcal{L}(\hat{\pi}_o^{\text{LMB}} \setminus \hat{\pi}_r^{\text{LMB}}) : r^{(\ell_1)} > r^{(\ell_2)} \wedge r^{(\ell_1)} > \vartheta_P^{\text{LMB}} \quad (3.30c)$$

F. State extraction

The final step of the filter is to extract deterministic state estimates from the corrected LMB density. This can be seen as the filter’s “best guess” of the MO state, based on all the information it has received. There are multiple possible estimators for the LMB, but the default BASIL implementation takes the cardinality \hat{n}_k as the maximum a-posteriori (MAP) estimate of its distribution

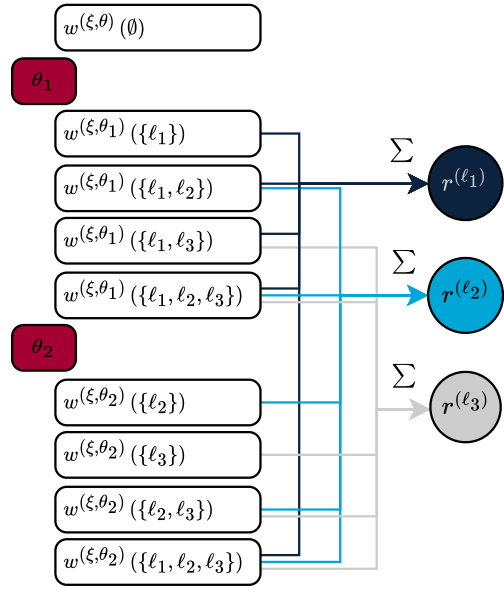


Figure 3.7: Computing existence probabilities for LMB approximation using a very simple example with a total of 3 labels and two association hypotheses.

$$\hat{n}_k = \arg \max_n [\rho_k^{\text{LMB}}(n)], \quad (3.31)$$

then chooses the \hat{n}_k labels with the highest existence probability

$$\hat{L}_k \subseteq \mathbb{L}_k \text{ with } |\hat{L}_k| = \hat{n}_k, \text{ and } \forall \begin{cases} \ell_i \in \hat{L}_k \\ \ell_j \in \mathbb{L}_k \setminus \hat{L}_k \end{cases} : r^{(\ell_i)} > r^{(\ell_j)}, \quad (3.32)$$

and estimates each of the corresponding states as the expected value of their SO distributions

$$\hat{X}_k = \{(\mathbb{E}_{\hat{p}(\cdot, \ell)}[\mathbf{x}], \ell) : \ell \in \hat{L}_k\}. \quad (3.33)$$

Reuter et al. [104, p. 10] use an alternative that passes each label with existence probability above a lower threshold ϑ_l , as long as it has once exceeded a higher threshold ϑ_u ,

$$\hat{X}_k = \{(\hat{x}, \ell) : r_{\max}^{(\ell)} > \vartheta_u \wedge r_k^{(\ell)} > \vartheta_l\}, \quad \text{with } \vartheta_l < \vartheta_u. \quad (3.34)$$

However, this introduces two more tuning parameters that can be avoided with the other approach. This last tuning argument is the only reason to prefer the estimator based on MAP cardinality as the default BASIL implementation. However, for cases like SSA where target death is very uncommon, this threshold alternative may be relatively intuitive and well-behaved. Extensive estimator comparison is left outside the scope of the present work.

Example (cont. from p. 13)

Using the same 1D linear example as before (first introduced on p. 11), suppose there are now a finite but unknown number of trains on parallel rails, so that there is no interaction between them. If there is previous knowledge for the position and velocity of one train, this might be represented by the LMB in Figure 3.8. The prediction step propagates this forward with $P_S = 0.98$ and some birth process generates a new label (another potential train); The union of those surviving and newborn components is the predicted LMB (Figure 3.10).

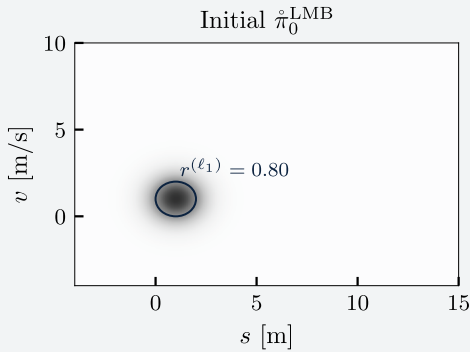


Figure 3.8: 1D linear example – initial LMB with 1 component. The ellipses in all these plots indicate the $d_M = 1$ covariance bound of a Gaussian distribution.

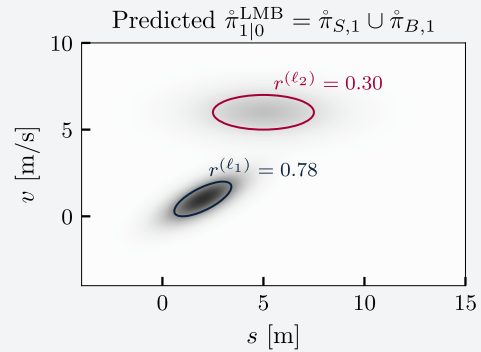


Figure 3.9: 1D linear example – predicted LMB with 1 surviving component (blue) and 1 birth component (red).

This predicted LMB is then represented as GLMB in Figure 3.10 with 4 hypotheses – one for each label subset $L \subseteq \mathbb{L}_1$. The LMB existence probabilities $r^{(\cdot)}$ are incorporated in the hypothesis weights $w(L)$ (hence the name *joint existence probability*). Each prior hypothesis

gives rise to a predicted measurement set, which is to be associated with the single observation in the measurement scan $Z_1 = \{z^{(m_1)}\}$.

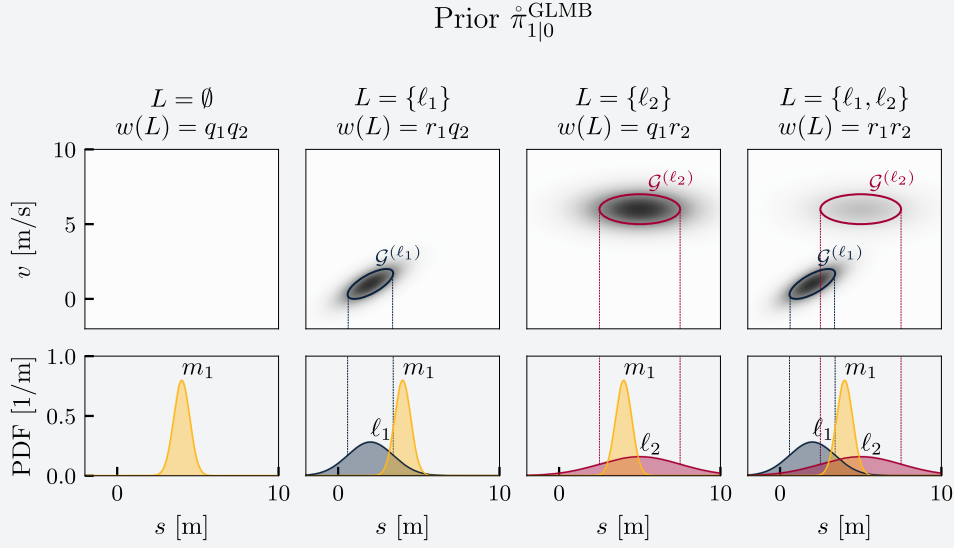


Figure 3.10: 1D linear example – prior GLMB and predicted measurement distributions for ℓ_1 and ℓ_2 . The yellow Gaussian is the observed measurement distribution.

The set of all possible association solutions then constitutes the corrected GLMB, shown in [Figure 3.11](#). The weights are omitted here but can be computed by multiplying the prior hypothesis weights $w(L)$ by the association likelihoods (recall [Eq. \(3.23\)](#) and [Figure 3.6](#)).

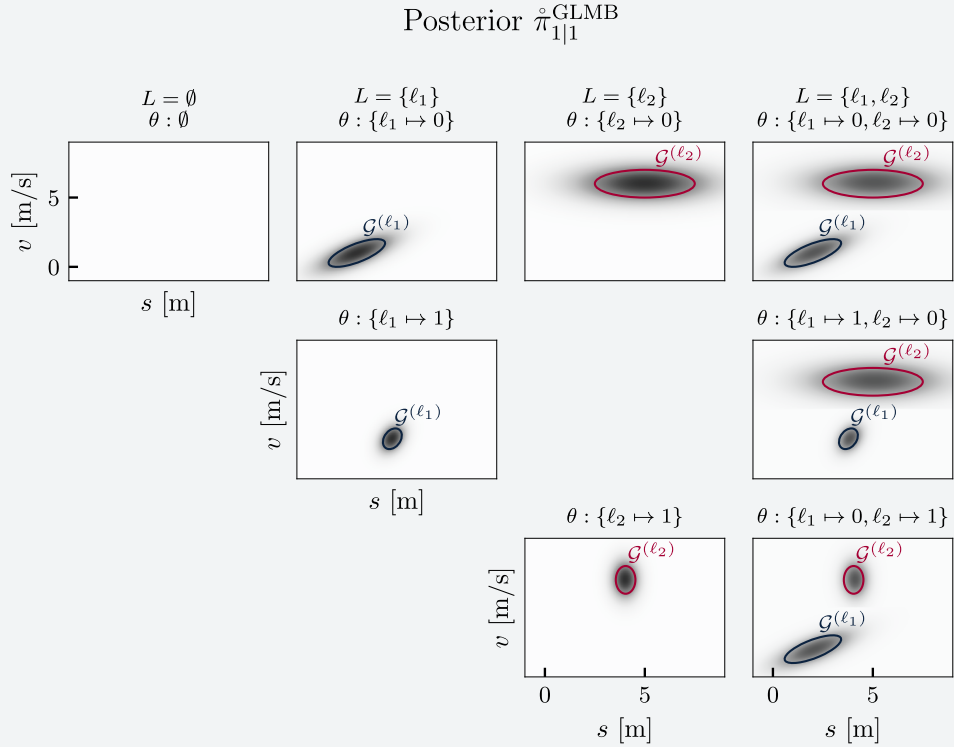


Figure 3.11: 1D linear example – prior GLMB and predicted measurements

The weighted superposition of all hypotheses as described in Eq. (3.29) then results in the LMB approximation of Figure 3.12 and extracting the mean states for the most likely cardinality gives the estimated state (Figure 3.13).

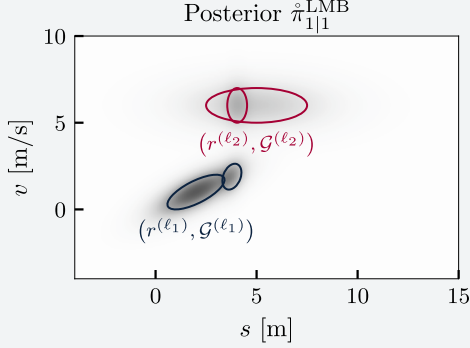


Figure 3.12: 1D linear example – posterior LMB with the weighted contributions of all association hypotheses




Figure 3.13: 1D linear example – estimated multi-object state

If Figure 3.12 is now substituted for Figure 3.8, the posterior LMB becomes the new prior and another iteration can be performed for the next measurement scan.

G. Adaptive birth

Recall from Section 3.3.1 that introducing new labels into the MO filter is part of the prediction step. However, it is possible to make the birth process depend on the observed measurements at the previous time step, using the association information in the GLMB update step (recall Figure 3.6).

First, the probability is computed that a measurement \mathbf{z} was not associated to any of the previously existing labels: [104][134]

$$r_{U,k-1}(\mathbf{z}) = 1 - \sum_{(\xi_\theta, L) \in \mathcal{H}_k} \mathbf{1}_\theta(\mathbf{z}) w_k^{(\xi_\theta)}(L), \quad (3.35)$$

where the indicator function $\mathbf{1}_\theta[\mathbf{z}]$ checks if a label was assigned to \mathbf{z} by the mapping θ . As addressed in Section 3.3.2, an unassociated measurement could either be a false positive (clutter) or it stems from an unknown object. To distinguish between the two, Reuter et al. [104] assume the number of birth components at time t_k is Poisson distributed with mean $\lambda_{B,k}$ and then impose a maximum $r_{B,\max}$ to compute the existence probabilities for birth components from each measurement

$$r_{B,k}(\mathbf{z}) = \min \left(r_{B,\max}, \frac{r_{U,k-1}(\mathbf{z})}{\sum_{\xi \in Z_{k-1}} r_{U,k-1}(\xi)} \lambda_{B,k} \right). \quad (3.36)$$

However, this doesn't work well if the birth cardinality does not really follow a Poisson distribution.

Example

If an MO filter is used to discover a population of objects and estimate the trajectories without any prior information (as will be the case for several test scenarios in [Chapter 7](#)), the expected number of birth components depends on how many objects are in the population, how many have already been found and how many are observable in a particular time step. The birth cardinality in this situation cannot be accurately modelled as Poisson.

Instead, this thesis uses the expected ratio between birth and clutter $\frac{\lambda_B}{\lambda_c}$ to scale the existence probabilities,

$$r_{B,k}(z) = \min \left(r_{B,\max}, r_{U,k-1}(z) \frac{\lambda_{B,k}}{\lambda_{c,k-1}} \right). \quad (3.37)$$

This still leaves a seemingly arbitrary parameter, but it is found to be easier to tune in the SSA case, where the expected number of born targets depends among others on the extent of the existing catalogue, and the sensitivity and position of the sensor.

To the author's knowledge, this approach is not previously used in the literature and it is worth discussing briefly why that might be the case. The tuning convenience of this birth-to-clutter ratio $\frac{\lambda_B}{\lambda_c}$ comes at the cost of an important theoretical weakness. This parameter couples the birth process to the clutter model, which is in principle not desirable. By definition, a clutter measurement is a false positive and thus does not belong to any real object. [Eq. \(3.37\)](#) makes birth probabilities dependent upon the clutter model. Notice that the original version in [Eq. \(3.36\)](#) decouples the two by first normalising the probabilities of non-association $r_{U,k-1}$ and therefore does not have the same issue.

The final step is to create an initial SO density from the information in the measurement. Since the state is usually only partially observable, this requires some assumptions depending on the application. In the context of orbit determination, a variety of approaches exist, including *initial orbit determination* and *admissible regions*. These are addressed in [Section 5.3](#) and [Appendix D](#).

3.3.5 Overview of the LMB parameters

This chapter has introduced a number of parameters which must be adjusted based on the specific tracking scenario under consideration or tuned to alter the behaviour of the LMB filter. A short overview of the most important parameters is provided in [Table 3.1](#) for reference.

Table 3.1: Summary of LMB parameters and their meaning

Category	Name	Symbol	Notes
Single-object	State transition function Eq. (2.22)	$\varphi_{k,k-1}$	Model that maps a target state from time t_{k-1} to t_k . These are the system state dynamics or equations of motion.
	Measurement mapping Eq. (2.23)	h_k	Model to predict the measurements based on the propagated target state at time t_k . This may be dependent on additional sensor information, such as its position or pointing.
	UT scaling parameters Eq. (2.30)	α, β, κ	These parameters influence the assumed shape of the distribution in the unscented transform. Refer to [139] for details.

Category	Name	Symbol	Notes
Adaptive birth Section 3.3.4–G	Max. existence probability	$r_{B,\max}$	This is a cap on the existence probability and can be used to moderate the confidence of new targets. If this is low, targets will require more time to get <i>confirmed</i> , but setting this value too low may result in estimates that should not exist.
	Birth rate Eq. (3.36)	$\lambda_{B,k}$	Expected number of birth components at a particular time t_k .
	Birth-to-clutter Eq. (3.37)	$\frac{\lambda_{B,k}}{\lambda_{c,k-1}}$	Expected ratio between birth and clutter rate. This alternative to λ_B is useful if the number of new objects per step varies a lot over time and there are only few false-positives.
Prediction Section 3.3.1	survival probability Eq. (3.24)	P_S	probability that a target still exists after a single prediction step. In the SSA context, this is usually ~ 1 , since RSOs don't tend to disappear (see Section 5.1.5).
	max. predicted hypotheses	$\vartheta_{C,k k-1}^{\text{GLMB}}$	Maximum number of components/hypotheses to retain after conversion of the predicted LMB to GLMB. Higher is more accurate, but slower.
Clutter Section 3.3.2–B	clutter rate	λ_c	Expected number of false positive measurements in a single observation scan (e.g. image).
	clutter intensity Eq. (3.21)	κ_c	Distribution that integrates to λ_c over the measurement space, describing the relative likelihood of finding a false-positive at any value of \mathbf{z} . Often modelled as uniform over some FOV (see Section 5.2.3).
Update Section 3.3.2 Section 3.3.4–C	detection probability	P_D	Probability that an object is detected, as a function of the predicted target state. For SSA, refer to Section 5.2.2 .
	gating probability Section 3.3.3–B	P_G	The <i>gate</i> is a region $\Delta\mathbf{z} < d_M^G$, where d_M^G is the Mahalanobis distance from the predicted measurement. $1 - P_G$ is the design probability that a measurement is wrongfully rejected. d_M^G and P_G are related by the χ^2 distribution (Appendix E).
	max. corrected hypotheses Section 3.3.4–C	$\vartheta_C^{\text{GLMB}}$	Maximum number of components/hypotheses to retain after the GLMB update. This influences how many measurement associations are considered for each prior hypothesis.
LMB reduction Section 3.3.4–E	pruning threshold	ϑ_P^{LMB}	A higher threshold may reduce runtime at the expense of accuracy (premature target death).
	capping threshold	ϑ_C^{LMB}	Should be significantly larger than the expected number of objects being tracked.

Category	Name	Symbol	Notes
GM reduction Section 2.1.2	pruning threshold Eq. (2.15)	ϑ_P^{GM}	This threshold is a trade-off between runtime and accuracy. A high value results in fewer GM components, which reduces the correctness of the mixture representation.
	merging threshold Algorithm 1	ϑ_M^{GM}	Consider that merging is done using the Mahalanobis distance, and its χ^2 distribution determines the <i>confidence</i> of the merging operation (see Appendix E).
	capping threshold Section 2.1.2–D	ϑ_C^{GM}	Should be significantly larger than the expected number of objects

3.3.6 Multi-object divergence metrics

The *distance* and *divergence* concepts introduced in [Section 2.1.3](#) can be extended to finite sets of vectors. As discussed, one of the applications is to evaluate and compare filter performance. Several metrics are briefly introduced here, but the reader is referred to [Mahler \[84, Sec. 6.2\]](#), who discusses this at length.

A. Optimal sub-pattern assignment (OSPA) metric

In MO tracking scenarios, an important question is how to evaluate the quality of the filter estimate. Since the result is a set of vectors, it is not immediately obvious how it can be compared to a ground truth. If the filter was initialised with prior information, it would be valuable to compare the corresponding states based on the label, but in the general case, it is not clear which target label in the filter belongs to which object in the ground truth.

One solution to this problem is provided by the optimal sub-pattern assignment metric (OSPA), a consistent distance measure that was first proposed by [Schuhmacher and Xia \[113\]\[114\]](#) to compare finite sets. The OSPA metric of order $p \in [1, \infty[$ with cut-off $c > 0$ is here defined between two arbitrary sets $X, Y \subset \mathfrak{Y}$.

Given some single-target distance metric $d(\mathbf{x}, \mathbf{y})$, such as the Euclidean distance, define the associated *cut-off metric* to be $d_c(\mathbf{x}, \mathbf{y}) = \min\{c, d(\mathbf{x}, \mathbf{y})\}$ and let Π_k denote the set of all permutations on $\{1, \dots, k\}$. The OSPA distance is then defined as

$$d_{p,c}^{\text{OSPA}}(X, Y) = \left[\frac{1}{|Y|} \left(\min_{\pi \in \Pi_{|Y|}} \sum_{i=1}^{|X|} d_c(\mathbf{x}_i, \mathbf{y}_{\pi(i)})^p + c^p(|Y| - |X|) \right) \right]^{1/p}, \quad (3.38)$$

provided that $|X| \leq |Y|$. However, if $|X| > |Y|$, the metric can be computed by simply swapping the inputs: $d_{p,c}^{\text{OSPA}}(X, Y) \triangleq d_{p,c}^{\text{OSPA}}(Y, X)$ [\[114\]](#). Finally, $d_{p,c}^{\text{OSPA}}(\emptyset, \emptyset) = 0$ by convention. Notice that c can be seen as the *cost* for cardinality differences and p drives the level of “punishment” for outliers.

In practice, the OSPA is computed in three steps:

- (i) Find the optimal subpattern assignment π , that minimises the summed distance between the elements from X and a selection of $|X|$ vectors from Y (i.e. find the minimum for the first term of [Eq. \(3.38\)](#)).
- (ii) Using the optimal sub-assignment π , set α_i for every \mathbf{y}_i , with the distance d_c if \mathbf{y}_i was assigned and c otherwise:

$$\alpha_i = \begin{cases} d_c(\mathbf{x}_i, \mathbf{y}_{\pi(i)}) & \text{for } 1 \leq i \leq |X| \\ c & \text{for } |X| < i \lesssim |Y| \end{cases}. \quad (3.39)$$

(iii) Compute the p-th order average $(\frac{1}{|Y|} \sum_{i=1}^{|Y|} \alpha_i^p)^{1/p}$.

B. Temporal OSPA (TOSPA) metric

Notice that the OSPA metric can only be used to evaluate the error of the multi-object state at a single epoch. To extend this to the quality of trajectories, [Ristic et al. \[107\]\[108\]](#) introduced a *temporal* version of OSPA (TOSPA or OSPA⁽²⁾), which allows evaluating the performance of the track, including accuracy of localisation, target numbering and target labelling. For more details, the reader is referred to the papers by Ristic et al. and the implementation of this metric is left for future work.

3.4 Implementations in BASIL

All MOT elements in the previous and current chapter are implemented from scratch for this thesis in the Bayesian Association and State Inference Library (BASIL), a new Java library which has Hipparchus⁹ as its only core dependency. The choice to implement this in Java is mainly to ensure compatibility with the ORbits Extrapolation KIT (Orekit), a widely used astrodynamics library. However, the core models in BASIL are entirely independent of Orekit and agnostic to the specific dynamics, state representations, measurement models, sensor assumptions, and so forth. To leverage the MO filters in concrete scenarios, users must implement several models for specific functionality. An overview of the core building blocks for the BASIL MOT toolbox is provided in [Figure 3.14](#).

For the specific case of orbit determination, BASIL also provides a wrapper for Orekit, so that any state representation and measurement type implemented in Orekit can also be used for MOT. One point of caution is that BASIL provides no specific handling for multi-sensor observation sets and one of the core assumptions is that every measurement scan can contain at most one observation for each tracked object. If a user considers fusing the observations from multiple sensors, they must ensure that this assumption is not violated.

BASIL employs a modular approach where each of the elements in the MO filtering recursion (such as the single-object models, MO density, MO estimator, prediction and update methods, birth models and more) can be easily adjusted, extended and replaced. An *observer* pattern is used in some of the core components to attach building blocks that extract information for separate handling. One specific use of the latter is to extract association probabilities (see [Chapter 4](#)) or to feed information back to an adaptive birth model (recall [Section 3.3.4–G](#)).

⁹<https://www.hipparchus.org/>

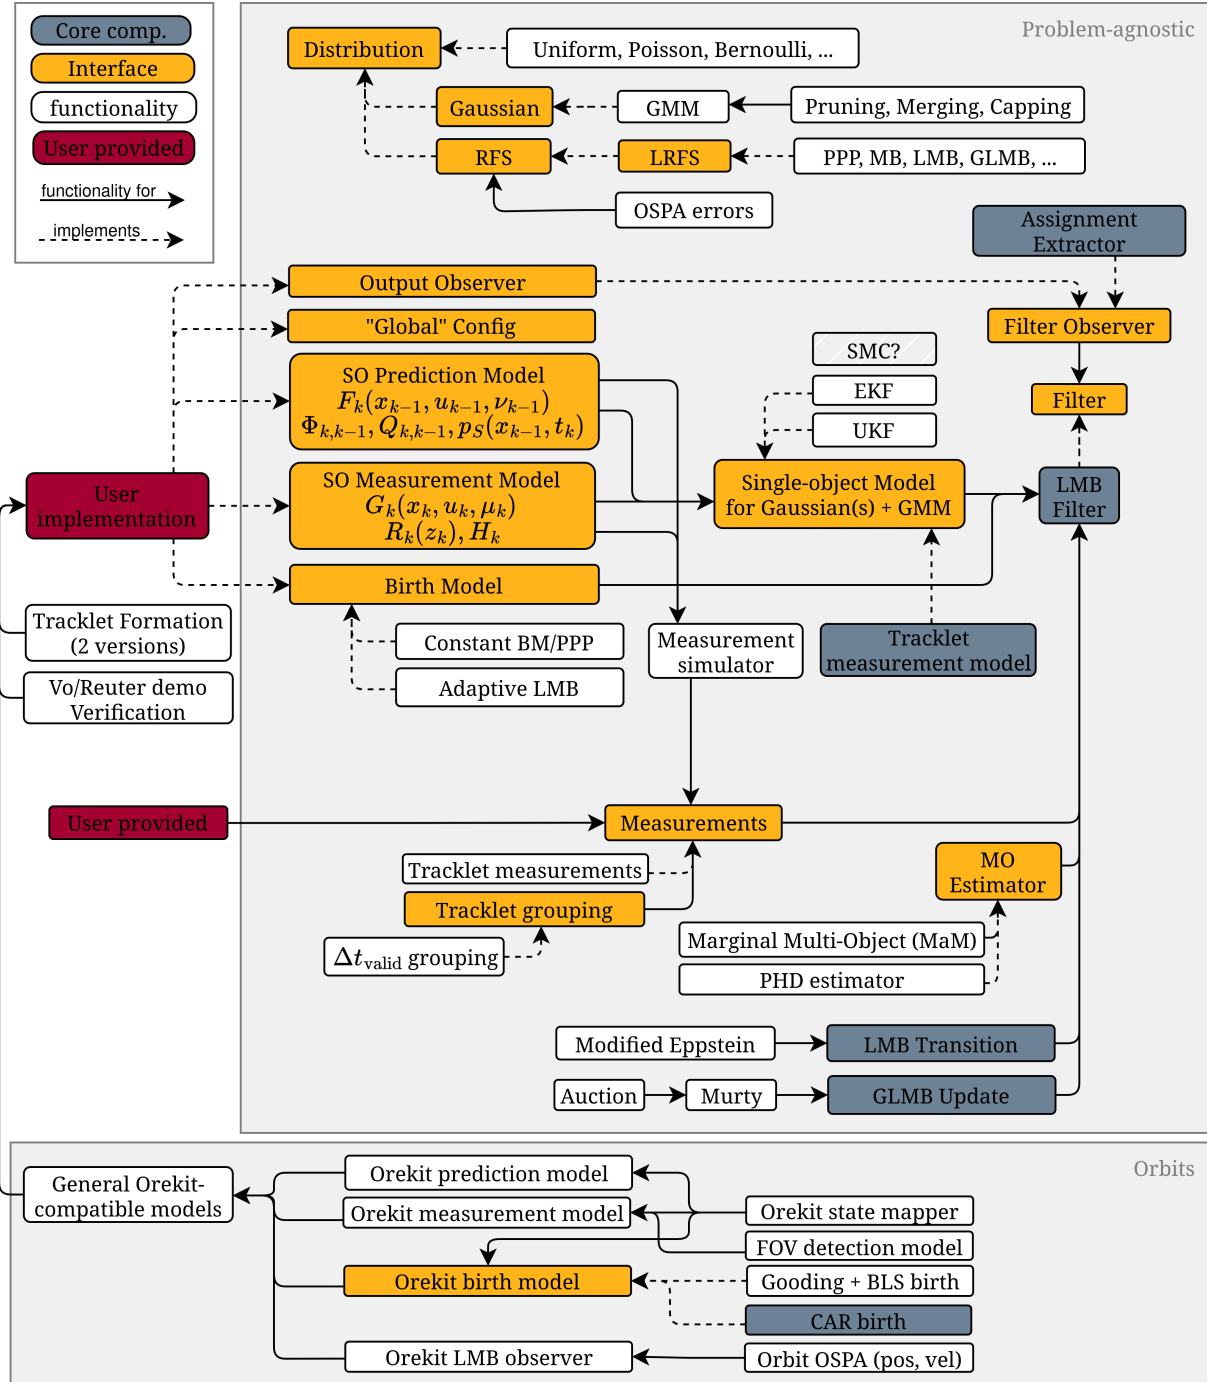


Figure 3.14: Overview of selected MOT components implemented in BASIL. Blue boxes are core blocks with extensive logic that cannot be easily replaced. Notice that the lower box lists orbit-specific functionalities (see [Chapter 5](#)) and the upper box contains some novel features introduced in [Chapter 4](#) and [Section 6.1](#)

EXTRACTING ASSOCIATIONS

As discussed, a core feature of MO trackers is dealing with the measurement association problem in order to obtain an accurate MO trajectory. In that process, however, the actual association results are often disregarded or lost. For many MO density representations, including PHD and LMB, this information is even lost in the process of moment matching, where the contributions of multiple hypotheses are combined to form the target distribution.

However, in the context of SSA and cataloguing it is useful to know which measurements belong to which RSOs, such that they can be grouped together to perform more specific single-object analysis and post-processing.

This section proposes a method to extract the *time-marginal association probabilities* from the GLMB update, i.e. the likelihood that any target was associated to any measurement at a particular time step. In addition, we formulate how one can *accumulate* the probabilities over time if the same measurement label occurs more than once.¹⁰

These two new probabilities can then be used to form a deterministic conclusion on the measurement association and enable storing those associations for future use. Particularly in SSA, where the historical ephemeris of individual objects is often still of interest, extracting that information can be required. As mentioned before, traditional methods obtain this information before even updating the orbits in the catalogue, but the methods of this chapter allow using the full rigour of FISST for the tracklet association purpose.

4.1 Time-marginal association probability

Up to this point, measurement assignments have been done based on their index in the observation scan Z_k (see e.g. Eq. (3.18)), consistent with the existing literature. However, one can also assign an identifier to each measurement to monitor its associations independently of the other measurements. Given an observation vector $\mathbf{z} \in \mathbb{Z}$ and a *measurement label* $\mathbf{m} \in \mathbb{M}$, the labelled single-object measurement is $\hat{\mathbf{z}} \triangleq (\mathbf{z}, \mathbf{m})$. Additionally, define $\mathbb{M}_k \subset \mathbb{M}$ as containing all measurement labels occurring in the measurement set \hat{Z}_k and recall that \mathbb{L}_k represents the set of all existing target labels up to and including time t_k ¹¹:

$$\mathbb{M}_k = \left\{ \mathbf{m} \mid (\mathbf{z}, \mathbf{m}) \in \hat{Z}_k \right\} \subset \mathbb{M}, \quad \text{and} \quad \mathbb{L}_k = \bigcup_{j=1}^k \mathbb{B}_j. \quad (4.1)$$

Then, every posterior GLMB hypothesis $(\xi, L) \in \mathcal{H}_k$, corresponds to an association mapping $\gamma_k^{(\xi, L)} : \mathbb{L}_k \rightarrow \mathbb{M}_k \cup \{-1, 0\}$, where 0 means undetected (*missed*) and -1 means non-existent (*dead*) in the hypothesis ($\ell \notin L$) [134]. Note that this is a slightly extended version of the association hypothesis $\theta_k(\ell)$ from Eq. (3.18), but it accounts for the case that a target label is

¹⁰In other words, if some pre-processing stage already grouped measurements that certainly belong together. This is common in SSA (*tracklets*) and is addressed in Chapter 6.

¹¹Note that the GLMB/LMB recursion never formally removes labels, but approximation by truncation can lead to labels disappearing over time. In practical implementations, \mathbb{L}_k can be interpreted as all target labels that have not been pruned out.

not part of the domain of θ_k . This can happen because different hypotheses could have different label sets or labels could be introduced by the birth process at a later time.

The time-marginal association probability $\alpha_k^{(\ell, m)}$ can then be computed by summing the weights of all GLMB hypotheses at time t_k where the target ℓ was associated to the measurement label m ¹² as

$$\forall \ell \in \mathbb{L}_k, m \in \mathbb{M}_k \cup \{0\} : \quad (4.2a)$$

$$\alpha_k^{(\ell, m)} = \sum_{(\xi, L) \in \mathcal{H}_k} \mathbf{1}_L(\ell) \cdot \delta_{\gamma_k^{(\xi, L)}(\ell)}(m) \cdot w^{(\xi)}(L), \quad (4.2b)$$

where $\mathbf{1}_L$ indicates existing labels (Eq. (2.7)) and the weights w are pre-normalised such that $\sum_{(\xi, L) \in \mathcal{H}_k} w^{(\xi)}(L) = 1$. As a result, the probabilities of non-association are complementary and determined by

$$\forall \ell \in \mathbb{L}_k : \quad \alpha_k^{(\ell, -1)} \triangleq 1 - \sum_{m \in \mathbb{M}_k} \alpha_k^{(\ell, m)} - \alpha_k^{(\ell, 0)} \quad (4.3a)$$

$$= \Pr(\ell \text{ non-existent}), \quad (4.3b)$$

$$\forall m \in \mathbb{M}_k : \quad \bar{\alpha}_k^{(m)} \triangleq 1 - \sum_{\ell \in \mathbb{L}_k} \alpha_k^{(\ell, m)} \quad (4.3c)$$

$$= \Pr(z_k^{(m)} \text{ clutter} \vee m \text{ target unknown}). \quad (4.3d)$$

Notice that an unassociated m still has an ambiguous origin, as it might mean that the corresponding measurement at time t_k is clutter *or* that the target to which m belongs does not (yet) exist in the MO filter. Figure 4.1 illustrates this by a sample association result and Algorithm 4.1 details how the values for Eq. (4.2) and Eq. (4.3) can be obtained.




Figure 4.1: Example label association result. The widths of the connection represent the association probability. Note that these sum to 1 for each $m \in \mathbb{M}_k$ and each $\ell \in \mathbb{L}_k$, but not for the non-association nodes ($m \in \{0, -1\}$ and $\ell \notin \mathbb{L}_k$).

In case the observations used for the filter are *unlabelled*, i.e. there is no prior knowledge on any correlation between the measurements, then these *time-marginal* probabilities constitute all the information the GLMB filter has on which measurement-to-object associations are likely and which are not. They can then be used to draw a definitive conclusion on the measurement assignment, for example using a MAP estimator like in Eq. (4.8). Knowing which measurements belong to each other is often useful in the space cataloguing context, as it allows for single-object processing such as precise orbit determination and conjunction assessment.

¹²Notice that $\alpha_k^{(\ell, 0)}$ is the probability that target ℓ was not detected at time t_k .

It should be noted that the association probabilities extracted from the GLMB recursion are based on the joint likelihood of consistent MO measurement assignments, and not just on the likelihood that each single target produced each single measurement (as traditional approaches do). The probabilities depend on all observations in the scan, together with the models for false positives and missed detections.

Algorithm 4.1: Marginal association probability mapping at time t_k

```

ALPHA( $\mathcal{H}_k, \gamma_k$ ):
1   $\alpha_k^{(\ell, m)} = 0, \quad \alpha_k^{(\ell, -1)} = 1, \quad \bar{\alpha}_k^{(m)} = 1 \quad \triangleright \forall \ell \in \mathbb{L}_k, m \in \mathbb{M}_k$ 
2   $\alpha_k^{(\ell, 0)} = 0 \quad \triangleright \forall \ell \in \mathbb{L}_k$ 
3  for  $(\xi, L) \in \mathcal{H}_k$ 
4       $w = w^{(\xi)}(L)$ 
5      for  $\ell \in L$ 
6           $m = \gamma_k^{(\xi, L)}(\ell) \quad \triangleright \text{Note } m \neq -1$ 
7           $\alpha_k^{(\ell, m)} \leftarrow \alpha_k^{(\ell, m)} + w \quad \triangleright \ell \text{ and } m \text{ associated}$ 
8           $\alpha_k^{(\ell, -1)} \leftarrow \alpha_k^{(\ell, -1)} - w \quad \triangleright \ell \text{ exists}$ 
9          if  $m \in \mathbb{M}_k$ 
10              $\bar{\alpha}_k^{(m)} \leftarrow \bar{\alpha}_k^{(m)} - w \quad \triangleright m \text{ not clutter/unknown}$ 
11          end if
12      end for
13 end for
14 return  $\alpha_k^{(\cdot, \cdot)}, \bar{\alpha}_k^{(\cdot)}$ 

```

4.2 Accumulated (tracklet) association probability

If, instead, some prior knowledge indicates that measurements from different scans belong to the same object, then this information can be used to refine the association probabilities over multiple steps in the filter. Such prior knowledge can arise, for example, if close-spaced observations are already collected into *tracklets* by some external method (as is often the case in SSA). It is important to stress that nothing here has an influence on the LMB/GLMB recursion itself. The state estimates are still updated based on single observations, but the accumulated probabilities allow cataloguing *tracklet associations* instead of single measurement associations, reducing the influence of outliers and of edge-effects introduced by the models for detection probability and clutter¹³.

It is here assumed that a process like *tracklet formation* has assigned the same label m to all measurements (over multiple consecutive observation scans) that certainly belong to the same object. This section details a proposed approach to *accumulate* the (time-marginal) association probabilities for all measurements with the same label and effectively obtain the *tracklet association probability* $\alpha_{1:k}^{(m)}(\cdot)$. Several considerations complicate this process.

- (i) Some measurement labels might occur only once, and others many times.
- (ii) The time-marginal associations probability is conditioned on the existing label set \mathbb{L}_k , so an unassociated measurement label m could correspond to a target ℓ that has yet to be included in the filter birth process.

¹³This is later demonstrated by a numerical simulation in [Section 7.2](#)

- (iii) The accumulated quantity must define a PMF for every tracklet label \mathbf{m} , meaning that the sum of probabilities for all possible assignments of that label is

$$\sum_{\ell \in \mathbb{L}_k} \alpha_{1:k}^{(\mathbf{m})}(\ell) + \bar{\alpha}_{1:k}^{(\mathbf{m})} = 1, \quad (4.4)$$

where $\alpha_{1:k}^{(\mathbf{m})}(\ell)$ is the probability that \mathbf{m} is assigned to target label ℓ and $\bar{\alpha}_{1:k}^{(\mathbf{m})}$ represents the false-positive probability (similar to Eq. (4.3)), all based on information up to and including time t_k .

To approach this, it is useful to reiterate the assumption that every measurement label \mathbf{m} can originate from at most one target ℓ , but a single ℓ can give rise to any number of measurement labels as long as they are in different time steps of the filter. For this reason, the accumulated probability $\alpha_{1:k}^{(\mathbf{m})}(\ell)$ is a separate mapping $\mathbb{L}_k \rightarrow [0, 1)$ for every measurement \mathbf{m} .

For convenience, define the *extended weight* $\beta_k^{(\mathbf{m})}(\ell)$ related to the marginal association probability at time t_k but including all possible target labels in its domain as

$$\beta_k^{(\mathbf{m})}(\ell) = \begin{cases} \alpha_j^{(\ell, \mathbf{m})} & \text{if } \ell \in \mathbb{L}_k \\ \bar{\alpha}_j^{(\mathbf{m})} & \text{otherwise} \end{cases}, \quad (4.5)$$

Note that these are *not* probabilities, as there could be an infinite number of future labels that are not yet included in \mathbb{L}_k . These weights can be accumulated for every \mathbf{m} over all the times where \mathbf{m} was present by

$$\beta_{1:k}^{(\mathbf{m})}(\ell) = \prod_{t_j \in T_{1:k}^{(\mathbf{m})}} \beta_j^{(\mathbf{m})}(\ell), \quad \text{and} \quad \bar{\beta}_{1:k}^{(\mathbf{m})} = \prod_{t_j \in T_{1:k}^{(\mathbf{m})}} \bar{\alpha}_j^{(\mathbf{m})}, \quad (4.6a)$$

$$\text{with} \quad T_{1:k}^{(\mathbf{m})} = \{t_j \mid \mathbf{m} \in \mathbb{M}_j\}_{j=1}^k. \quad (4.6b)$$

Using these accumulated weights, the accumulated association probabilities per measurement label are computed as

$$\forall \mathbf{m} \in \mathbb{M}_k : \quad \alpha_{1:k}^{(\mathbf{m})}(\ell) = \frac{\beta_{1:k}^{(\mathbf{m})}(\ell)}{\Sigma \beta_{1:k}^{(\mathbf{m})}}, \quad \text{and} \quad \bar{\alpha}_{1:k}^{(\mathbf{m})} = \frac{\bar{\beta}_{1:k}^{(\mathbf{m})}}{\Sigma \beta_{1:k}^{(\mathbf{m})}}, \quad (4.7a)$$

$$\text{where} \quad \Sigma \beta_{1:k}^{(\mathbf{m})} = \sum_{\ell' \in \mathbb{L}_k} \left[\beta_{1:k}^{(\mathbf{m})}(\ell') \right] + \bar{\beta}_{1:k}^{(\mathbf{m})}. \quad (4.7b)$$

This ensures that the accumulated probabilities sum to 1 for each measurement label¹⁴. To avoid overflow/underflow complications in this normalisation step, one can use the *log-sum-exp trick* described in [Appendix E.2](#).

[Algorithm 4.2](#) shows a recursive implementation for the accumulation of probabilities in [Eq. \(4.7\)](#), as used inside an MO filter of this thesis.

¹⁴Contrary to the time-marginal association probabilities, these accumulated values are *not* probabilities on \mathbb{L}_k , since it is possible for multiple measurement labels \mathbf{m} to be associated to a single target ℓ , as long as they exist at different times.

To illustrate the validity of this approach intuitively, consider these simplifying edge cases as a sanity check:

- (i) If no targets are ever created ($\mathbb{L}_k = \emptyset$), then $\forall m \in \mathbb{M}_k : \bar{\alpha}_{1:k}^{(m)} = 1$.
- (ii) If a measurement label m exists at exactly one time step t_j , then $\forall \ell \in \mathbb{L}_j : \alpha_{1:k}^{(m)}(\ell) = \alpha_j^{(\ell, m)}$ and $\forall \ell \notin \mathbb{L}_j : \alpha_{1:k}^{(m)}(\ell) = \bar{\alpha}_j^{(m)}$.

Algorithm 4.2: Accumulated association probability mapping up to time t_k . This implementation assumes that pruned-out labels can never return, which is equivalent to assuming the birth process generates unique labels.

```

ACCUMULATED-ALPHA( $\alpha_k, \beta_{1:k-1}, \mathbb{L}_k, \mathbb{M}_k$ ):
1   $\mathbb{M}_{1:k} = \mathbb{M}_{1:k-1} \cup \mathbb{M}_k$ 
2  for  $m \in \mathbb{M}_k$ 
3      if  $m \notin \mathbb{M}_{1:k-1}$ 
4           $\log[\bar{\beta}_{1:k-1}^{(m)}] = 0$ 
5      end if
6      for  $\ell \in \mathbb{L}_k$ 
7          if  $\ell \in \mathbb{B}_k$  ▷ or  $\ell \notin \mathbb{L}_{k-1}$ 
8               $\log[\beta_{1:k-1}^{(m)}(\ell)] = \log[\bar{\beta}_{1:k-1}^{(m)}]$ 
9          end if
10          $\log[\beta_{1:k}^{(m)}(\ell)] = \log[\beta_{1:k-1}^{(m)}(\ell)] + \log[\alpha_k^{(\ell, m)}]$ 
11     end for
12      $\log[\bar{\beta}_{1:k}^{(m)}] = \log[\bar{\beta}_{1:k-1}^{(m)}] + \log[\bar{\alpha}_k^{(m)}]$ 
13      $\alpha_{1:k}^{(m)}(\cdot), \bar{\alpha}_{1:k}^{(m)} \leftarrow \text{normalise}(\beta_{1:k}^{(m)}(\cdot); \bar{\beta}_{1:k}^{(m)})$  ▷ Eq. (4.7)
14 end for
15 return  $\alpha_{1:k}, \bar{\alpha}_{1:k}$ 

```

The *association conclusion* then consists of the most likely originator for each measurement label and can be obtained by a MAP estimator:

$$\forall m \in \mathbb{M}_{1:k} : \hat{\ell}^{(m)} = \begin{cases} \arg \sup_{\ell \in \mathbb{L}_k} \alpha_{1:k}^{(m)}(\ell) & \text{if } \sup_{\ell \in \mathbb{L}_k} \alpha_{1:k}^{(m)}(\ell) > \bar{\alpha}_{1:k}^{(m)} \\ \text{clutter/unknown} & \text{otherwise} \end{cases} \quad (4.8)$$

Finally, measurements associated to the same target could be collected and stored for future single-object estimation processes:

$$\mathcal{M}(\ell) = \{m \in \mathbb{M}_{1:k} \mid \hat{\ell}^{(m)} = \ell\}. \quad (4.9)$$

In conclusion, the quantities proposed in this chapter allow for the principled extraction of complete measurement histories for each of the targets in a GLMB/LMB filter, along with the confidence of the assignments. The time-marginal association probabilities from Eq. (4.2) do this for single unlabelled measurements, whereas the accumulated probabilities from Eq. (4.7) can be used to obtain a single assignment for complete pre-formed tracklets. These quantities can either be used for further SO processing or to determine the confidence of new components in a tracklet-based adaptive birth model (see Section 5.3.3).

MULTI-OBJECT TRACKING IN SPACE

The previous sections discussed the relevant background and formulation of RFS-based multi-object tracking. When using these to solve orbital tracking problems in SSA, several domain-specific challenges arise. This section introduces the main building blocks to use the LMB filter for orbital tracking of RSOs. The present implementation of everything related to spacecraft dynamics relies on the open-source software package Orekit [85], but the MOT logic is completely separate and could be used for other representations of the dynamical and observation models.

5.1 Orbital motion model

The main part of the RSO target state consists of six orbital components. The most straightforward representation of that state combines simply its position vector \mathbf{r} and velocity vector $\mathbf{v} = \dot{\mathbf{r}}$, expressed in some Cartesian inertial reference frame:

$$\mathbf{x} = \begin{pmatrix} \mathbf{r} \\ \dot{\mathbf{r}} \end{pmatrix} = (x \ y \ z \ \dot{x} \ \dot{y} \ \dot{z}). \quad (5.1)$$

Ignoring process noise for now, the motion model (recall Eq. (2.21)) for this Cartesian, orbital state is

$$\dot{\mathbf{x}} = \begin{pmatrix} \dot{\mathbf{r}} \\ \ddot{\mathbf{r}} \end{pmatrix}, \quad (5.2)$$

where the acceleration vector over time $\ddot{\mathbf{r}}$ is determined by the equations of motion and applied force models (see Section 5.1.2).

Note that it is very common in orbit determination to extend the spacecraft state by parameters of the relevant force models (e.g. drag and radiation pressure coefficients), so that these can be estimated as well, instead of relying on prior knowledge only. The orbital model in BASIL does not presently support adding estimated dynamical parameters and limits itself to 6D orbital state estimation. However, this would require only minor modification, as Orekit has extensive support for parameter estimation and only a robust mapping between the state representations of Orekit and BASIL is lacking.

5.1.1 Reference frames

Many reference frames exist in which to express the state and associated equations of motion. Here, only the relevant reference frames are addressed as used in the thesis.

A. Earth-centred frames

The state of geocentric orbits is most easily expressed in a coordinate system that also has the Earth as reference. By convention, these frames always have their \mathbf{u}_z axis aligned with the Earth's axis of rotation, such that \mathbf{u}_x and \mathbf{u}_y are in the equatorial plane. However, the orientation of the latter two depends on the specific coordinate system. In Earth-centred inertial (ECI) frames, \mathbf{u}_x usually points to the *vernal equinox* \mathcal{V} and \mathbf{u}_y completes the right-handed system, resulting in a *quasi-inertial* reference frame. Figure 5.1 shows this family of frames as the (X, Y, Z) axes. This figure schematically represents the Earth-Sun two-body motion from a geocentric perspective, so that the vernal equinox can be seen as the ascending node of the ecliptic orbit ($\mathbf{u}_x = \mathbf{u}_z \times$

\mathbf{h}). The difference between various ECI frames is usually how the equatorial and ecliptic planes are defined, since these vary over time due to perturbations. This thesis uses the Earth mean equinox reference frame at Julian epoch J2000 (EME2000) [92].




Figure 5.1: The Sun-Earth two-body motion is represented in geocentric perspective to illustrate the family of Earth-centred reference frames (X', Y', Z') in relation to the celestial parameters. The orbital plane defined by the Sun-Earth angular momentum \mathbf{h} is the *ecliptic* the intersections with the *equatorial* plane are the *vernal* and *autumnal* equinox. (X, Y, Z) represents the ECI frame and ECEF is shown as (x, y, z) [124, p. 30].

When expressing positions relative to objects of interest on Earth (such as ground stations and observers), it is often more straightforward to use a rotating frame with the same angular velocity as the Earth. In such Earth-centred Earth-fixed (ECEF) frames, the \mathbf{u}_x basis vector is usually directed towards the Greenwich meridian and \mathbf{u}_y again completes the system – the frame is represented as (x, y, z) in Figure 5.1. Once more, slight variations exist between the references of various frames. ECEF positions in this work are expressed in the International Terrestrial Reference Frame (ITRF) with International Earth Rotation Service 2010 (IERS2010) conventions [12][60].

B. Object-centred frame

Analysing estimation errors is usually inconvenient and not very insightful in Earth-centred frames. It is often more useful to formulate the error with respect to a frame that follows the truth or *reference orbit*. One common choice is referred to as the radial, transverse, normal (RTN) frame and is visualised in Figure 5.2. The basis vector \mathbf{u}_R is aligned with the reference position vector, \mathbf{u}_N follows the angular momentum \mathbf{h} of the reference orbit and \mathbf{u}_T completes the right-handed, orthogonal system.




Figure 5.2: RTN reference frame [124, p. 42]

5.1.2 Equations of motion

The equations of motion for orbital states describe how the state evolves over time. In the case of the Cartesian state representation (Eq. (5.1)), this means defining an acceleration model depending on the position and velocity of the object and as a function of time.

The main contribution for an object in Earth's orbit is the point mass gravity contribution, modelled as

$$\ddot{\mathbf{r}} = -\frac{\mu_E}{\|\mathbf{r}\|^3}\mathbf{r}, \quad (5.3)$$

where $\mu_E = Gm_E \approx 3.986 \times 10^{14} \text{ m}^3 \text{ s}^{-2}$ is the standard gravitational parameter.

In reality, various perturbations complicate the orbital motion. Since BASIL relies on Orekit [85] to perform all orbital propagation, the models are not developed in detail here, but the main perturbations are listed for reference.

A. Spherical harmonic gravity

Since the Earth is in reality not a perfect sphere with uniform density but rather an oblate ellipsoid commonly referred to as the *geoid*, the gravity is commonly described in terms of its potential

$$U_E = \frac{\mu_E}{\|\mathbf{r}\|} \sum_{\ell=0}^{\infty} \left(\frac{a_E}{\|\mathbf{r}\|} \right)^{\ell} \sum_{m=0}^{\ell} P_{\ell,m}(\sin(\phi)) [C_{\ell,m} \cos(m\lambda) + S_{\ell,m} \sin(m\lambda)], \quad (5.4)$$

where a_E is the semi-major axis of the Earth's reference ellipsoid [93], ϕ and λ are the latitude and longitude in an ECEF reference frame, $C_{\ell,m}$ and $S_{\ell,m}$ are the *spherical harmonic coefficients* and $P_{\ell,m}$ is the *associated Legendre polynomial* of degree ℓ and order m .

The potential can be approximating by truncating the double series in Eq. (5.4) to the desired degree and order and the resulting acceleration caused by the Earth's gravity is then the gradient of that potential field

$$\ddot{\mathbf{r}}_{G,E} = \nabla U_E. \quad (5.5)$$

Note that this simplifies to Eq. (5.3) for $\ell_{\max} = m_{\max} = 0$.

B. Third-body gravity perturbations

Apart from the Earth, other celestial bodies such as the Sun and the Moon cause additional gravitational accelerations. These are usually modelled as point-mass gravity effects, taking into account that Earth also experiences an influence from the same perturbing bodies.

C. Drag

Even though the atmospheric density reduces exponentially with altitude, there is still a non-negligible drag force causing an acceleration opposing the spacecraft velocity. Using some underlying density model as a function of position and time, the acceleration can be written as

$$\ddot{\mathbf{r}}_D = -\frac{1}{2}C_b\rho(\mathbf{r},t)\|\mathbf{v}_{\text{rel}}\|\mathbf{v}_{\text{rel}}, \quad (5.6)$$

where $C_b = \frac{C_d S_D}{m}$ is the ballistic coefficient, composed of the dimensionless drag coefficient C_d , the effective drag area S_D and the object mass m . Note that the velocity $\mathbf{v}_{\text{rel}} = \dot{\mathbf{r}} - \mathbf{v}_{\text{atm}}$ is the velocity relative to the atmosphere, which has its own motion due to the rotation of the Earth and wind effects [125]. In LEO orbits, this is usually the main source of propagation errors over time, due to the high complexity of atmospheric density modelling.

D. Solar radiation pressure

Photons emitted by the Sun constantly interact with objects in space and thereby transfer some of their momentum. Since the photons all come from the same direction, the combined effect is a resultant force called solar radiation pressure (SRP). The *Cannonball model* approximates the space object as a sphere and formulates the corresponding acceleration as [146]

$$\ddot{\mathbf{r}}_{\text{SRP}} = \frac{P_0 R_0^2}{c} C_{\text{SRP}} \frac{\mathbf{r}_s}{\|\mathbf{r}_s\|^3}, \quad (5.7)$$

where $P_0 \approx 1367 \text{ W/m}^2$ is the Solar flux at a distance of $R_0 = 1 \text{ AU}$, c is the speed of light and the radiation pressure coefficient $C_{\text{SRP}} = \frac{C_r S_{\text{SRP}}}{m}$ consists of the reflection coefficient C_r , the SRP effective area S_{SRP} and the mass m . Finally, $\mathbf{r}_s = \mathbf{r} - \mathbf{r}_{\text{Sun}}$ is the vector from the Sun to the object.

5.1.3 Evolution of uncertainty

It has already been discussed that the Gaussian approximations are not always valid in non-linear scenarios. To illustrate that the propagation of Cartesian coordinates is non-Gaussian, even with the most simple dynamics, Figure 5.3 visualises the uncertainty evolution under Keplerian motion in LEO. The blue dot represents the initial state for an orbit at 700 km perigee altitude and eccentricity $e = 0.1$ with a Gaussian uncertainty where $\sigma_r = 10 \text{ km}$ and $\sigma_v = 10 \text{ m/s}$ in all directions. Propagating 1000 samples for 1 day results in the distribution of red dots, which is clearly non-Gaussian. A linearised prediction using the EKF equations would result in the yellow covariance ellipsoid, which is not at all representative for the true distribution.

To resolve this problem, one might choose to use a GM representation of the initial distribution as discussed in Section 2.1.2, which definitely improves the fidelity of the predicted uncertainty. However, it can be observed that the samples mainly spread along the orbit while the shape of that orbit is not heavily affected. If the state is instead represented in orbital elements, it becomes clear that the two-body dynamics are almost linear and only affect the uncertainty in true anomaly.




Figure 5.3: Uncertainty evolution in LEO under Keplerian dynamics. Initial uncertainty is Gaussian (blue dot) and 1000 samples are propagated analytically for 1 day (14 revolutions). The Gaussian properties are lost in Cartesian coordinates (left) but remain in the orbital element representation (right).

5.1.4 State noise compensation

Choosing suitable models to represent process noise is a complicated task and often requires extensive tuning with application-specific test cases and background knowledge. In this work, the process noise assumes that an object in orbit is a constant velocity particle in the RTN frame with a white noise perturbation on the acceleration $\nu_a \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_a)$; the fact that this is not an inertial reference frame is ignored. Eq. 43 from Reid and Term [103] is then used to obtain the continuous-time process noise matrix for position and velocity in the RTN frame as

$$\mathbf{Q}_{k,k-1}^{\text{RTN}} = \begin{pmatrix} \frac{\Delta t^3}{4} \mathbf{Q}_a & \frac{\Delta t^2}{2} \mathbf{Q}_a \\ \frac{\Delta t^2}{2} \mathbf{Q}_a & \Delta t \mathbf{Q}_a \end{pmatrix}, \quad \text{where} \quad \mathbf{Q}_a = \begin{pmatrix} q_R & 0 & 0 \\ 0 & q_T & 0 \\ 0 & 0 & q_N \end{pmatrix}. \quad (5.8)$$

Here Δt is the propagation time between two measurement scans and q_R, q_T, q_N are tuning parameters. To avoid overly large noise after long propagations, the time gap is artificially capped at Δt_{\max} , so that $\Delta t = \min(t_k - t_{k-1}, \Delta t_{\max})$. This RTN process noise matrix can then be transformed to the appropriate state space and used in the SO prediction step.

5.1.5 Survival model

The final component to make this a complete MO prediction model is the probability of survival (recall Section 3.3.1). In principle, the only way RSOs could disappear is if they reach end of life and burn up in the atmosphere or depart Earth orbit. For this reason, a common choice is to use a constant $P_S \lesssim 1$. However, it is also possible to exploit this mechanism to enforce constraints on the orbits of interest. Sometimes, incorrect measurement associations can give rise to unrealistic orbit predictions, so the survival model used here is

$$p_S(\mathbf{x}, \ell) = \begin{cases} P_S & \text{if } \mathbf{x} \in \mathcal{C} \\ 0 & \text{if } \mathbf{x} \notin \mathcal{C} \end{cases} \quad (5.9)$$

Where $\mathcal{C} \subseteq \mathbb{X}$ is the constrained region defined by any number of orbital constraints. To avoid hyperbolic orbits and sub-orbital trajectories, define

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{X} : a(1 - e) > R_E \wedge e < 1\}, \quad (5.10)$$

where a and e are the semi-major axis and eccentricity of the orbit associated to \mathbf{x} and R_E is the equatorial radius of the Earth. If a constrained admissible region approach is used for the birth

process (see Section 5.3), then it could also make sense to apply the same constrained admissible region (CAR) constraints here.

Implementation note

If \mathbf{x} is distributed as a GM $p(\mathbf{x}, \ell) = \mathcal{G}$, then the survival probability can be computed as

$$p_S(\mathbf{x}, \ell) = \sum_{(w, \boldsymbol{\mu}, \mathbf{P}) \in \mathcal{G}} w p_S(\boldsymbol{\mu}, \ell). \quad (5.11)$$

5.2 Angles-only optical observation model

Within the scope of this thesis, observations for SSA are optical measurements $\mathbf{z} = (\alpha, \delta)$, representing the topocentric *right ascension* and *declination* in some ECI coordinate system (here: EME2000), as determined from an observer with state $\mathbf{x}_o = (\mathbf{r}_o^T, \mathbf{v}_o^T)$. For a target object with known state $\mathbf{x} = (\mathbf{r}^T, \mathbf{v}^T)^T$, the *slant vector*¹⁵ $\boldsymbol{\rho} = \mathbf{r} - \mathbf{r}_o$ fully determines the simplified measurement model:

$$\alpha(\mathbf{x}) = \text{atan2}(\langle \boldsymbol{\rho}, \mathbf{u}_j \rangle, \langle \boldsymbol{\rho}, \mathbf{u}_i \rangle), \quad (5.12a)$$

$$\text{and } \delta(\mathbf{x}) = \arcsin(\langle \boldsymbol{\rho}, \mathbf{u}_k \rangle), \quad (5.12b)$$

where $(\mathbf{u}_i \ \mathbf{u}_j \ \mathbf{u}_k) = \mathbf{I}_{3 \times 3}$ represent the unit vectors of the ECI coordinate system and $\langle \cdot, \cdot \rangle$ represents the dot product. Note that the slant range $\|\boldsymbol{\rho}\|$ is unknown, so only part of the position vector is observable with these measurements. Other measurement types like RADAR can also incorporate that extra information.

In any case, the model in Eq. (5.12) does not take into account that the target and observer move during the time the light travels from one to the other. This is often referred to as the light time correction. In other words, one must compute the angles such that they correspond to where the satellite was at the time the observed light was reflected on it. This effect is visualised in Figure 5.4.

For a satellite in GEO orbit, which is at least 35 790 km away from the ground observer, and moves at about 3.1 km/s, the light travels for about $\Delta t \approx \frac{\|\boldsymbol{\rho}\|}{c} = 0.12$ s. This corresponds to a displacement of ~ 370 m or about 2.1". Since measurement noise is often in the order of several arcseconds (1σ), this effect is not negligible.

For measurements where the observer is ground-based, the implementation of this model is provided by Orekit¹⁶. Using a simple wrapper, the Orekit implementation can be used as *measurement model* for the filters in BASIL (see Section 3.4), without any adaptations.

5.2.1 Space-based observers

Apart from ground-based telescopes, Vyoma will rely on its constellation of space-based telescopes to collect observations. This offers a multitude of advantages, as discussed in Appendix A, and since the measurements are still (α, δ) , a very similar measurement model can be used. The main difference is that the observer position is no longer fixed in an ECEF reference frame, and must be obtained by dedicated orbit determination (OD) routines. In the present work, it is

¹⁵This is also often referred to as the line of sight (LOS) vector, with respect to the observer

¹⁶The AngularRaDec measurement type is documented at <https://www.orekit.org/static/apidocs/org/orekit/estimation/measurements/AngularRaDec.html>.

assumed that this is done ahead of time and accurate observer ephemeris is available¹⁷.

In addition to the light-time effect (depicted in Figure 5.4), optical observations are also distorted by *stellar aberration*. This is a relativistic effect caused by the observer's velocity with respect to the inertial reference frame. Due to the movement of the observer relative to the incoming light, the inertial star background is slightly distorted, and the angle between the observer velocity and line of sight to the star is observed (θ_{obs}) smaller than in reality (θ_{true}). Since the right ascension and declination of the tracking observations are obtained by comparison with inertial star background, this stellar aberration must be incorporated in the measurement model.

$$\tan\left(\frac{\theta_{\text{obs}}}{2}\right) = \sqrt{\frac{1 - v_o/c}{1 + v_o/c}} \tan\left(\frac{\theta_{\text{true}}}{2}\right), \quad (5.13)$$

where $v_o = \|\mathbf{v}_o\|$ is the observer's velocity in the ECI frame and c is the speed of light. The influence of aberration is visualised in Figure 5.5. Notice that the effective observer velocity in the celestial frame consists of (1) its quasi-inertial velocity w.r.t. the Earth and (2) the movement of the ECI frame relative to the Sun. Since the latter has a period of 1 year, it is commonly referred to as the *annual aberration*. The former, short-period component has a *daily* variation for ground-based observers, whereas space-based sensors introduce a once-per-orbit fluctuation.

- To compute the maximum annual aberration, consider that the Earth moves at about 29.8 km/s w.r.t. the Sun. For an inertial angle of 90° , where the effect is largest, the offset becomes $\Delta\theta_{\text{annual}}^{\text{max}} \approx 20.5''$.
- For a ground-based observer, the maximum ECI velocity occurs at the equator with about 465 m/s. Again evaluating the offset for $\theta_{\text{true}} = 90^\circ$, the daily aberration is limited by $\Delta\theta_{\text{daily}}^{\text{max}} \approx 0.32''$.
- Finally, an observer in circular LEO orbit, e.g. at 500 km altitude, has a velocity of 7.6 km/s. This causes an aberration up to $\Delta\theta_{\text{LEO}}^{\text{max}} \approx 5.2''$.

Depending on the measurement noise level of the sensor, it might be acceptable to ignore the daily aberration, but the annual and once-per-orbit effects must certainly be accounted for to avoid that model limitations dominate the measurement errors.




Figure 5.4: Exaggerated visualisation of light-time effect for space-based observer




Figure 5.5: Exaggerated visualisation of aberration effect for space-based observer

¹⁷It is assumed that the influence of observer position errors on the observations is at least an order of magnitude below the measurement noise.

5.2.2 Probability of detection

Another essential part of the MO filtering model is quantifying the detection probability. Recall from Eq. (2.23) that the detection model is part of the MO measurement model and has an implicit influence on the existence probability of the targets in the filter. For example, if a target has a very high detection probability but none of the likely association solutions assign a measurement to this target, then its probability of existence will reduce.

Since the sensors in the SSA case have a very limited FOV compared to the complete state space, most targets will generally be unobservable most of the time. For this reason, it is common to use a constant detection probability P_D if the target is inside the field of view, and 0 otherwise.

Define the local reference frame of a sensor s using the non-inertial basis vectors $(\mathbf{b}_s, \mathbf{j}_s, \mathbf{k}_s)$, where \mathbf{b}_s is the *boresight* vector, \mathbf{j}_s is the image horizontal and \mathbf{k}_s completes the system. Then define the field of view of sensor s as $\text{FOV}_s = \{\psi_s\} \subseteq [0, \pi] \times [0, 2\pi)$ where

- $\psi_s \triangleq (\theta_s, \phi_s)^T$ is the angular position of a line of sight in the sensor frame.
- θ_s is the offset angle between the line of sight vector $\boldsymbol{\rho}_s = \mathbf{R}^{s \leftarrow \text{ECI}}(\mathbf{r} - \mathbf{r}_s)$ and the sensor's boresight vector \mathbf{b}_s .
- ϕ_s is the angle between the projected line of sight and the image horizontal \mathbf{j}_s .

Given a state \mathbf{x} , this angular position can be computed for the sensor as

$$\psi_s(\mathbf{x}, t) = \begin{pmatrix} \theta_s(\mathbf{x}, t) \\ \phi_s(\mathbf{x}, t) \end{pmatrix} = \begin{pmatrix} \arccos(\langle \hat{\boldsymbol{\rho}}_s, \mathbf{b}_s \rangle) \\ \text{atan2}(\langle \hat{\boldsymbol{\rho}}_s, \mathbf{k}_s \rangle, \langle \hat{\boldsymbol{\rho}}_s, \mathbf{j}_s \rangle) \end{pmatrix}. \quad (5.14)$$

The detection probability is then

$$P_D(\mathbf{x}, \ell, t_k | s) = \begin{cases} P_D & \text{if } \psi_s(\mathbb{E}[\mathbf{x}], t_k) \in \text{FOV}_s \\ 0 & \text{otherwise} \end{cases}. \quad (5.15)$$

However, if the previous state was very uncertain (for example because it was the result of a birth process and has not yet been *confirmed* by more measurements), significant portions of the PDF may lie both inside and outside the FOV (illustrated in Figure 5.6). If a high P_D is assigned and the target is not observed, its existence probability will drop and it might never get confirmed. Conversely, if it is assigned $P_D = 0$ while it was in fact observed, this measurement will be effectively ignored and cannot help to confirm the target. For this reason, a more rigorous procedure is applied for the detection probability.

Given that $\hat{\mathbf{x}}$ has a predicted distribution $p_{k|k-1}(\cdot, \ell)$, one can project this distribution onto the 2D ψ -space using an unscented transform

$$P_{k|k-1}(\psi_s, \ell) \stackrel{\text{UT}}{\leftarrow} p_{k|k-1}(\mathbf{x}, \ell), \quad (5.16)$$

after which this density can be integrated over the entire FOV – this results in the probability that the target is inside the field of view at time t_k . Finally, multiply by the constant P_D to obtain the realistic detection probability

$$P_D(\mathbf{x}, \ell, t_k | s) = P_D \int_{\text{FOV}_s} p_{k|k-1}(\psi_s, \ell) d\psi_s. \quad (5.17)$$

Since the distribution resulting from the unscented transform is a bivariate Gaussian (mixture), an efficient numerical algorithm by Genz [51] is implemented to easily obtain the cumulative probability. The reader is referred to the original paper for algorithm details.




Figure 5.6: Detection probability based on field of view. Projecting the probability distributions onto the image plane and integrating over the FOV yields the probability that the target was in the FOV at the time of the measurement.

5.2.3 Clutter model

The clutter measurements for the optical sensor are modelled as in Eq. (3.21), using a Poisson assumption on the number of clutter measurements per image and allowing the clutter measurement to appear anywhere in the FOV with uniform intensity. This means that $\kappa_c(z) = 0$ if $z \notin \text{FOV}$.

Since zero clutter intensity is detrimental to the filter implementation¹⁸, it is best to filter out any measurement that falls outside the configured field of view, as it is by definition never expected to occur.

5.2.4 Optical tracklets

Because SSA sensors have a limited FOV compared to the orbital state space, it is common to see a particular object multiple times in quick succession before it leaves the field of view, returning only hours or even days later¹⁹. In those short *passes* through the FOV, the topocentric right ascension/declination measurements (α, δ) often follow linear or curvilinear motion in only 2 dimensions and it is therefore possible to group them into *tracklets*, without even resorting to orbital dynamics. This process of *tracklet formation* is well-established and commonly used in traditional space cataloguing approaches. Due to the lower dimensionality and simplified motion model, outliers and false-positives can already be removed before considering them in the orbital tracking setting. In Section 5.3, these tracklets will be used to drive the birth model. Later, in Chapter 6, *tracklets* are leveraged further by considering them in full individual measurements in the *tracklet LMB filter*.

5.3 Initial orbit determination - space birth models

In many other applications, a common choice of birth model for the LMB filter is to distribute the means of new Bernoulli components uniformly over the state space, or over a part of it where new objects might enter the filter.

¹⁸Recall the clutter associations are represented implicitly by dividing all association scores in Eq. (3.23) by the clutter intensity and cancelling them out afterwards. This results in $\log(0)$ or division by 0 operations and undefined behaviour.

¹⁹Note that this holds specifically if the observer is in *surveillance mode*, tracking as many objects as possible without any preference for individual satellites.

For example, consider maritime tracking using RADAR. The birth model could place new LMB components along the boundary of the detection range, as this is where new ships are expected to appear. The velocity mean could point towards the observer and the state covariance can account for all options. With a limited number of components, it is possible to represent all possible new objects.

This works if the birth region is well-bounded and of limited dimensionality, and if newborn objects are expected to appear with significant separation (a single labelled birth component can represent at most 1 real object). Most importantly, this requires relatively dense measurements, so that the initial uncertainty has limited time to grow before the components are accepted or rejected. For these reasons, it is generally infeasible in SSA to use an uninformed multi-Bernoulli birth density over the entire domain; the space of possible orbital states is so large that such an approach is either insufficient or intractable. Instead, the unassociated measurements from a previous time step can be used to create an *adaptive* initial distribution for new targets [41][48][65][105]. The creation of an orbit state from a limited number of measurements is commonly referred to as *initial orbit determination (IOD)*.

Many different approaches exist to address this, some of which were mentioned in Section A.3. It is possible to initialise birth densities by combining uncorrelated tracklets (UCTs) at multiple time steps [14][116], but the focus in this work is on IOD based on a single UCT, i.e. a tracklet that has a low association probability for all of the existing targets in the MO filter. The reason multi-tracklet IOD methods are not used here is that it is hard to avoid introducing bias into the RFS framework. Since the same tracklet might contribute to multiple birth components, coordinating the relevant uncertainties and existence probabilities is more involved and considered beyond the scope of this thesis.

In the numerical simulations of Chapter 7, the advantages and shortcomings of two different orbit birth methods are highlighted – they are briefly introduced below. Section 5.3.1 describes a method using batch least-squares optimisation, which can be suitable for relatively long tracklets with low measurement noise. The other approach relies on the linearity of (α, δ) in short tracklets and uses constraints on the orbit geometry to obtain an initial distribution for the newborn target state (Section 5.3.2–A). Since both methods are well-established in the SSA literature, these sections are kept short and the reader referred to external sources for extensive details.

5.3.1 Gooding and batch least squares

One approach that is expected to work for IOD with relatively long tracklets uses the first, last and central observation in the tracklet to obtain an initial orbit estimate by Gooding’s method [52][53]. After that, the complete tracklet is used in a batch least squares (LS) optimisation to compute a single state and covariance matrix.

For this thesis, these methods were used as implemented by Orekit, and they are well described in existing literature (e.g. [26][48][124]) so only a brief overview of the main steps is provided here.

- (i) Choose three $(\alpha, \delta)^T$ measurements in the tracklet for the Gooding initial guess. These should be separated as far as possible.
- (ii) Find the position of the observer $\mathbf{r}_{\text{obs}(t)}$ at each of the measurement epochs
- (iii) Guess the initial slant range ρ from the observer to the target object at the first and last measurements (note that $\mathbf{r}_{\text{obs}}^T$ and $(\alpha, \delta, \rho)^T$ constrain the position of the target object).
- (iv) Now that two position vectors and the time of flight are known, use a Lambert solver [62] to fully constrain the orbit estimate.

- (v) Predict the remaining (central) measurement based on the current orbit estimate and compute the measurement residual.
- (vi) Update the slant range guesses iteratively in a Newton-Raphson optimisation until convergence of the central measurement prediction. An illustrative case is visualised in Figure 5.7.
- (vii) Feed the resulting initial orbit to a batch LS estimator that uses the complete dynamical model and all measurements in the tracklet to obtain a state and covariance matrix at the central time. Note that the covariance is usually approximated as the inverse Fisher information matrix, which constitutes the lower bound of the uncertainty for an unbiased estimator [69]. If the dynamics are imperfect, the measurement noise is larger than expected, or the observation errors in the tracklet are not Gaussian, the estimated covariance will be underestimated.

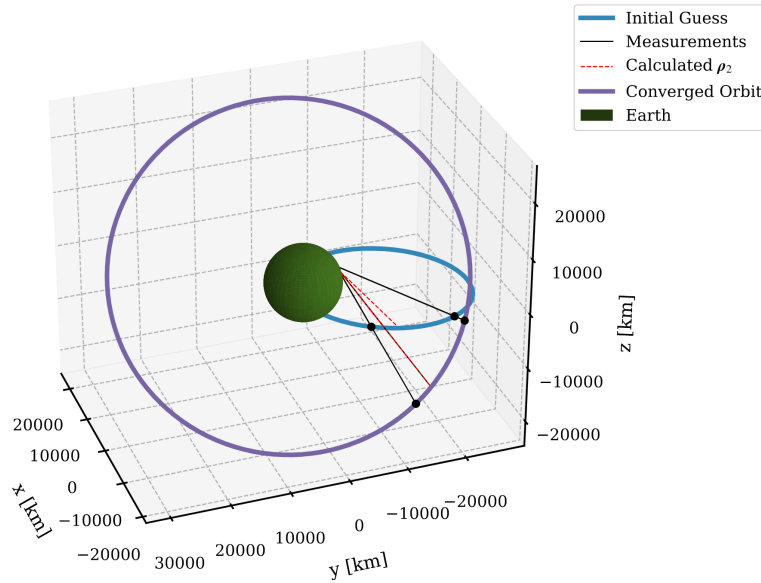


Figure 5.7: Example of initial orbit determination using Gooding's method. Note that the geometry in this example is unrealistic, as tracklets are generally not expected to be this long.

For sufficiently long tracklets where the evolution of (α, δ) is not linear over time, this approach is expected to produce a realistic orbit estimate (see Section 7.3). However, if tracklets are very short (only a small fraction of the orbit is observed) in combination with measurement noise, many potential orbits could have created the same tracklet. Since the batch estimator does not impose any constraints on the orbit geometry of the solution it produces, it often leads to very large uncertainties and even hyperbolic orbits due to the limited angular arc and lack of curvature information in these short tracklets. In such cases the single covariance matrix will not be an accurate representation of the orbit probability distribution, especially when propagated further in time.

In summary, this method can be a suitable choice for particular observation strategies, but should generally be treated with care as a birth process in the tracklet LMB filter.

5.3.2 Attributable and constrained admissible region

In those cases where the Batch approach does not have sufficient information to obtain a well-constrained orbit, the second approach implemented for this thesis can be a suitable choice. This method relies on the linearity of α and δ in short tracklets and uses pre-defined orbit restrictions

to define a space of possible states (*region*). An approximation by Gaussian mixture then ensures an accurate propagation of the probability distribution to be refined by future tracklets.

Given a sufficiently short tracklet of topocentric right ascension/declination measurements, it is possible to approximate it with a single *attributable* \mathbf{a} , located at the centre epoch of the original tracklet and containing the estimated angles, along with their rates. The attributable is denoted as

$$\mathbf{a} = \begin{pmatrix} \alpha & \delta & \dot{\alpha} & \dot{\delta} \end{pmatrix}^T, \quad (5.18)$$

where α and δ are the right ascension and declination, respectively. A linear LS estimation can be used to obtain a mean μ_a and covariance \mathbf{P}_a for the combined attributable measurement. However, to obtain an orbit density, there are two more unresolved dimensions. To completely define an orbit, an additional estimate is required for the *slant range* ρ and its rate $\dot{\rho}$ – that is, the unknown distance between the observer and the observed object.

Figure 5.8 visualises by example why a single attribute does not fully constrain the orbit of an observed RSO. In particular, different combinations of semi-major axis a and eccentricity e can result in the same angles and angular rates at specific time.

By putting bounds on the acceptable values of a and e , one can limit the number of possible solutions for ρ and $\dot{\rho}$ and thereby obtain a 6D initial orbit distribution.

A. Constrained admissible regions

The procedure of constraining the geometry of new orbits densities by putting bounds on the orbital parameters results in a *constrained admissible region (CAR)* [88].

The full development on how to compute the orbital elements for a given attributable and $(\rho, \dot{\rho})$ assumption is provided in Appendix D, along with how to invert this mapping to obtain bounds on the slant range and its first derivative based on constraints in a and e . An example CAR for an attributable observed by a ground-based telescope is shown in Figure 5.9, where the semi-major axis is constrained to be smaller than that of a GEO orbit and the eccentricity is limited to < 0.4 .

Plotting a sample of the orbits that are accepted by the CAR results in Figure 5.10, where it is clear that the inclination is very well constrained, but many different combinations of a and e cause in large position uncertainty over time. The distribution of semi-major axis and eccentricity for the solutions inside the CAR is visualised in Figure 5.11.




Figure 5.8: Identical attributable for very different orbits. In this example (not to scale), ρ is the same for both orbits and only $\dot{\rho}$ differs.




Figure 5.9: Example constrained admissible region with ground-based observer. This example follows the settings as discussed in [28, p. 1326] the observer is located at 30deg latitude and the attributable observation is $\mathbf{a} = (10 \text{ deg}, -2 \text{ deg}, 15 \text{ deg/h}, 3 \text{ deg/h})$.




Figure 5.10: Example ground-based CAR – The distribution of orbits sampled from the admissible region and seen from two points of view.

B. Gaussian mixture approximation

Provided that no other information is available on the distribution of orbital parameters, the target state density is assumed to be uniform in $(\rho, \dot{\rho})$ over the entire CAR. To be able to use the closed-form Kalman filter equations (recall Section 2.2.4), it is useful to approximate the uniform region as a GM. The process to achieve this is detailed by DeMars and Jah [28], but can be briefly summarised as follows:

- (i) Compute the range-marginal PDF $p_\rho(\rho) = \int_{-\infty}^{\infty} p_{\rho, \dot{\rho}}(\rho, \dot{\rho}) d\dot{\rho}$, where $p_{\rho, \dot{\rho}}$ is the complete, uniform PDF of the constrained admissible region.
- (ii) Based on a desired slant range uncertainty σ_ρ , choose the appropriate number of range components L_ρ from [28, Table 1]²⁰ such that $\tilde{\sigma}_\rho \lesssim \sigma_\rho$, distribute them equally over the

²⁰Note that Table 1 in [28] is unitless, so the design uncertainties σ_ρ and $\sigma_{\dot{\rho}}$ need to be divided by the extent of the domain and later multiplied again. Additionally, if the domain is very large or the desired σ very small, the maximum $L = 15$ is used and the standard deviation adjusted accordingly.




Figure 5.11: Example ground-based CAR – distribution of semi-major axis and eccentricity for accepted $(\rho, \dot{\rho})$.

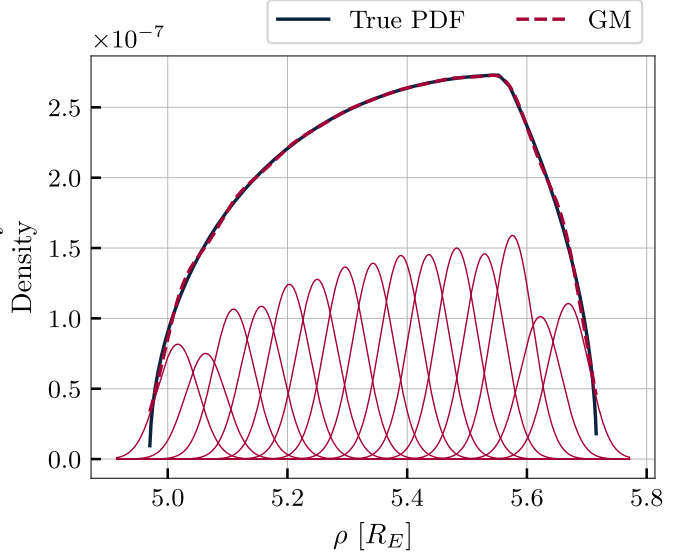


Figure 5.12: Example ground-based CAR – range-marginal probability density function and its approximation by GM for $\sigma_\rho = 220$ km.

ρ domain and estimate the GM weights $\{w_\rho^{(i)} : 1 \leq i \leq L_\rho\}$ by linear least squares to approximate the range-marginal PDF as shown in Figure 5.12.

- (iii) For each of the slant range means ρ_i , compute the local domain of the range rate and create $L_{\dot{\rho}}^{(i)}$ components with equal spacing for a desired $\sigma_{\dot{\rho}}$ such that $\sigma_{\dot{\rho}}^{(i)} \lesssim \sigma_{\dot{\rho}}$. This time, the distribution is conditioned on ρ_i and thus uniform, so all GM weights are equal to $w^{(i,j)} = \frac{1}{L_{\dot{\rho}}(\rho_i)}$.
- (iv) The complete 2D GM is constructed by combining all the components as

$$\mathcal{G}_{\rho, \dot{\rho}} = \left\{ \left[w_\rho^{(i)} w_{\dot{\rho}}^{(i,j)}, \begin{pmatrix} \rho_i \\ \dot{\rho}_{i,j} \end{pmatrix}, \begin{pmatrix} \tilde{\sigma}_\rho & 0 \\ 0 & \tilde{\sigma}_{\dot{\rho}}^{(i)} \end{pmatrix} \right] \right\}_{(i,j) \in L_\rho \times L_{\dot{\rho}}^{(i)}}, \quad (5.19)$$

and depicted for the present example in Figure D.8.

By combining this GM with the estimated attributable density $\mathcal{N}(\boldsymbol{\mu}_a, \mathbf{P}_a)$, a complete 6D GM distribution is obtained

$$\mathcal{G}_{a, \rho, \dot{\rho}} = \left\{ \left[w_{\rho, \dot{\rho}}^{(i,j)}, \begin{pmatrix} \boldsymbol{\mu}_a \\ \boldsymbol{\mu}_{\rho, \dot{\rho}}^{(i,j)} \end{pmatrix}, \begin{pmatrix} \mathbf{P}_a & \mathbf{0}_{4 \times 2} \\ \mathbf{0}_{2 \times 4} & \mathbf{P}_{\rho, \dot{\rho}}^{(i,j)} \end{pmatrix} \right] \right\}_{(i,j) \in L_\rho \times L_{\dot{\rho}}^{(i)}}, \quad (5.20)$$

which can in turn be mapped to the orbital state space of choice using an unscented transform (recall Eq. (2.30)).

This method is implemented in BASIL following the detailed descriptions by DeMars and Jah [28][30] and Fujimoto and Scheeres [44], and based on the implementation in the Multitarget Estimation, Tracking and Information Synthesis (METIS) Python library by Gehly [47][49]. Two validation cases that compare to these references are also provided in Appendix D.

Since the CAR approach allows for the use of a single uncorrelated tracklet to form a birth target and one can represent the state PDF by a GM, this model is consistent with the current LMB

implementation in BASIL and chosen as the main method to model the birth process in the numerical simulations of [Chapter 7](#).

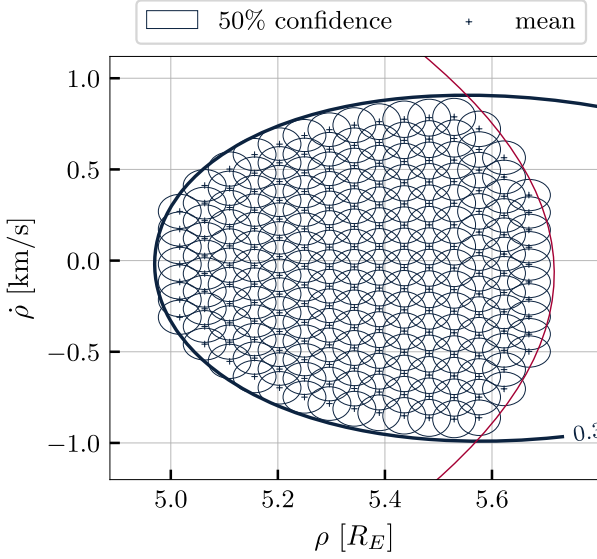


Figure 5.13: Example ground-based CAR
– Gaussian mixture approximation with
 $\sigma_\rho = 220$ km and $\sigma_{\dot{\rho}} = 80$ m/s.




Figure 5.14: Example ground-based CAR
– Gaussian mixture approximation with
 $\sigma_\rho = 220$ km and $\sigma_{\dot{\rho}} = 80$ m/s.

C. Probabilistic admissible regions

Although the CAR approach already provides a relatively robust framework to get an initial orbit density, [\[110\]](#) points out that the uniform region with hard constraints is not the best representation of our knowledge on the distribution of RSOs. Instead, the known statistics on the spread of a , eccentricity e , inclination i and right ascension of the ascending node Ω can be combined with the uncertainty in the (α, δ) observations and to form a more informed probabilistic admissible region (PAR). [Cament et al. \[17\]](#) uses this method inside a Poisson LMB filter with promising results.

As discussed in the next chapter, CARs do not always perform well for LEO targets observed by space-based sensors. Adapting to a probabilistic admissible region (PAR) approach might improve the behaviour in such scenarios but the implementation and analysis of such solutions for the tracklet LMB filter is left for future work.

5.3.3 Adaptive birth with tracklets

The previous two sections described how a single-object density can be obtained from an optical tracklet. What remains is to assign a label and existence probability to all components, so that they are completely defined labelled Bernoullis. For simplicity, let the birth component produced by a tracklet with label \mathbf{m} be labelled ℓ_m . The existence probability $r_B^{(\ell_m)}$ can be assigned using the *adaptive birth* approach ([Section 3.3.4–G](#)), by either [Eq. \(3.36\)](#) or [Eq. \(3.37\)](#).

Note that both equations rely on $r_{U,k-1}(\mathbf{z})$ from [Eq. \(3.35\)](#) – the probability that a measurement \mathbf{z} was unassigned in the previous time step. Using the development from [Chapter 4](#), there are two options to generalise this to a tracklet with label \mathbf{m} , depending on the type of the LMB filter:

- (i) Either tracklets are processed *as measurements* in the filter (this is the topic of [Chapter 6](#)), in which case $r_{U,k-1}$ is exactly equivalent to the *time-marginal association probability* $\bar{\alpha}_{k-1}^{(m)}$ from [Eq. \(4.2\)](#).
- (ii) Otherwise, if the traditional LMB filter is used, a tracklet should only be added to the birth model as soon as the filter has processed all its measurements (i.e. all measurements with label m). The probability that this tracklet is a UCT and thus not associated to any existing label is then the complement $\bar{\alpha}_{1:k-1}^{(m)}$ of the *accumulated association probability* from [Eq. \(4.7\)](#). Here t_{k-1} is the time of the last measurement in the tracklet and the birth component is introduced at t_k . It should be mentioned that none of the test cases in [Chapter 7](#) use a birth model in combination with the single-measurement LMB filter. Thus, this option is only theoretical and not experimentally verified in the scope of this work.

To limit the number of birth components, the birth model is only run for tracklets with $r_B^{(\ell_m)}$ above some minimum threshold $r_{B,\min}$. At a minimum, this parameter is equal to the LMB pruning threshold ϑ_P^{LMB} and birth components with $r_B^{(\ell_m)}$ below that value can be *pre-pruned*, before computing the corresponding SO density.

THE TRACKLET LMB FILTER

In the case of RSOs, it is expected to see the same object multiple times in close succession, after which it leaves the telescope’s field of view, only to return hours or days later²¹. As discussed in [Section 5.2.4](#), it is common to first group the measurements from single passes through the FOV into *tracklets*, and then associate the tracklets to existing RSO orbits or to each other (e.g. [\[16\]](#) [\[43\]](#) [\[99\]](#) [\[116\]](#) [\[147\]](#)).

In the conventional MOT setting, measurements are chronologically fed to the filter in the form of scans Z_k , where each $z_k \in Z_k$ is valid at the exact same time t_k and the measurements themselves contain no information that could associate them with observations at different times. This means that measurement labels obtained by tracklet formation methods are entirely ignored. Recall that tracklet formation is usually reliable as it can be solved in fewer dimensions and with (near-)linear dynamics.

Some efforts to introduce the benefits of tracklets in MOT were made by [Frueh et al. \[41\]](#). Their paper describes a multi-stage PHD filter where measurements in fixed intervals are associated into tracklets and the tracklets are used as a whole to obtain the posterior PHD. That method incorporates the tracklet formation and orbit determination as two stages of the same algorithm, and assumes that all tracklets overlap completely with a clear start and end time. However, enforcing time intervals on all the tracklets inevitably results in some of them being split unnecessarily. The method proposed by Frueh also does not label the estimated targets and does not retain the association result.

In what follows, a novel approach is introduced to adapt the LMB filter so that it can accept complete, pre-formed *tracklets as measurements*. Additionally, the method adapts its time steps dynamically based on the available tracklets. The expected benefit compared to the traditional LMB is threefold:

- (i) Tracklet formation can be done in the image space, effectively reducing the complexity of short-term measurement association to a linear problem in two dimensions.
- (ii) Small variations in time stamps due to shutter time effects are more easily handled, as measurements do not need to be at the exact same time [\[41\]](#).
- (iii) The effect of individual outliers is reduced, as entire tracklets are associated rather than individual detections.

[Section 6.1](#) first details how all elements in the LMB iteration can be adapted to facilitate *groups of tracklets* instead of *scans of observations*. After that [Section 6.2](#), addresses how these tracklet groups can be created to ensure that the tracklets in each group represent distinct RSOs, while maximising the number of tracklets in every step.

6.1 Tracklets as measurements

This approach leaves as much as possible from the original (G)LMB formulation intact, but aims to substitute the measurement vectors $z \in Z_k$ by tracklets $\mathfrak{Z}_{a,b}^{(m)}$ (i.e. labelled sequences of measurements). This sacrifices some of the mathematical rigour of the GLMB, particularly

²¹The time between observations is highly dependent on the pointing strategy of the observer, as well as the target’s orbit and the observation geometry.

in relation to modelling clutter and target death, but offers an effective framework for tracklet association.

Let $\mathfrak{T}_{a,b}^{(m)} \triangleq (\mathfrak{T}_{a,b}, m) \triangleq \{z_a^{(m)}, \dots, z_b^{(m)}\}$ be a labelled tracklet with sorted observation set $\mathfrak{T}_{a,b}$ and label $m \in \mathbb{M}_k$, such that $t_i \in [t_a, t_b]$ for every observation $z_i^{(m)} \in \mathfrak{T}_{a,b}^{(m)}$ and assume that the duration of such a tracklet $t_b - t_a$ is always very short with respect to the targets orbital period.²²

Now assume it is possible to form a group of these tracklets that are guaranteed to originate from different targets and denote it as

$$\mathfrak{Z}_k = \left\{ \mathfrak{T}_{a,b}^{(m)} \right\}_{m \in \mathbb{M}_k}, \quad (6.1a)$$

$$\text{s.t. } \forall m \in \mathbb{M}_k : [t_a^{(m)}, t_b^{(m)}] \subset]t_{k-1}, t_k], \quad (6.1b)$$

where \mathbb{M}_k is the set of all measurement labels in \mathfrak{Z}_k and no other available tracklet has measurements in this time interval. [Section 6.2](#) suggests a simple method to form valid groups in most cases, but for now assume the mentioned requirements are satisfied. If this is the case, the *tracklet group* \mathfrak{Z}_k can be used as substitute for the measurement scan Z_k from the original LMB filter.

Since there are now measurements at multiple epochs, the *predicted* and *corrected* filtering densities will be defined at different points in time. The following three times are relevant for the tracklet LMB filter:

- t_{k-1} is the time of the previous corrected MO density
- $t_k^* = \min_{m \in \mathbb{M}_k} t_a^{(m)}$ is the time of the first observed measurement in the tracklet group, here referred to as the *prior time*.
- $t_k = \max_{m \in \mathbb{M}_k} t_b^{(m)}$ is the time of the last observed measurement in the tracklet group and thereby also the epoch of the next corrected MO density. It is here called the *posterior time*.

The LMB recursion for this tracklet filter can then be summarised with some minor modifications compared to the original LMB filter:

- (i) Given the previous LMB density $\hat{\pi}_{k-1}^{\text{LMB}}$ at time t_{k-1} .
- (ii) Predict the surviving LMB at the prior time t_k^* : $\hat{\pi}_{S,k^*|k-1}^{\text{LMB}}$ (see [Section 3.3.4-A](#)).
- (iii) Add birth components $\hat{\pi}_{B,k^*}^{\text{LMB}}$ from an adaptive birth model (see [Section 5.3](#)).
- (iv) Convert to GLMB: $\hat{\pi}_{k^*|k-1}^{\text{GLMB}}$ (see [Section 3.3.4-B](#)).
- (v) Compute the detection probabilities for all targets $\ell \in \mathbb{L}_k$ during the interval $[t_k^*, t_k]$ (see [Section 6.1.1](#))
- (vi) For every viable combination of measurement m and target ℓ , compute the tracklet measurement log likelihood (see [Section 6.1.2](#))
- (vii) Compute the weights $w^{(\xi_\theta)}(L)$ for these updates using the normal GLMB update step, but employing the adapted detection, clutter and likelihood models.
- (viii) Compute the posterior densities for each hypotheses by running a single-object filter (e.g. EKF/UKF) through all the observations and propagating further to the posterior time t_k . Note that undetected targets must also be propagated to t_k to ensure the posterior MO density $\hat{\pi}_k^{\text{GLMB}}$ is defined at a single point in time.
- (ix) Convert the posterior to LMB: $\hat{\pi}_k^{\text{LMB}}$ (see [Section 3.3.4-D](#)).
- (x) Set $k := k + 1$ and return to (i).

²²For optical tracklets obtained during general surveillance, the duration is often more than two orders of magnitude shorter than the target orbital period.

6.1.1 Tracklet detection probability

The procedure in [Section 5.2.2](#) showed how the detection probability is obtained for a discrete observation time t_k , but since the update now happens in a time window $[t_k^*, t_k]$, the detection probability must be adapted accordingly.

Assuming the sensor still collects measurement scans at discrete times, then the illustration in [Figure 5.6](#) can be extended to [Figure 6.1](#). The relevant quantity is the probability that each target was observed during *at least one* of these scans.




Figure 6.1: Multi-scan detection probability with field of view. This shows the target densities projected onto the FOV at all measurement scan epochs. The

If $P_D \lesssim 1$ for an object inside the FOV, the most straightforward solution to obtain the multi-scan detection probability is to use the maximum value

$$P_D(\mathbf{x}_{k^*}, \ell, \{t_k^*, \dots, t_k\}) = \max_{t_j \in \{t_k^*, \dots, t_k\}} P_D(\varphi_j(\mathbf{x}_{k^*}), \ell, t_j) \quad (6.2)$$

where φ_j represents the predicted state (along with its uncertainty) to time t_j and $P_D(\mathbf{x}, \ell, t)$ follows [Eq. \(5.17\)](#).

However, for lower P_D , this method will underestimate the detection probability of targets that remain in the FOV longer. Instead, compute the probability that the target is outside the field of view during *each* of the measurement scans. The complement is then the probability that the target is observed *at least once*,

$$P_D(\mathbf{x}_{k^*}, \ell, \{t_k^*, \dots, t_k\}) = 1 - \prod_{j=k^*}^k [1 - P_D(\varphi_j(\mathbf{x}_{k^*}), \ell, t_j)]. \quad (6.3)$$

Example – Consider target ℓ_2 in [Figure 6.1](#) and assume the detection probability in the FOV is $P_D = 0.5$. During this time span, there are five observation opportunities and the predicted target mean is inside the FOV at 4 of those times. According to [Eq. \(6.2\)](#), the detection probability would now still be $P_D(\mathbf{x}, \ell_2, \{t_1, \dots, t_5\}) \approx 0.5$, which is not realistic. The probability that it was not detected in any of the scans is $(0.5)^4 \approx 0.06$, so the probability of detection according to [Eq. \(6.3\)](#) becomes 0.94.

Recall from Eq. (3.23) that the measurement assignment score has $1 - P_D$ in the denominator, so $P_D(\cdot) = 1$ results in a singularity and should be avoided. Therefore, an upper bound P_D^{\max} (e.g. 0.99) is used so that

$$P_D(\dots) = \min(P_D^{\max}, P_D(\dots)), \quad (6.4)$$

which also explains why Eq. (6.3) does not make much difference if P_D is already close to 1. Repeated presence of a target in the FOV would anyway result in $P_D(\cdot) = P_D^{\max}$.

6.1.2 Tracklet log likelihood

To compute the *measurement likelihood* for tracklet measurement given a target ℓ and its prior distribution, one approach is to evaluate all measurements in that tracklet separately and then use the average log likelihood.

$$\log g(\mathfrak{T}^{(m)} | \mathbf{x}_{k^*}, \ell) = \frac{1}{|\mathfrak{T}^{(m)}|} \sum_{\mathbf{z} \in \mathfrak{T}^{(m)}} \log g(\mathbf{z} | \mathbf{x}_{k^*}), \quad (6.5)$$

where $g(\mathbf{z} | \mathbf{x}_{k^*})$ is a minor misuse of notation to indicate that the prior state density is predicted to the time of \mathbf{z} after which the likelihood is computed. For Gaussian (mixture) distributions, this effectively computes the Kalman filter equations without actually updating the distribution. Note that Eq. (6.5) is conceptually similar to estimating the likelihood based on the mean squared Mahalanobis distance (recall Eq. (2.8) and Eq. (2.19)).

6.1.3 Tracklet clutter model

Modelling false positives and measurements from unknown objects is the most challenging part of the tracklet filter. To fit the core assumptions of the MO measurement model from Eq. (2.23), the clutter must be modelled as a Poisson RFS. However, since a tracklet is in itself a set of vectors, the tracklet clutter intensity function $\kappa_c^{\mathfrak{T}}$ will have different units depending on the number of measurements in the tracklet. It is therefore necessary to also model the size of clutter tracklets.

Moreover, it is not immediately clear when a tracklet is *clutter*. One possible definition could be: “a tracklet is clutter if it consists *exclusively* of clutter measurements”, but no less valid would be “the tracklet is clutter if it contains *at least* 1 clutter measurement” and many other options. Neither of these two is ultimately used in this work, but for the sake of argument, the first definition is briefly developed.

Let $\rho_c^{\mathfrak{T}}$ be a cardinality PMF, such that $\rho_c^{\mathfrak{T}}(n)$ is the probability that a *clutter tracklet* has n measurements. Further, assume that individual clutter measurements follow a PPP distribution with intensity κ_c and clutter rate λ_c . Then, the tracklet clutter intensity can be modelled as

$$\kappa_c^{\mathfrak{T}}(\mathfrak{T}^m) = \rho_c^{\mathfrak{T}}(|\mathfrak{T}^{(m)}|) \prod_{\mathbf{z} \in \mathfrak{T}^{(m)}} \kappa_c(\mathbf{z}), \quad (6.6)$$

and the corresponding tracklet clutter rate is

$$\lambda_c^{\mathfrak{T}} = \sum_{n=1}^{\infty} \left[\int_{\mathbb{Z}^n} \kappa_c^{\mathfrak{T}}(\mathfrak{T}') \delta \mathfrak{T}' \right] \quad (6.7a)$$

$$= \sum_{n=1}^{\infty} \rho_c^{\mathfrak{T}}(n) (\lambda_c)^n, \quad (6.7b)$$

where the set integral is as discussed in Eq. (3.2). If clutter tracklet length is itself modelled by a Poisson distribution with mean $\bar{n}_{\mathfrak{T}}$, then Eq. (6.7) simplifies to

$$\lambda_c^{\mathfrak{T}} = \sum_{n=1}^{\infty} \frac{\exp(-\bar{n}_{\mathfrak{T}}) (\bar{n}_{\mathfrak{T}})^n}{n!} (\lambda_c)^n = \exp(-\bar{n}_{\mathfrak{T}}) \sum_{n=0}^{\infty} \frac{(\bar{n}_{\mathfrak{T}} \lambda_c)^n}{n!} \quad (6.8a)$$

$$\stackrel{\text{Taylor}}{=} e^{-\bar{n}_{\mathfrak{T}}} (e^{\lambda_c \bar{n}_{\mathfrak{T}}} - 1) \quad (6.8b)$$

It is important to stress that $\bar{n}_{\mathfrak{T}}$ *does not* represent the mean length of any tracklet, but of a *clutter tracklet* as defined by the first definition above. It is the mean number of clutter measurements that would be linked together by the preprocessing step.

Although this representation is mathematically consistent and suitable as a direct substitute following the assumptions of the GLMB filter, it has several very important drawbacks:

- (i) It only models the specific *failure mode* where a complete tracklet is full of clutter measurements, but this is not sufficient. For example, the tracklet formation process might combine measurements from different objects in the same (long) tracklet, which would not be captured by Eq. (6.6).
- (ii) It is incompatible with adaptive birth (see Section 3.3.4–G), which relies on the mechanism of association probabilities to select tracklets for newborn targets. Since the MO measurement model does not distinguish between clutter measurements and measurements belonging to an unknown object, the *clutter model* must capture both scenarios.

The last drawback of course also applies to a model that considers tracklets with ≥ 1 clutter measurement as false positive (second definition above). In addition, such a model would be unreasonably quick to discard longer tracklets and has other downsides that make it undesirable.

In the context of this work, the clutter model is mainly used to accommodate the measurement-driven birth model and a heuristic method is used as a first attempt, where the tracklet clutter intensity is set to the geometric mean of its measurement clutter intensities and the tracklet clutter rate is set to the clutter rate:

$$\lambda_c^{\mathfrak{T}} = \lambda_c, \quad \kappa_c^{\mathfrak{T}} = \exp \left[\frac{1}{|\mathfrak{T}|} \sum_{z_i \in \mathfrak{T}} \log \kappa_c(z_i) \right]. \quad (6.9)$$

This is a sub-optimal approach for many reasons; most importantly because it violates²³ the definition of the PPP clutter model and therefore introduces a bias in the GLMB update. Rigorously quantifying this bias is also not possible without clearly defining all the types of false-positive tracklet measurements. The problem of accurately and consistently modelling tracklet clutter is now briefly introduced but further investigation is considered beyond the scope of this thesis and will be left for future work.

Aside from the mentioned limitations in mathematical consistency, it should be stressed that the tracklet formation process should significantly reduce the clutter rate, and the expected number of false-positives is already relatively low for optical observations in SSA (compared to e.g. RADAR). For this reason, it is expected that the effective clutter rate with tracklets as measurements will be close to zero and the clutter model can be exploited mostly to represent

²³Eq. (6.9) does not have units $[\frac{1}{\mathbb{Z}^n}]$ (where $n = |\mathfrak{T}|$) and thus does not define a proper intensity function. This breaks the additivity and total expectation property from Eq. (6.7a).

unknown objects that will later be introduced by the adaptive birth model from [Section 3.3.4–G](#) and [Section 5.3](#).

6.2 Tracklet grouping

To build the LMB filter with tracklets as measurements, it was assumed that tracklet groups \mathfrak{Z}_k can be formed such that

- (i) they contain tracklets which certainly belong to different objects and
- (ii) no two tracklets in different groups have overlapping measurement times.

This section proposes a simple algorithm to obtain such groups, based on the sole assumption that there is some minimum time span Δt_{valid} between tracklets that originate from the same target. In other words, given two tracklets with labels m_1 and m_2 , if

$$\left[t_a^{(m_1)}, t_b^{(m_1)} + \Delta t_{\text{valid}} \right] \cap \left[t_a^{(m_2)}, t_b^{(m_2)} + \Delta t_{\text{valid}} \right] \neq \emptyset, \quad (6.10)$$

then m_1 and m_2 are guaranteed to belong to different objects and can be in the same tracklet group. This *validity padding* Δt_{valid} is the only design parameter for the grouping method and can be must be adjusted to context of the problem. Note that the validity padding is dependent on the specific tracking scenario and particularly on the sensor tasking. This thesis is limited to single-observer cases, so Δt_{valid} mostly depends on the target objects regime, the observer position/orbit and its pointing strategy.

Given a set of tracklets $\mathfrak{T} = \left\{ \mathfrak{T}_{a,b}^{(m)} \right\}_{m \in \mathbb{M}_{1:k}}$ ordered by $t_a^{(m)}$, [Figure 6.2](#) visualises how they are then robustly collected into measurements groups. In general, the first tracklet defines t_k^* , and tracklets are added to the group as long as they satisfy [Eq. \(6.10\)](#) with the first one. If not, a new group is started. In case the first excluded tracklet overlaps with a tracklet in the group – violating requirement (ii), the tracklet after the last gap marks the start of a new group.




Figure 6.2: Tracklet grouping example. m_1 defines the start of \mathfrak{Z}_1 . m_2 and m_3 are added because their start time overlaps with the validity padding of m_1 . m_4 does not overlap and therefore starts a new group \mathfrak{Z}_2 . m_5 , m_6 and m_7 could be added to that group, however m_8 cannot and overlaps with m_7 . Since this would violate the requirement (ii), m_7 marks the start of a new group.

The complete procedure is also reflected in [Algorithm 6.1](#), where the following abbreviations are used in the notation for convenience:

$$\mathfrak{T} \triangleq \mathfrak{T}_{a,b}^{(m)}, \quad t_a(\mathfrak{T}) \triangleq t_a^{(m)}, \quad \text{and} \quad t_a(\mathfrak{Z}) \triangleq \min_{\mathfrak{T} \in \mathfrak{Z}} t_a(\mathfrak{T}). \quad (6.11)$$

Algorithm 6.1: Simple tracklet grouping

```

GROUP( $\mathfrak{T} \triangleq \{\mathfrak{T}_{a,b}\}, \Delta t_{\text{valid}}$ ):
1   $\mathfrak{Z} = \emptyset$                                 ▷ groups
2   $\mathfrak{Z}_c = \emptyset$                             ▷ confirmed tracklets
3   $\mathfrak{Z}_p = \emptyset$                             ▷ pending tracklets
4  sort  $\mathfrak{T}$  by start time  $t_a(\mathfrak{T})$ 
5  for  $\mathfrak{T} \in \mathfrak{T}$ 
6      if  $\mathfrak{Z}_d = \emptyset$                         ▷ Initialise
7           $\mathfrak{Z}_c = \{\mathfrak{T}\}$ 
8          continue
9      end if
10
11     if  $t_a(\mathfrak{T}) < t_b(\mathfrak{Z}_c)$                 ▷ confirmed overlap
12          $\mathfrak{Z}_c = \mathfrak{Z}_c \cup \{\mathfrak{T}\}$ 
13     else if  $t_a(\mathfrak{T}) < t_b(\mathfrak{Z}_p)$             ▷ pending overlap
14          $\mathfrak{Z}_p = \mathfrak{Z}_p \cup \{\mathfrak{T}\}$ 
15     else                                        ▷ New gap; all pending confirmed
16          $\mathfrak{Z}_c = \mathfrak{Z}_c \cup \mathfrak{Z}_p$ 
17          $\mathfrak{Z}_p = \{\mathfrak{T}\}$ 
18     end if
19
20     if  $t_a(\mathfrak{T}) < t_a(\mathfrak{Z}_c) + \Delta t_{\text{valid}}$     ▷ New group needed
21         if  $\mathfrak{Z}_p = \emptyset$                     ▷ Continuously overlapping
22             throw algorithm fails ▷ grouping ambiguous
23         end if
24          $\mathfrak{Z} = \mathfrak{Z} \cup \{\mathfrak{Z}_c\}$                 ▷ register confirmed group
25          $\mathfrak{Z}_c = \mathfrak{Z}_p$                         ▷ pending → new group
26          $\mathfrak{Z}_p = \emptyset$                     ▷ No gaps yet
27     end if
28
29 end for
30  $\mathfrak{Z}_c = \mathfrak{Z}_c \cup \mathfrak{Z}_p$                     ▷ Collect all remaning
31 if  $\mathfrak{Z}_c \neq \emptyset$ 
32      $\mathfrak{Z} = \mathfrak{Z} \cup \{\mathfrak{Z}_c\}$                 ▷ Register final group
33 end if

```

Using this grouping method, sets of tracklets can robustly be incorporated as substitutes for the traditional *measurement scans* in the new tracklet LMB filter, leveraging both the rigour of the LRFS framework and the reliability of existing tracklet formation pipelines.

SIMULATIONS

Various test cases are investigated to analyse the effectiveness of the association extraction methods and the tracklet LMB filter. Every test case has several points of focus for investigation, as stated at the start of the section. Still, the main objectives are to demonstrate that measurement association results can be extracted from the LMB filter as worked out in [Chapter 4](#) and to illustrate the properties of the adapted LMB filter using angles-only tracklets as measurements. Each scenario’s case setup is briefly summarised in [Table 7.1](#).

Table 7.1: Overview of simulation cases

Focus	Observer type	Pointing strategy	# Objects	Regime	Section
Association extraction	ground (TFRM)	2 h GEO fences	8	GEO	7.1
Closely spaced objects	$i = 0^\circ$ (500 km)	Zenith	7	GEO	7.2
GEO object discovery	$i = 97^\circ$ (500 km)	GEO equinox line	200	GEO	7.3
LEO object discovery	$i = 97^\circ$ (500 km)	Above the poles	10	LEO	7.4

7.1 Tracklet to tracklet association in GEO

The first test case borrows from the setup discussed by [Pirovano \[98, p. 59 et seq.\]](#), who in turn adapted it from [Zittersteijn et al. \[148\]](#). This scenario with measurements from a ground-based telescope serves a triple purpose:

- (i) to evaluate the algorithm’s ability to associate tracklets from a small number of objects in the geostationary regime,
- (ii) to confirm the usefulness of the CAR birth model for the LMB filter with objects in GEO and
- (iii) to validate the filter convergence when detection is certain.

7.1.1 Setup

The object population for this case is a set of 8 satellites with ground truth initial states as reported in [Appendix G](#). All measurements are simulated for January 2016 for a single ground-based observer at the location of the Telescope Fabra ROA Montsec (TFRM) observatory, where the pointing strategy is to use a *fence* that keeps the right ascension fixed for 2 hours, after which the fence is moved 2 hours in right ascension. The process repeats 4 times every night during three consecutive nights. Measurement scans are collected every 30 seconds for a range of declination values in such a way that every object is certainly observed in every fence and no two tracklets in a single fence can originate from the same object. Given that the telescope has a field of view of 1 degree, the resulting tracklets contain about 7 observations (3 minutes). The fences and all observations are visualised in [Figure 7.1](#). The measurement noise is assumed Gaussian with $\mathbf{R} = (2'')^2 \mathbf{I}$.




Figure 7.1: Observation geometry for ground-based GEO observations with fences and all 3 nights of observations (based on [18][98][148])

This case has several advantages for an initial performance check of the tracklet filter. First of all, it avoids the need for dedicated grouping, because this pointing strategy immediately satisfies the tracklet group requirements listed in Section 6.2, so that the measurements can simply be grouped by observation fence. The fact that every object is observed in every group also allows using a constant detection probability $P_D = 0.99$.

The filter is initialised without any prior information and the birth process uses a constrained admissible region approach with $0.95 a_{\text{GEO}} < a < 1.05 a_{\text{GEO}}$ and $e < 0.1$ and GM approximation with desired uncertainty in slant range and range rate of $\sigma_\rho = 500$ km and $\sigma_{\dot{\rho}} = 20$ m/s, respectively. The existence probabilities are computed by the adaptive method from Eq. (3.37) with $\frac{\lambda_B}{\lambda_c} = 1$ and $r_{B,\max} = 0.3$. For consistency, the survival model uses these same CAR constraints as the birth model and assigns $p_S(\mathbf{x}, t_k) = 1$ if the state fulfils the constraints and 0 otherwise (recall Section 5.1.5). In case the state is a GM, the survival probability is a weighted sum. In principle, RSOs do not disappear and should therefore not die in the filter, but the mechanism of target death is, for example, needed to prevent that a failed birth component results in a persisting hypothesis if it never returns to the field of view.

7.1.2 Results

Because no initial information was provided to the filter, all 8 tracklets from the first fence immediately feed into the CAR birth process, which produces new labelled Bernoullis at the start of the second tracklet group. Since the constraints for the admissible region in this case are more stringent (particularly for the eccentricity) than in the example discussed in Section 5.3.2–A, the CAR generates, on average, 14 GM components per Bernoulli in the birth mixture and all but one or two of those are pruned out after only a few measurement updates.

The second-order OSPA error results are presented in Figure 7.2. Note that position and velocity components are separated to ensure consistent units in Eq. (3.38). Furthermore, the plot distinguishes between *estimated errors*, which occur after a filtering step of a single tracklet group, and *propagated errors* which are obtained simply by predicting the last estimated MO state forward and comparing it to the ground truth.

After the first tracklet group, there are no estimates because there are no target labels yet. Therefore the cardinality penalty is applied for each of the objects and the OSPA starts at cut-off values of 100 km and 100 m/s for position and velocity, respectively. One filter step later, the position OSPA is still close to the maximum, but now because the orbit estimates have not converged yet. To confirm this, Figure 7.3 shows the Euclidean position errors. The median estimated error is still above 200 km and none of the estimates are closer than ~ 40 km to the

truth. At the end of the first observation night, all errors are below 3 km, but since these are then propagated for about 18 hours, the OSPA also increases again. During the next 4 observation fences (the second night), all states are reduced enough to avoid growing errors in the second observation gap and the final position errors are at most ~ 400 m (Figure 7.3). These results demonstrate that the tracklet LMB filter integrates well with CAR birth and converges for a small number of objects in GEO.




Figure 7.2: OSPA errors in position and velocity for the small population in GEO with a ground-based telescope




Figure 7.3: Euclidean position errors of the corrected states w.r.t. ground truth. Note that this plot starts with the estimates after the second tracklet group, since there were no labels in the filter at the first time step.

Additionally, the single-object uncertainty bounds, estimated by the filter, correctly capture the estimation errors for each RSO. This is shown for object 26038 in Figure 7.4, but all others exhibit similar behaviour. That plot shows the RTN errors after every tracklet group and thus represents the update based on a complete tracklet. The covariance is large at first because it represents the uncertainty of the CAR. It then shrinks quickly as the estimates become better but no filter saturation occurs, indicating that the process noise is well-tuned. Recall that there are 12 tracklets for each object – one for all 4 fences in 3 consecutive nights, but the first night in Figure 7.4 contains only 2 estimates. As mentioned before, there cannot be an estimate for the first step, because the filter was initialised with an empty MO density, so the first tracklet must go to the adaptive birth model and only contributes in the prediction step for the next iteration of the filter (see Section 3.3.1). After the first update, its existence probability was still insufficient to be included in the state extraction (recall that $r_{B, \max} = 0.3$ was used here to avoid an overconfident birth process) and only after the second measurement update (i.e. the third tracklet group) did the filter consider this target *confirmed*. Remember that the decision to include targets in the MO state is made by estimating the number of objects \hat{n} based on the

mode of the cardinality distribution (MAP). Only the \hat{n} labels with highest existence probability $r^{(\ell)}$ result in an estimate (Eq. (3.31) and Eq. (3.32)). A target is considered *confirmed* if it is consistently included in the MO state estimate.

This is also consistent with the high OSPA in the beginning (Figure 7.2), due in part to the underestimated cardinality. To obtain accurate state estimates from the initial epochs onwards, one could perform a smoothing operation to go back in time with the newly obtained information. However, for this thesis, the focus is instead on collecting all the correct tracklet associations, such that a dedicated single-object estimation model can later be used to improve the orbit fit if required.

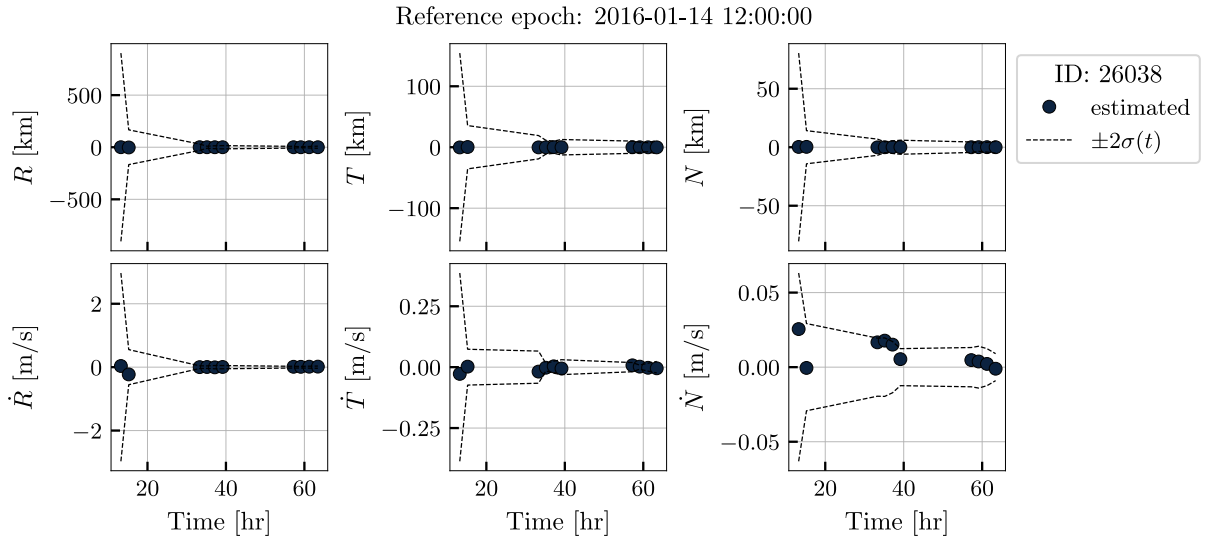


Figure 7.4: RTN filter estimation errors and uncertainty bounds for object 26038. Note that the uncertainty bounds are only valid at the estimation times and correspond to the *posterior* uncertainty. The apparent decrease in uncertainty during the measurement gaps is a plotting artefact.

Extracting the association results is done as discussed in Section 4.1. Since measurement labels are unique to each tracklet and the tracklet filter is used in this case, there is no need to accumulate the association probabilities. In contrast to direct tracklet-to-tracklet association methods, where the tracklets are evaluated two-by-two, this type of association is indirect as the observations are assigned to targets in the filter and therefore implicitly associated to all other measurements of that target label.

The complete output with association probabilities for every combination of tracklet label m and target label ℓ can be found in Figure F.1 of the appendix. Here, the performance is summarised by Figure 7.5. Since the tracklet labels contain the true ID of the source object²⁴, and the target labels are borrowed from the tracklet used for their IOD, this *confusion matrix* gives information on the correctness of the association estimates. In this case, no target has measurements from different sources and all tracklets from each source are assigned to a single target, so the association is perfect here.

²⁴This information is of course not used by the filter and only serves to evaluate the output and behaviour of the algorithm.




Figure 7.5: Confusion matrix for GEO objects with ground-based observer. All 12 tracklets are correctly assigned for each of the objects




Figure 7.6: Overview of components in the LMB filter

The average runtime of this test case based on 25 samples was 20.1 seconds, where 80% of the time was spent processing the first 3 out of 12 tracklet groups. The reason is that most GM components resulting from the birth process are pruned out after several measurement updates. This fact can be appreciated from Figure 7.6, where the cardinalities of several key sets in the LMB filter are summarised. Since the initial state is empty (no prior knowledge), all tracklets from the first group result in a GM, introducing a significant number of Gaussian components. In this case, the total number of predicted GM components is 121 (not depicted in the figure). Since all targets have $P_D \approx 1$, the Kalman correction is computed for every combination of Gaussian component and tracklet (i.e. 968 updates). After the next measurement update, 80 of those remain (see Figure 7.6) and after half of the second night, each target has a single Gaussian left, allowing the final associations to complete in a fraction of a second.

7.1.3 Summary

This test case borrowed its setup from existing literature. It showed that the tracklet LMB filter can find and estimate the state of a small number of GEO objects based on sparse, ground-based measurements and correctly extract measurement associations from the sequential estimation process.

7.2 Closely spaced GEO satellites with space-based observer

The next test case considers a small number of GEO objects in close proximity to each other, observed by a space-based telescope. The case is used to investigate:

- (i) The performance of the proposed tracklet LMB filter as compared to the single-scan implementation, both for state estimation and measurement association.
- (ii) The features of the CAR IOD method with a space-based observer.

7.2.1 Setup

For this scenario, the object population consists of Arabsat 4B, 5B, 6B and 7B, ES'HAIL 1 and 2 and SKYNET 5B. These seven satellites have very similar, near-circular orbits at geosynchronous altitude with a longitude between 25° and 26.2° East in the ECEF reference frame. The two-line elements (TLEs) used as a reference for the simulation are listed in Table G.2 of the appendix.

The observer is a space-based telescope in a circular orbit at 500 km altitude, the viewing direction is constantly pointed towards zenith, and the FOV is 4 by 4 degrees. To reduce the data density,

the observer orbit has a 10-degree inclination. Although this type of orbit and pointing strategy is not very realistic for a space-based observer, it does ensure that most tracklets are about the same length and avoids telescope pointing logic that might introduce its own errors in the analysis.

If the observer is at 500 km with $i_{\text{obs}} = 0^\circ$ the centre of the FOV shifts by about $228.3^\circ/\text{h}$, while GEO satellites move at $15^\circ/\text{h}$. Given a $4^\circ \times 4^\circ$ FOV, the expected tracklet length is $\frac{4^\circ}{(228.3-15)^\circ/\text{h}}$ or slightly more than 1 minute. For increasing observer inclination, this length reduces.

Observations are simulated for a period of one week with an observation frequency of 1 Hz, resulting in the complete geometry shown in Figure 7.8. Note that tracklets are found when the observer crosses the equator and with declination δ roughly between -2° and 2° , which corresponds to the half-angle of the FOV. The variation in right ascension originates from the movement of the space-based observer and the observed satellites, while the FOV overlaps with the RSO orbits. Most objects are in the FOV for ~ 55 seconds (see Figure 7.7), but some are longer depending on the inclination of the observed object.




Figure 7.7: Distribution of tracklet duration




Figure 7.8: Observation geometry of GEO satellites observed from LEO with $i_{\text{obs}} = 10^\circ$ for 7 observation days. Every marker represents a complete tracklet

Notice from Figure 7.8 that SKYNET 5B (32294) is more often visible than the rest of the objects due to its inclination of 4.3° . The FOV crosses this satellite significantly more often than all the others. Therefore, this object is expected to be relatively easy to identify compared to the rest of the RSOs.

The tracking problem is solved using both the classical LMB filter (Section 3.3.4) and the proposed variation with tracklets as measurements (Chapter 6). To ensure a fair and insightful comparison between the filters, most of the common parameters are kept constant and are as listed in Table 7.2. Notice that the survival probability is set to 1.0. This is, in particular, necessary for the single-measurement filter because if an object is not observed during a specific pass, this corresponds to ~ 60 propagation steps and thus an effective survival probability of $(P_S)^{60}$. This unfairly affects objects that are less often seen and results in premature target death. As mentioned before, target death is not expected in the case of SSA and merely serves to stabilise the filter if the birth process produces too many targets. This is addressed in more detail for the current case in Section 7.2.4.

Table 7.2: Common parameters for the case with closely spaced GEO objects

Parameter	Value	Unit
Measurement noise $\sigma_\alpha = \sigma_\delta$	5	"
GM pruning ϑ_P^{GM}	10^{-5}	
GM cap $n_{\text{max}}^{\text{GM}}$	10	
LMB pruning ϑ_P^{LMB}	10^{-5}	
GLMB prior hypotheses cap $ \pi_{k k-1}^{\text{GLMB}} _{\text{max}}$	1000	
GLMB posterior hypotheses cap $ \pi_{k k}^{\text{GLMB}} _{\text{max}}$	1000	
Survival P_S	1.0	
Gating P_G	0.99999	
Clutter rate λ_c	10^{-4}	
Clutter intensity κ_c	$\frac{\lambda_c}{\int_{\text{FOV}} d\mathbf{y}} = \frac{\lambda_c}{16}$	deg^{-2}
Process noise \mathbf{Q}_a (Section 5.1.4)	$10^{-16} \mathbf{I}_{3 \times 3}$	m/s^2
Dynamics	Keplerian	
State representation	Cartesian	

7.2.2 Single-measurement LMB filter with prior knowledge

The first approach to address this tracking and association scenario is to use the classical LMB filter, where the update steps consider a set Z_k of individual measurements at a particular time t_k . Here, the MO state is initialised based on prior knowledge of the objects, 4 hours before the first measurements. An offset is applied to the ground truth as sampled from an uncorrelated Gaussian distribution where the initial position and velocity errors follow $\mathcal{N}(\mathbf{0}, (10 \text{ km})^2 \mathbf{I}_{3 \times 3})$ and $\mathcal{N}(\mathbf{0}, (10 \text{ m/s})^2 \mathbf{I}_{3 \times 3})$, respectively.

Since the objects are in GEO, 4 hours is much less than 1 orbital period, and the predicted uncertainty can be obtained using the Clohessy-Wiltshire equations [22][26]. The resulting RTN position covariance is shown in Figure 7.9, and its trace is $(259 \text{ km})^2$. Note that this rapid uncertainty growth is mostly due to the initial velocity error. For comparison, Figure 7.10 shows the inter-satellite distance between the tracked objects at the time of the first measurements. Apart from SKYNET-5B (32294), most satellites are at distances similar to the prior position uncertainty.




Figure 7.9: Predicted state uncertainty at time of first measurements using Clohessy-Wiltshire equations




Figure 7.10: Euclidean distance between the simulated objects just after the first tracklet group

As shown in Figure 7.11, the filter converges after about 60 hours of observation time, which corresponds to about 4 tracklets per satellite (except SKYNET 5B, which is observed 6 times by then). Because passes occur where not all objects have measurements, the errors of the unobserved objects keep increasing and the OSPA distance can therefore also increase between different estimated steps.

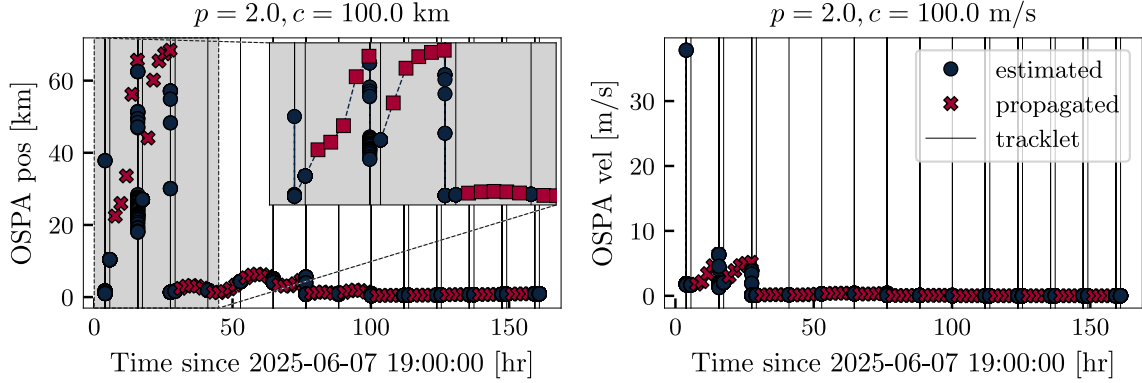


Figure 7.11: OSPA for classical LMB filter with 7 closely spaced GEO objects

The association results are first extracted based on the maximum of the individual (time-marginal) association probabilities from Eq. (4.2) in Section 4.1, and the resulting confusion matrix is shown in Figure 7.12. No false associations are found, but 225 out of 5014 measurements (4.5 %) are incorrectly categorised as clutter. Without changing anything in the filtering approach (i.e. still using the classical, single-measurement LMB filter), one can refine these association results considering that the measurements belong to tracklets and are therefore labelled. Even if the filter does not use this information, it is known which observations should belong to the same object and therefore, the *tracklet association* probabilities can be computed by accumulating the results as discussed in Section 4.2 by Eq. (4.7). The result is shown in Figure 7.13.

In other words, the difference between the two association results is that Figure 7.12 has independent assignment decisions for every single observation, whereas in Figure 7.13, all measurements from an entire tracklet are assigned as one to an object or to the clutter/unknown category.

For 4 of the 7 satellites, this accumulation step manages to pick up the previously rejected measurements so that all tracklets are correctly linked. The three others are each left with one tracklet that could not be associated. These are left in the last column of Figure 7.13.




Figure 7.12: Confusion matrix for classical LMB filter with time-marginal probabilities




Figure 7.13: Confusion matrix for classical LMB filter with accumulated probabilities

Remembering that the tracklets consist on average of 55 measurements (Figure 7.7), the total number of missed observations went from 225 in Figure 7.12 to ~ 165 (or 3.3 %) in Figure 7.13. It is essential to mention that the missed measurements using the marginal association probability in Figure 7.12 do not all originate from the same tracklet. They were individual measurements from various tracklets. In case a tracklet has relatively few incorrect associations, the accumulation step incorporates them based on all the correct conclusions from other measurements with the same label. However, this effect also works the other way. If many individual measurements in a tracklet are incorrectly rejected, the associated measurements are also considered clutter. So even though the total number of correctly assigned measurements is larger with the accumulated result from Figure 7.13, some observations that were correctly assigned in Figure 7.12 are now marked as clutter/unknown. This is a limitation of the method because even if the total number of associated measurements is higher, they might come from fewer tracklets. In a sparse data environment, rejecting all measurements from a tracklet might not be desirable.

To further investigate why some tracklets correctly pick up missed measurements and others are wrongly assigned to clutter, the existence and detection probabilities over time are shown for Arabsat 6B (41029) and 7B (56757) in Figure 7.14 and Figure 7.15, respectively. These plots indicate the target's existence probability before every measurement update with a triangle and use crosses to represent the detection probability. In addition, all tracklets for the relevant object are indicated, along with the duration of the tracklet group of which they are a part. Notice that tracklet groups are not used in the single-measurement LMB filter, but they show that the tracklets overlap extensively. First of all, notice that both targets start with an existence probability of 5, but are immediately confirmed by the first measurements, reaching $r^{(\ell)} \approx 1$ and remaining there indefinitely, as expected. At the times far away from the tracklets, the detection probability is zero, which corresponds to the FOV-based detection model from Section 5.2.2.

Focusing on the zoomed area in Figure 7.14, notice that the detection probability drops to ~ 0 before the end of the tracklet, making it very unlikely that the filter would associate the corresponding measurements correctly. This is an undesirable effect of the sharp edges of the FOV. However, because the majority of the measurements in this 41029 tracklet are still correctly associated, the missed observations are incorporated by the accumulation step and the complete

tracklet is assigned to the appropriate label (recall that there are no clutter assignments for 41029 in Figure 7.46). In contrast, if the detection probability is underestimated *most of the time*, the effects also build up and push the filter towards concluding that the target should have been undetected. In that case, the target could not have caused the tracklet, and the association is missed. Figure 7.15, the plot for Arabsat 7B, is a fitting example. Here, the modelled detection probability is only close to 1 for a small fraction of the tracklet (the predicted orbit may pass through a corner of the FOV instead of through the middle), and hence the tracklet is rejected in Figure 7.13.

Notice that the mismodelled detection probability mostly occurs after the object has not been detected for a while (several days in the examples of Figure 7.14 and Figure 7.15). After this time, the state error has increased enough to be outside the FOV most of the time.

Finally, it should be mentioned that the accumulation step for the association probabilities (recall Section 4.2) occurs after the filter update and therefore does not influence the convergence of the state estimates. It only serves to improve the extraction of measurement and tracklet associations for later evaluation.




Figure 7.14: Detection probabilities for Arabsat 6B. Here, the majority of the tracklet falls inside the predicted FOV, and the tracklet is correctly associated.




Figure 7.15: Detection probabilities for Arabsat 7B. In this case, the predicted FOV is so far off that the accumulated association probability decides in favour of clutter.

7.2.3 Tracklet LMB filter with prior knowledge

If the same problem is to be addressed by the new tracklet LMB filter ([Chapter 6](#)), several additional settings must be provided to the filter. Since the observer is space-based with a period of about 90 min and all objects are in GEO, the tracklet grouping can be performed with a validity padding of $\Delta t_{\text{valid}} = 1$ h. Because all objects are very close to each other in this case, the grouping algorithm has no difficulty with identifying tracklets that belong to the same filter update. Furthermore, the birth model is disabled here since it is not expected to discover new objects anyway. Finally, the same errors as before are introduced in the initial state.

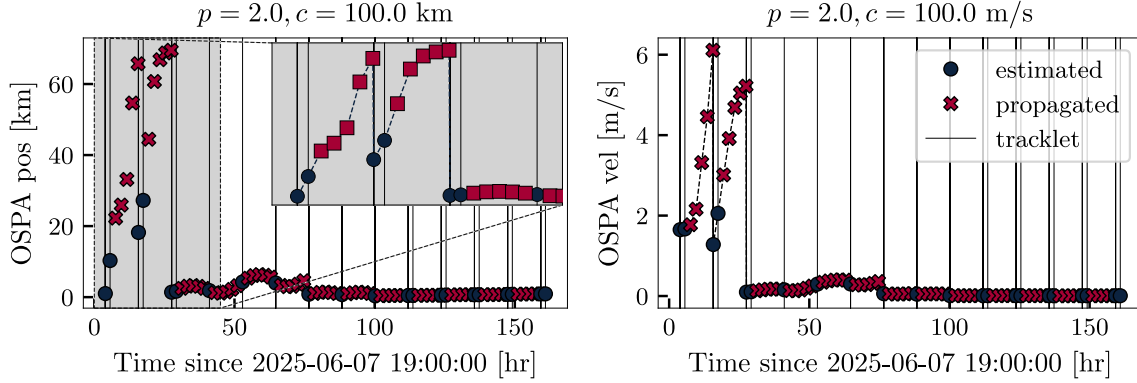


Figure 7.16: OSPA for tracklet LMB filter with 7 closely spaced GEO objects

The tracklet filter converges as quickly as the classical LMB filter did ([Figure 7.16](#) shows the OSPA errors for comparison, but a similar discussion applies as for the SO result from [Figure 7.11](#)) and the association conclusion is shown in [Figure 7.17](#). This method only leaves a single tracklet unassociated while all the others previously marked clutter are resolved. Again, the reason for misassociation is the detection model. The detection probability for ES'HAIL 2 (43700) is predicted close to zero for the first tracklet after a measurement gap of almost 40 hours. Therefore, the tracklet was not associated with the satellite but with a clutter process or some object unknown to the filter.




Figure 7.17: Confusion matrix for closely spaced GEO objects with tracklet LMB filter

For this case, the estimation errors over time are shown in Figure 7.18. Notice that the errors first increase because of the initial velocity error, which takes a while to correct with optical observations. Then, there is a first significant reduction in the error about 20 hours after the reference epoch and then the median error again increases for two days. Recall from Figure 7.14 and Figure 7.15 that most objects were not observable during that time and therefore did not have a possibility to obtain an improved state estimate. For insight in the accuracy of the estimated covariance bounds, Figure 7.19 visualises the Mahalanobis distance for all objects. The fact that all errors are far below the theoretical bound of 90 % confidence, demonstrates that no filter saturation occurred. The applied process noise could arguably have been even slightly lower, but this is not considered a problem as it does not interfere with the convergence of the filter and also avoids filter saturation.




Figure 7.18: Post-fit estimation errors using the tracklet LMB filter.




Figure 7.19: Mahalanobis distance between truth and estimated results using the tracklet LMB filter.

Although no extensive runtime comparison was performed, it is worth mentioning that the tracklet filter completes this problem in 11 seconds on average, whereas the single-measurement filter generally takes more than 4 times as much for the current scenario. It is important to mention that many parameters influence the runtime, particularly those related to hypothesis pruning and GM reduction. Furthermore, the duration of the tracklets heavily influences the complexity of the tracklet update step. This impact is likely further increased if the underlying propagation method is more complex than Keplerian or TLE dynamics.

7.2.4 Tracklet LMB filter without prior knowledge

To evaluate the working of CAR IOD with a space-based observer, the same case is addressed but without prior knowledge. The goal is for the tracklet LMB filter to discover all targets and, as before, to extract the association results. The adaptive birth process models existence probability based on the expected birth-to-clutter ratio in Eq. (3.37) and uses the single-tracklet CAR approach from Section 5.3.2. The relevant parameters, design uncertainties and CAR constraints are mentioned in Table 7.3. Additionally, the survival model is also conditioned on similar but less stringent constraints to ensure the results of bad associations get pruned out, just like in the first test case (see Section 7.1.1).

Table 7.3: Admissible region settings for closely spaced GEO objects. Here, $a_{\text{GEO}} \approx 42.2 \times 10^6$ m represents the theoretical geostationary semi-major axis for a point mass gravity model

Parameter	Birth value	Survival value	Unit
a_{\min}	0.99	0.95	a_{GEO}
a_{\max}	1.01	1.05	a_{GEO}
e_{\max}	0.03	0.1	—
Desired σ_ρ	10	—	km
Desired $\sigma_{\dot{\rho}}$	20	—	m/s
Birth-to-clutter $\frac{\lambda_{B,k}}{\lambda_{c,k}}$	0.5	—	—
Maximum existence r_B^{\max}	0.5	—	—

At first glance, the resulting association output (Figure 7.20) is not very successful. However, this confusion matrix merits closer evaluation as it provides insight into when the IOD process is successful and what it takes for a target to get confirmed in the filter. The complete matrix of association probabilities from which this figure is derived can be found in Figure F.2 of the appendix. It provides a more detailed log of which tracklets are misassociated and when each target is introduced.

In Figure 7.20, every column represents a separate target label in the filter, *born* based on distinct tracklets, but only the object ID found in that tracklet label is shown here. This association result has no column to represent clutter/unknown because those associations give rise to the birth components and are thus associated accordingly.




Figure 7.20: Confusion matrix for track let LMB filter without prior knowledge.




Figure 7.21: Confusion matrix for track let LMB filter without prior knowledge, rearranged by final association and dropping labels that are not part of the final estimated state

Notice that there are multiple columns in [Figure 7.20](#) with the same NORAD ID (*parent object*), which means that the birth process sometimes created various targets over time, competing to represent the same object. In each case where this occurs, however, a clearly dominating target collected most of the subsequent tracklets. For the satellites SKYNET 5B (32294), Arabsat 5B (36592), and ES'HAIL 2 (43700), a single target was created based on the first available tracklet and all other tracklets were correctly associated with it. For the other four satellites, the firstborn target was unsuccessful and resulted in a confusing association step, after which most of these initial hypotheses are pruned out and the second target converges. The one exception is Arabsat 7B (56757), which doesn't seem to have a target with many measurement associations. This is because all but the first of its tracklets were associated with the firstborn label of SKYNET 5B (39233). Thus, two targets started from a 39233 tracklet but converged on different satellites over time. So the filter was able to discover and converge on the states of all the individual satellites, but the close proximity of the RSOs, the initial uncertainty of the birth process and the sparsity of the tracklets resulted in *track switching* in the early stages of each target's existence. This can be visually confirmed by rearranging the columns of the confusion matrix and removing those labels with fewer than 3 associated tracklets. The result is [Figure 7.21](#) and it shows clearly that each object has a target, but the one for 56757 was born from a tracklet that actually belonged to 39233 and immediately switched.

This conclusion is also supported by the evolution of the OSPA metric in [Figure 7.22](#), which is agnostic to the correct labelling but confirms that the estimated number of states is accurate. Each RSO has a corresponding target with an absolute state error below 100 m and 0.1 m/s (Euclidean distance).

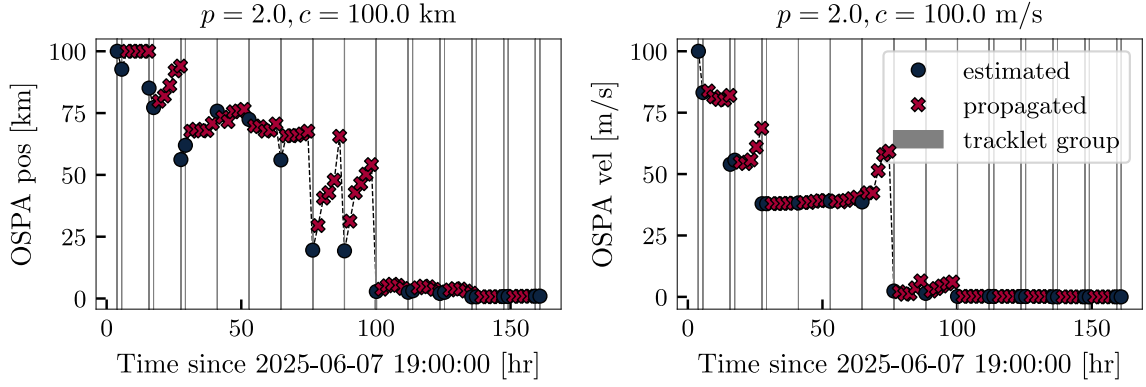


Figure 7.22: OSPA for tracklet LMB filter without any initial information. Reference epoch is 2025-06-07T19:00:00Z

Suppose one is interested in retrieving corrected associations for the early tracklets. In that case, it is possible to feed all tracklets to the filter in reverse order, since the scenario now corresponds to that where an initial state is known and a result similar to [Section 7.2.3](#) can be expected. Note that this introduces a bias because the same information is considered twice and a more rigorous smoothing approach could be desirable. Still, for the present case, this bias does not affect the results and MO smoothing in the tracklet LMB filter is left to be explored in future work.

To illustrate the process of discovering the targets, it is worth focusing on Arabsat 6B, for which 3 different labels were born. [Figure 7.23](#) shows the evolution of their existence and detection probabilities (before each measurement update), and all association probabilities can be consulted in [Figure F.2](#) for more context. The first tracklet from Arabsat 6B creates the target $\ell_1 = 41029-07T22:55$, which is undetected 90 minutes later, but has some components passing through the FOV after half a day (15.5 h after the reference epoch). Notice that $P_D \approx 0.5$ results from the birth uncertainties increasing over time – some of the GM components pass through the FOV whereas others do not and the resulting detection probability is a weighted product, accumulated




Figure 7.23: Tracklet detection probabilities of Arabsat 6B with tracklet LMB filter. The different colours represent different targets in the filter, born from the first three tracks. The triangles and crosses represent existence and detection probabilities, respectively.




Figure 7.24: Predicted densities for Arabsat 4B and 6B at the start of three different tracklet groups. At 15.5 hours after the reference epoch, both densities overlap very much and the number of GM components is insufficient to correctly capture the difference.

over the duration of the tracklet group (recall Eq. (6.3)). At this point, the filtering density of Arabsat 6B overlaps nearly completely with that of Arabsat 4B, and the filter makes the wrong association conclusion. The overlapping GM densities are depicted in Figure 7.24. At the same time, the correct tracklet creates a new label $\ell_2 = 41029-08T10:44$ for Arabsat 6B with very low existence probability. The wrong association of ℓ_1 results in several components that do not pass the survival constraints and ℓ_1 is pruned out, while $r^{(\ell_2)}$ increases enough for the target state to converge. A final label $\ell_3 = 41029-08T22:33$ is created from the third tracklet because $r^{(\ell_2)}$ was still very low, but it quickly loses credibility and later also disappears.

The above discussion captures part of what is going on in the tracklet filter, but ignores an important sensitivity in the filter parameters. As mentioned in the test case definition and Table 7.2, the same basic LMB filter settings were used for all solution methods and in particular, the settings dictate that the posterior density for each label must be a GM with at most $n_{\max}^{\text{GM}} = 10$ components. This makes sense if prior information is available on the object states but is incompatible with the CAR birth model used here. This becomes very clear when the birth density is plotted for a typical attributable in this scenario. Figure 7.25 shows the GM approximation for the CAR with design parameters from Table 7.3, which often has more than 100 components, and Figure 7.26 depicts the associated range-marginal PDF, which is mostly well-behaved and only shows some instability towards the concave edge of the region. Since all components have very similar weights, capping this mixture to only 10 components is detrimental and should be avoided.




Figure 7.25: CAR birth GM approximation with space-based observer.




Figure 7.26: CAR birth range-marginal PDF with space-based observer

Looking back at [Figure 7.24](#), it is also clear that the predicted density at $t = t_0 + 15.5$ h has only 10 components for each object, even though neither of the birth densities was ever updated by a tracklet. This prediction cannot fairly represent the mixture from [Figure 7.25](#).

Indeed, if the GM capping threshold is removed, the track switching no longer occurs and the tracklet filter manages to obtain a near-perfect association of all tracklets, without any prior knowledge on the number of objects and their states. The full association result is shown in [Figure F.3](#) of the appendix, which also makes clear that the only wrong association has only a 41% association probability, indicating correctly that this result should be treated with care. The resulting confusion matrix is provided in [Figure 7.27](#) for comparison.

It is important to stress that the imperfect results in [Figure 7.20](#) are thus mostly caused by non-optimal tuning of the parameters (or even a single parameter) in the filter, rather than by a deficit in the filter itself.




Figure 7.27: Confusion matrix for tracklet LMB filter with more GM components. Refer to [Figure 7.20](#) for comparison.

7.2.5 Evaluating association performance

As an alternative to the confusion matrices, it is also useful to define some notion of true positive (TP), false positive (FP) and false negative (FN) associations to evaluate the quality of the

measurement association result as a whole. If those are defined, it is also possible to compute widely used metrics such as the *positive predictive value (PPV)* and *true positive rate (TPR)*, also referred to as *precision* and *recall*:

$$\text{precision} = \text{PPV} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (7.1a)$$

$$\text{recall} = \text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}. \quad (7.1b)$$

For the purpose of this analysis, the following procedure is used to count the number of correct and incorrect associations:

- For every true object, initialise $\text{TP} = \text{FP} = \text{FN} = 0$.
 - Loop over a sorted list of measurements that belong to this object (ground truth association)
 - If the current and previous measurement are assigned to the same target label: $\text{TP} + 1$
 - If the current measurement is associated to clutter: $\text{FN} + 1$
 - If the current measurement is assigned to a different target than the previous: $\text{FP} + 1$
- The total counts are found by the sum over all objects.
- Note that if the assignment deviates a single time and then immediately returns to the previous target label, the above rules would count 2 FPs. To avoid such double counting of outliers, a single FP is considered in such a scenario. For example, if the measurements for object A are assigned chronologically to $\ell_1, \ell_1, \ell_1, \ell_2, \ell_1, \ell_1$

It is important to mention that *precision* and *recall* are often used as properties of a method or algorithm. However, the results from individual scenarios do not necessarily apply in other situations. *Instead*, this thesis uses the concepts to compare the accuracy of individual association outputs – these should not be read as claims on the the overall performance.

By the above rules, the association results from the discussed case compared for all of the considered approaches and the result is summarised in [Table 7.4](#).

Table 7.4: Overview of association performance for GEO objects from LEO observer

Initial state	Birth	Meas. type	TP	FP	FN	PPV	TPR	Confusion matrix
Yes	None	(α, δ)	4789	0	225	1.00	0.96	Figure 7.12
Yes	None	Labelled (α, δ)	87	0	3	1.00	0.97	Figure 7.13
Yes	None	(α, δ) tracklet	89	0	1	1.00	0.99	Figure 7.17
No	CAR	(α, δ) tracklet	89	1	0	0.99	1.00	Figure 7.20

7.2.6 Summary

This test case showed that the LMB filter in both presented forms is able to converge on the states of closely spaced objects in GEO and that the measurement association probabilities can be extracted from the GLMB update as proposed in [Chapter 4](#). It was found that the extraction of measurement associations is sensitive to *correct modelling of the detection probability*, definitely in scenarios where $P_D \lesssim 1$ inside a small field of view and nearly 0 otherwise. It was shown that using the accumulated association probabilities can alleviate this issue, but only if a significant portion of the predicted trajectory passes through the field of view. The tracklet LMB filter introduced in [Chapter 6](#) of this work does not have the P_D sensitivity, as long as the predicted trajectory passes through the FOV during at least one of the observation epochs. Further, it was found that the CAR birth model can be suitable for a space-based telescope observing GEO

objects. However, the ability of the LMB filter to confirm the components can be heavily affected by incompatible filter settings, in particular related to overly aggressive component reduction. Such inconsistencies are hard to detect without good knowledge of the algorithms.

7.3 Discovering GEO population from LEO

The following case aims to test the tracklet LMB filter for a *larger number* of GEO objects with measurements from a space-based observer in a *polar* LEO orbit. The purpose of this scenario is twofold:

- (i) to analyse how the filter handles a larger number of targets, and
- (ii) to evaluate the applicability of CAR IOD for a more complex observation geometry.

7.3.1 Setup

For this scenario, the RSO population is a set of objects in GEO, where the ground truth for the simulation is the Space-Track²⁵ TLE catalogue on 23 May 2025. Since the current implementation is not focused on performance and memory optimisation, the case considers a sample of 200 objects to simulate measurements from²⁶. The resulting distribution of semi-major axis and eccentricity is shown in Figure 7.28 and the ECEF longitudes of the sample are reflected in Figure 7.29.




Figure 7.28: Distribution of a and e 200 for considered GEO objects




Figure 7.29: Longitude distribution for considered GEO objects

This time, the observer is in a circular, near-polar orbit with $a_{\text{obs}} = R_E + 500 \text{ km}$, $\Omega_{\text{obs}} = 90^\circ$ and $i_{\text{obs}} = 97^\circ$. The attitude is defined such that observer line-of-sight $\hat{\rho}_{\text{LOS}}$ intersects with the equatorial GEO orbit at vernal equinox (i.e. $\langle (a_{\text{GEO}}, 0, 0)^T - \mathbf{r}_{\text{obs}}, \hat{\rho}_{\text{LOS}} \rangle = 0$) and the image horizontal is parallel to the equatorial plane. Again, the FOV is square-shaped with a half-angle of 2° . A visualisation of this pointing method is shown in Figure 7.30. This observation strategy is expected to be effective in discovering new GEO objects with either low inclination or with right ascension of the ascending node (RAAN) close to $\Omega = 0$ or $\Omega = 180^\circ$.

²⁵<https://www.space-track.org/>

²⁶The TLE catalogue is first filtered on orbital period between 1435.9 and 1436.2 min. After that, 200 objects were randomly selected.




Figure 7.30: Observation strategy to discover GEO objects with polar observer. The line-of-sight vector $\hat{\rho}_{\text{LOS}}$ is equivalent to the x axis of the observer body frame and aligned with the centre of the FOV.

The filter uses the tracklet LMB method with similar parameters as in the previous test case (recall [Table 7.2](#)), but several adjustments are made to account for the larger number of expected targets and listed in [Table 7.5](#).

Table 7.5: Common parameters for case with 200 GEO objects. Settings not mentioned here are the same as in [Table 7.2](#).

Parameter	Value	Unit
GM cap $n_{\text{max}}^{\text{GM}}$	200	
GLMB prior hypotheses cap $\left \hat{\pi}_{k k-1}^{\text{GLMB}} \right _{\text{max}}$	2000	
GLMB posterior hypotheses cap $\left \hat{\pi}_{k k}^{\text{GLMB}} \right _{\text{max}}$	10000	
Tracklet validity padding Δt_{valid}	1.0	h

7.3.2 Birth model

This test case clearly illustrates why the birth model must be tailored to the specific case. This is not only limited to ground-based vs. space-based and target regime variations, but it can also be highly dependent on the observation strategy. In the previous case, the tracklets were in the order of a minute, and the attributable approximation used for CAR IOD was then completely justifiable. However, the situation is very different with the observation strategy from [Figure 7.30](#).

Assuming for simplicity that the FOV is constant, then a GEO satellite with $i = 0^\circ$ will produce a tracklet of roughly $\frac{4}{360} \cdot 24 \text{ h} \approx 16 \text{ min}$. During this time, the observer traverses nearly 20% of its orbit.

The tracklets are thus not expected to be linear in the (α, δ) space, and that is confirmed by [Figure 7.31](#), which shows a small number of tracklets as they are provided to the MO filter. The paths are curved, as dictated by the fast movement of the observer. Trying to fit an attributable vector through such a tracklet might result in reasonable angular rates because of symmetry




Figure 7.31: Shape of GEO tracklets with polar space-based observer. Clearly, these tracklets are not suitable for linear approximation by an attributable.

properties, but the estimated right ascension and declination are almost always too far off to result in a representative admissible region. Additionally, since there are now more measurements per tracklet and the variation over time far exceeds the measurement noise, this observation geometry lends itself well to a Batch LS IOD. To improve the chance of convergence, the initial guess can be obtained by Gooding’s algorithm (recall [Section 5.3.1](#)), where the first slant ranges are set to a_{GEO} . One potential disadvantage is that this method provides only a single Gaussian component; this might be insufficient if the uncertainty is large and must be propagated for a long time. For the present example, it is not a significant problem since most objects will be observed about once per revolution (1 day), as they pass through the nearly fixed FOV depicted in [Figure 7.30](#).

7.3.3 Filter results

Using the described observation strategy for a total duration of 2 weeks results in 2486 tracklets spanning 2-17 minutes each and originating from 171 different objects – the other 29 are unobservable due to their inclination. The grouping algorithm divides these tracklets into 280 groups with an average duration of 53 minutes, and none of these groups contains more than one tracklet from the same object. This confirms that the grouping method can handle larger target populations correctly.

Over the two observation weeks, the filter extracts state estimates based on tracklets from 170 different objects. Some of the targets do not get confirmed and the final MO state contains 147 converged estimates (or 86% of the observable population).

[Figure 7.32](#) shows that most objects are discovered within the first 100 hours (75 tracklet groups), during which time the birth process produces on average 5.2 new target labels per group. After that, the number of birth components keeps fluctuating between 0 and 5 per filter step, but the resulting target labels mostly have a low existence probability and are pruned out shortly after. This is also clear from the figure, which shows that the corrected LMB $\hat{\pi}_{k|k}^{\text{LMB}}$ usually contains about 10 target labels more than the estimated MO state \hat{X}_k . The total number of GM components (denoted $\sum_{\ell} |\mathcal{G}_{k|k}^{(\ell)}|$ in the plot) is barely larger than the number of target labels since the batch LS birth process introduces a single Gaussian component per object and the mixture only grows due to association ambiguity.




Figure 7.32: Component counts in the tracklet LMB filter, observing 171 objects in GEO. The batch LS birth process produces a single GM component per target, so the total number of GM components only grows due to association ambiguity. The number of LMB components is higher than the number of estimates, as birth components are created with $r^{(\ell)} < r_B^{\max}$ and need time to be confirmed or rejected.

In Figure 7.33, it is observed that the post-fit OSPA errors level off at 54 km and 52 m/s for position and velocity, respectively (with $p = 2$ and $c = 100$ in both cases). Correcting for the 29 unobservable objects that all contribute an unavoidable cardinality penalty, the OSPA distance compared to the observable population is 41 km and 39 m/s. To better understand the errors for those objects that the filter has discovered, the OSPA can also be computed without any cardinality penalty. Since the second-order ($p = 2$) distance is used here, this is effectively




Figure 7.33: OSPA errors over time as compared to various subsets of the ground truth population. The blue curve uses all 200 objects, the red curve reduces the cardinality penalty for unobservable objects and the yellow one considers the final estimated cardinality as the truth (With $p = 2$, this corresponds to the RMS error for the optimal assignment).

equivalent to the RMS error between the estimates and the best-matched object of the ground truth. All these results are summarised in [Table 7.6](#).

Table 7.6: OSPA distances as computed from [Eq. \(3.38\)](#) using different penalties for cardinality. Y indicates the ground truth MO state, $p = 2$, $c^{\text{pos}} = 100 \text{ km}$ and $c^{\text{vel}} = 100 \text{ m/s}$.

	All objects	Observable	Observed	Unit
$ Y $	200	171	147	–
Position $d_{p,c}^{\text{OSPA}}(Y, \hat{X})$	53.9	41.3	18.6	km
Velocity $d_{p,c}^{\text{OSPA}}(Y, \hat{X})$	52.2	38.6	10.1	m/s

In what follows, the *parent object* for a specific target label refers to the source of the tracklet used in the birth process to introduce this target. Just like before, it is found that not all targets converge on their respective parent object, and some quickly switch to another RSO after an initial association confusion. This is the same type of early-stage track switching that was also observed in the previous test case. The switching behaviour doesn't make the final state estimates less valid but prevents directly computing the errors with respect to the ground truth²⁷.

One way to evaluate how many targets converged on a different object is by inspecting the Mahalanobis distance. If this is high, it means that the error between the target and its parent object is far larger than the estimated covariance bound. That might suggest that the filter has consistently found measurements to support the target's existence and narrow down its state, but these measurements do not belong to the same object as the original birth tracklet. The middle plot of [Figure 7.34](#) shows that the errors for more than 90% of the final estimates are smaller than the 90% confidence interval from their estimated covariance (i.e $d_M < 3.26$). The median distance is within $d_M < 2$, proving that the errors are usually well-bounded by the filter uncertainty. 10 targets have $d_M > 5$ (note from [Section E.1](#) that this is outside the 99.97% confidence limit), so they must either have switched to a different object or diverged due to filter saturation. Because of the previous observation that the OSPA distance without cardinality penalty is below 20 km (recall [Figure 7.33](#)), it can be assumed that no filter saturation has occurred for all the targets in the estimated MO state.

For all targets that converged on their parent object ($d_M < 5$), the evolution of the error is represented in the lower plot of [Figure 7.34](#), showing that half of the targets have final estimation errors below 1.9 km and the rest are within 11 km of the truth. Since this represents an angular offset of about $58''$, it is safe to conclude that all targets indeed converged on separate objects. Notice that the errors in that figure are oscillating with a period of 1 day. This is most likely due to the fact that the objects are not uniformly distributed over the GEO orbit. [Figure 7.29](#) shows that the sample has far fewer objects in the Eastern hemisphere than in the Western hemisphere and the GEO regime is generally less populated above the Pacific ocean. During times when no objects are observed, the median error increases, whereas dense observation scans cause the median error to reduce – hence the observed trend.

A summary of all the object discovery results is shown in [Table 7.7](#). Finally, to better understand why 21 of the observable objects were not confirmed by the LMB filter, the actual and estimated distribution of RAAN and inclination are plotted in [Figure 7.35](#) and [Figure 7.36](#), along with

²⁷In principle, the most reliable way to obtain the errors would be to extract the sub-pattern assignment from the OSPA distance computation, but this information is not retained in the current implementation.

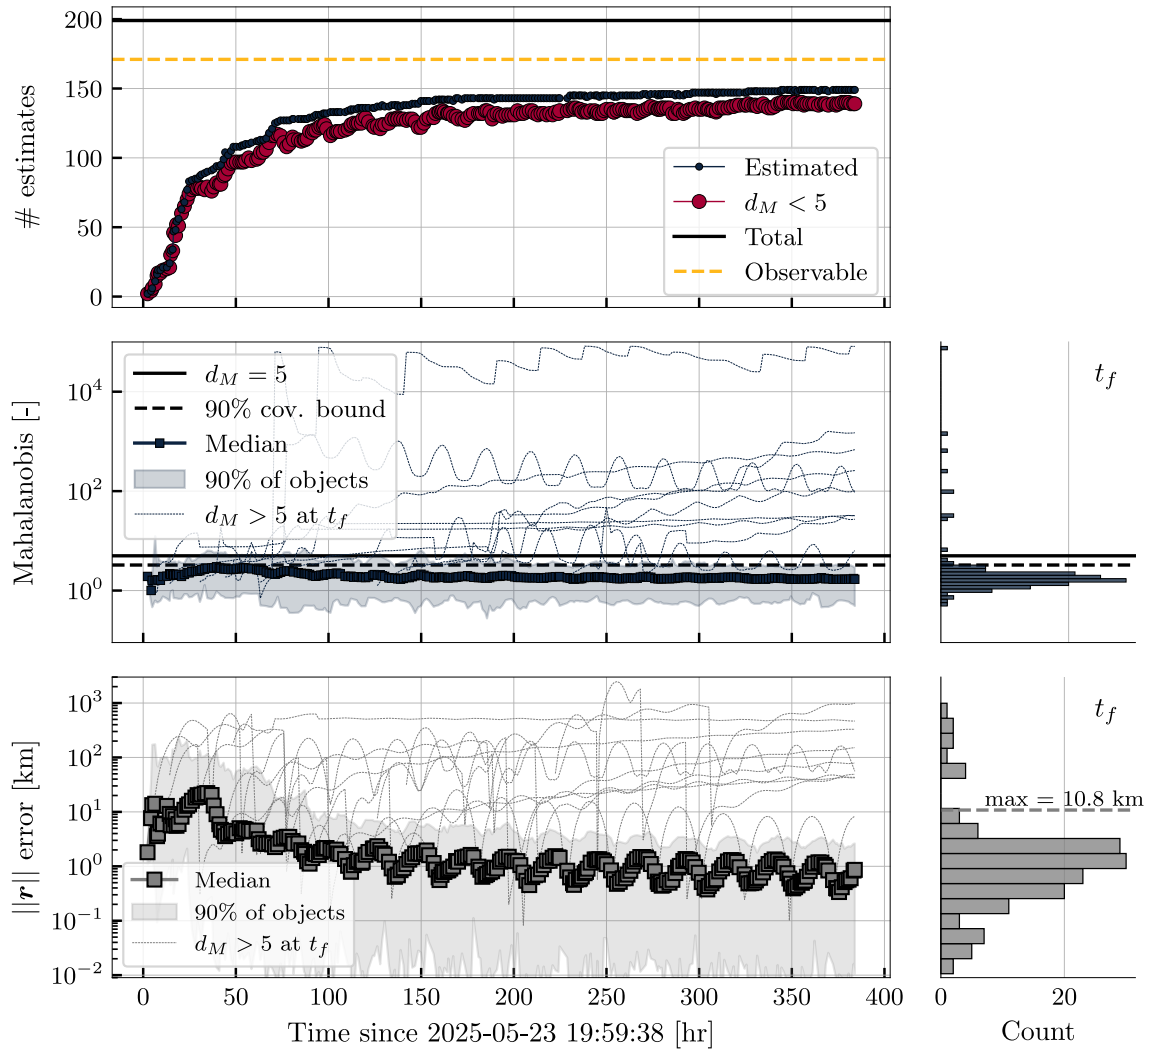


Figure 7.34: Evolution of estimates over time. *Top:* cardinality, both including and excluding targets that that did not converge on their parent object. *Middle:* Mahalanobis distance median with 0-90% percentile interval. Objects for which $d_M > 5$ at the end are shown individually. *Bottom:* Estimation errors.

the theoretical observability limit. These plots seem to indicate that the chosen set of objects is not uniform in RAAN and particularly dense at $\Omega \approx 90^\circ$. There does not seem to be a specific trend in hard-to-find GEO objects, and since the cardinality estimate is still increasing and the birth process still generating components at the end of the simulation (recall Figure 7.32), it is possible that the remaining objects would also be discovered given more time.

Table 7.7: GEO object discovery results from polar LEO observer

Total objects	# observable	# found	# confirmed	# converged on parent object
200	171	170	147	137




Figure 7.35: True distribution of inclination and RAAN for 200 GEO objects. Objects are marked red if estimated at the final filter epoch, blue if estimated at least once and yellow if never found at all.




Figure 7.36: Estimated distribution of inclination and RAAN at the final filter epoch.

The final estimation errors are shown in the RTN local orbit frame in Figure 7.37. Visual inspection suggests that these errors are nearly Gaussian and unbiased (centred around 0 in all coordinates). The position error is dominated by the along-track component (T) whereas the velocity error is mostly in the radial direction (\dot{R}). This is consistent with expectations; uncertainties in semi-major axis cause phase drift which mostly affects the along-track component of the orbital state (recall the example from Figure 5.3) The radial velocity error results from its geometric coupling with the transverse position error.

Furthermore, notice that the RMS errors are 1.83 km and 0.17 m/s, respectively. These are the errors for all objects that converged on their parent object ($d_M < 5$). Recall that the RMS errors were computed before in Figure 7.33, by removing the cardinality penalty from the second-order OSPA distance. Those errors were more than an order of magnitude larger (see Figure 7.33). That means that one or more of the targets with $d_M > 5$ did not switch to another object, but rather diverged altogether without representing any real satellite.

Finally, the association performance metrics from Section 7.2.5 are listed in Table 7.8 for the current case. The precision of 0.99 shows that – in this setting – the adapted LMB filter very rarely misassigns tracklets once a target is confirmed, but if newborn targets do not converge fast enough (i.e. if it takes too long before they are again in the FOV), this results in clutter assignments (FNs), causing lower recall value of 0.91.

Table 7.8: Association quality for GEO discovery case

TP	FP	FN	PPV	TPR
1613	14	166	0.99	0.91

7.3.4 Summary

This test case aimed to discover a subset of the GEO population from space-based measurements. The chosen observation strategy resulted in long tracklets and it was shown that the CAR approach is unsuitable when the observer moves too much during a single pass. The Gooding + Batch LS approach worked better in this scenario and the tracklet LMB filter converged on

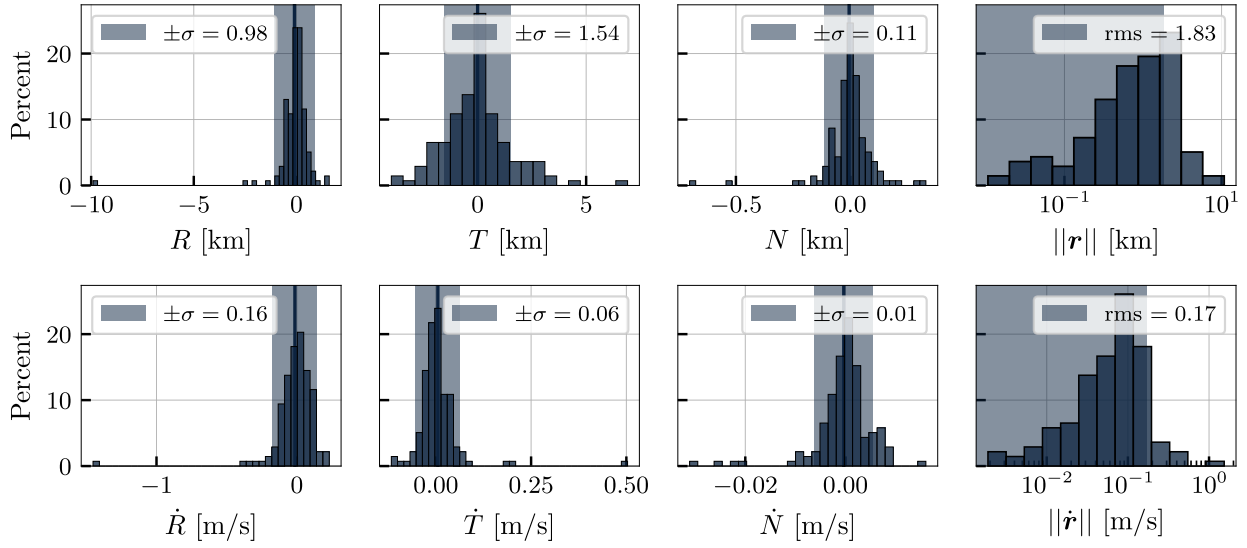


Figure 7.37: Final RTN error histograms for all objects that converged on the parent object.

reliable estimates for 86% of the observable RSOs, with all errors below 11 km and 1 m/s, but the majority almost an order of magnitude smaller. Finally, the tracklet grouping performed as expected, resulting in consistent and non-ambiguous measurement sets.

7.4 Polar payload separation

As a final test case, it is attempted to observe a payload separation event²⁸ in LEO and track the new RSOs using a space-based observer. The primary goal is to determine how all the previous findings in GEO translate to a different target orbit regime.

7.4.1 Setup

This case considers a rocket upper stage in a near-circular orbit at 800 km altitude with $i = 90^\circ$ and $\Omega = 25^\circ$. Its 9 rideshare payloads separate with a relative velocity between 0.5 and 3 m/s in the transverse (T) direction of the local orbit frame, each 1 minute after the previous one. The resulting orbital elements for the ground truth are listed in [Table G.3](#) of the appendix.

The observer is in the same near-polar orbit as before with $i = 97^\circ$ and $\Omega = 90^\circ$, but now employs a different observation strategy, depicted in [Figure 7.38](#). The telescope is pointed at the ECI point $(0, 0, R_E + 800 \text{ km})$ when it is above the equatorial plane and at $(0, 0, -R_E - 800 \text{ km})$ when it is below. The resulting tracklets ([Figure 7.39](#)) are very close to each other but are still relatively varied in right ascension and declination due to the space-based observer.

7.4.2 CAR birth model

First, the problem is attempted without any prior knowledge and using the CAR approach to discover the separate payloads. The constraints are set at $R_E + 700 \text{ km} < a < R_E + 800 \text{ km}$ and $e < 0.05$. Since both the target object and the observer are in LEO, however, the admissible region is less straightforward than before.

²⁸The case setup is inspired by this type of scenario, but it should be noted that conditions are not entirely realistic. Still, the setup sheds light on the kind of observations and filter behaviour that can be expected tracking closely spaced LEO objects from a space-based observer.

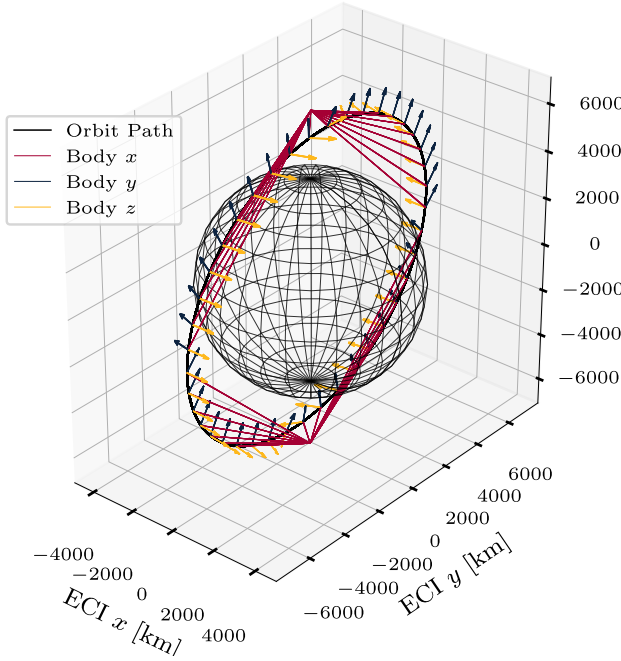


Figure 7.38: Observer attitude profile for polar payload separation. Note that the Earth is not drawn to scale (about 70%).

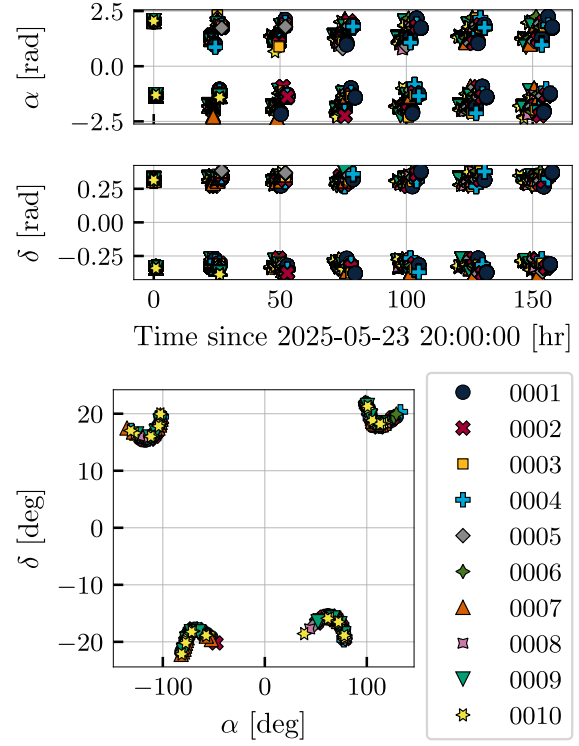


Figure 7.39: Tracklet distribution for polar targets from polar observer

Figure 7.40 depicts the mapping of semi-major axis and eccentricity in the $(\rho, \dot{\rho})$ space for an arbitrarily chosen but representative case where the observer state and attributable in ECI are

$$\mathbf{r}_{\text{obs}} = \begin{pmatrix} 815.3 \\ -1596.9 \\ 6640.3 \end{pmatrix} \text{ km}, \quad \mathbf{v}_{\text{obs}} = \begin{pmatrix} -215.4 \\ -7404.6 \\ -1754.3 \end{pmatrix} \text{ m/s}, \quad \text{and} \quad \mathbf{a} = \begin{pmatrix} 117.2^\circ \\ 18.4^\circ \\ 5.32^\circ/\text{min} \\ -3.26^\circ/\text{min} \end{pmatrix}. \quad (7.2)$$

Several key differences can be observed compared to the CAR behaviour for GEO objects. First, note that there are now two distinct regions where both constraints are satisfied. One of those contains the point $(0, 0)$, which corresponds to the observer orbit itself and, thus, the trivial solution. Other solutions in that region are targets that would have nearly the same trajectory as the observing satellite. Clearly, a discontinuous SO density is detrimental at the point of state estimation – if a mean estimator is used, a weighted sum of the two regions will most likely result in a value outside of both. For this reason, the region containing the observer is discarded for the present application; it is not expected to find an object there.

Furthermore, the second region is far more elongated than the CAR results previously found for GEO orbits. Recall Figure 7.32, where the maximum slant range in the region was less than 1% larger than the minimum and the allowable range rate varied by less than 250 m/s. For the current admissible region, these measures are nearly 50% and 1.8 km/s. The GM approximation of this CAR with design uncertainties of $\sigma_\rho = 30$ km and $\sigma_{\dot{\rho}} = 10$ m/s is shown in Figure 7.41, which highlights another shortcoming of the approach in this context: each of the Gaussian components assumes no correlation between ρ and $\dot{\rho}$, but it is apparent from visual inspection that the region




Figure 7.40: Constrained admissible region for polar observer and target

would be better represented if the covariance ellipses were at an angle. Incidentally, this would require fewer components for an adequate approximation.

In Figure 7.42, it is shown how this elongated CAR heavily affects the distribution of orbital elements. The trivial solution is found again in the upper right corner, overlapping with the observer. The other solutions are in a nearly uniform region that spans about 4° of inclination and nearly 15° in RAAN. This makes the IOD far more varied in geometry than it was for GEO orbits.

Not surprisingly, the tracklet LMB filter with CAR birth model does not produce satisfactory results for this scenario. The large uncertainties, combined with very closely spaced targets, result in low measurement likelihoods for each association hypothesis. In the beginning, most tracklet cause their own birth label and new components keep appearing until the objects are sufficiently




Figure 7.41: GM approximation of CAR in LEO, ignoring the *trivial solution*.




Figure 7.42: Distribution of RAAN and inclination resulting from a CAR in LEO.

separated. Only after 4 days (over 50 orbital revolutions) the final objects are confirmed, and the solution converges.




Figure 7.43: Association summary for polar targets from polar observer. Note tracklets marked “clutter” are at the start of the simulation.




Figure 7.44: Association summary, rearranged by final association. Unconfirmed targets are removed.

The confusion matrix in Figure 7.43, and the complete set of association probabilities in Appendix F show that early-stage track switching again occurs for three of the targets (columns 2, 8 and 10 in Figure 7.43 and Figure 7.44). After that, all associated tracklets (with three exceptions) belong to the same object. Note that the *clutter* column contains mostly the early tracklets and gives an indication of how long it took to find the corresponding object. Note that an extra target is present with 3 associated tracklets (column 4 in the above confusion matrices). These associations raise the existence probability enough to influence the cardinality estimate and the resulting orbit estimate never again passes through the field of view. Since $P_S \approx 1$, this *rogue target* is not easily removed, even though it is not supported by new tracklets. The final OSPA distance is summarised in Table 7.9.

Table 7.9: Final OSPA distances for payload separation with CAR birth. The actual OSPA includes the cardinality penalty, but removing it shows that all tracked objects are within 0.5 km of the truth.

	Actual	Ignoring rogue target	Unit
$ \hat{X} $	11	10	
d_r^{OSPA}	30.15	0.47	km
d_v^{OSPA}	30.17	1.22	m/s

Using the rules from Section 7.2.5 and the definitions from Eq. (7.1), the TP/FP/FN association counts are listed for all objects of the current case in Table 7.10, resulting in an overall precision of 89 % and recall of 63 %. This again confirms that the method has a tendency to overestimate the probability of clutter. It is further interesting to plot the cumulative evolution of the TP, FP and FN quantities per object as a function of the number of tracklets considered thus far. This is shown in Figure 7.45 and shows that the majority of clutter assignments (FNs) occur at the start of the simulation and the curve later flattens off. In contrast, there are essentially no TPs in the early filter steps, but once an object is confirmed, the correct assignments keep increasing.

Table 7.10: Association count for all objects. These correspond to the final values in Figure 7.45.

Object ID	TP	FP	FN	PPV	TPR
0001	13	1	16	0.93	0.45
0002	20	1	8	0.95	0.71
0003	18	2	10	0.90	0.64
0004	23	1	6	0.96	0.79
0005	20	3	6	0.87	0.77
0006	11	3	15	0.79	0.42
0007	17	2	11	0.89	0.61
0008	9	3	15	0.75	0.38
0009	21	1	4	0.95	0.84
0010	16	3	9	0.84	0.64
Total	168	20	100	0.89	0.63




Figure 7.45: Cumulative association counts per object over time. The x-axis indicates the number of tracklets that were already seen per object and the cumulative count of TP/FP/FN is shown on the y-axis.

Since this shows that it is possible for a tracklet to be associated with the birth density from a space-based CAR, it is worth revisiting why the early tracklets are more likely associated with the clutter process than with any of the existing targets. One probable hypothesis is that this is caused by the heavy simplifications in the tracklet clutter model. Recall from Section 6.1.3 that a heuristic is used to model the tracklet clutter intensity as the geometric mean of its measurement clutter intensities. This violates the PPP clutter assumption of the standard measurement model (Eq. (3.21)) since the clutter intensity will not integrate to λ_c over all possible tracklets but rather to some *larger* value. Therefore, the MO measurement likelihoods (Eq. (2.23)) are consistently *underestimated* and because the GLMB accounts for clutter implicitly (by the complement of other hypotheses), its likelihood is then *overestimated*. This effect only manifests if the measurement association likelihoods are low, which is the case for the large LEO CAR densities, but also in case the predicted target is outside the FOV and is therefore assigned a low detection probability.

To further support this hypothesis, it could be good to use the CAR in combination with the single-measurement LMB filter and test if the targets get confirmed earlier. The expectation is that this would indeed be the case, because clutter is modelled consistently and there is no coupling between clutter and birth. For adaptive tracklet birth using the traditional LMB filter, refer to Section 5.3.3.

7.4.3 Gaussian birth from rocket upper stage

Alternatively, it is possible to eliminate the birth model and initialise the filter at the time of the first tracklet based on prior tracking knowledge of the rocket upper stage. The initial LMB is then

$$\pi_{0|0}^{\text{LMB}} = \left\{ \left(r^{(\ell_i)}, \{ (1, \mathbf{x}_i, \mathbf{P}_0) \} \right) \right\}_{i=1}^{10}, \quad \text{with} \quad r^{(\ell_i)} = 0.8, \quad (7.3a)$$

$$\text{and } \mathbf{x}_i \stackrel{\text{sample}}{\leftarrow} \mathcal{N}(\mathbf{x}_0^{\text{rocket}}, \mathbf{P}_0), \quad \mathbf{P}_0 = \begin{pmatrix} (5 \text{ km})^2 \mathbf{I} & \mathbf{0} \\ \mathbf{0} & (5 \text{ m/s})^2 \mathbf{I} \end{pmatrix}. \quad (7.3b)$$

Keeping all other settings the same, the association results of the first 10 tracklets per object are shown in [Figure F.4](#) of the appendix, which indicates that the association probabilities for the first tracklet groups are very low and distributed over multiple targets. This is because of the large state uncertainty and close proximity of all the targets in the filter. As a result, there is some early track switching after which the following associations still have relatively low probability but consistently result in the correct assignments. These are summarised in [Figure 7.46](#) and rearranged for clarity in [Figure 7.47](#). It should be stressed that the birth model is disabled so the “target label” in these plots has no intrinsic meaning but was arbitrarily assigned by the filter while sampling the initial states. It is used here to ensure consistency with the detailed output in [Figure F.4](#).




Figure 7.46: Assignment summary for payload separation case with initial state from rocket body and no birth model

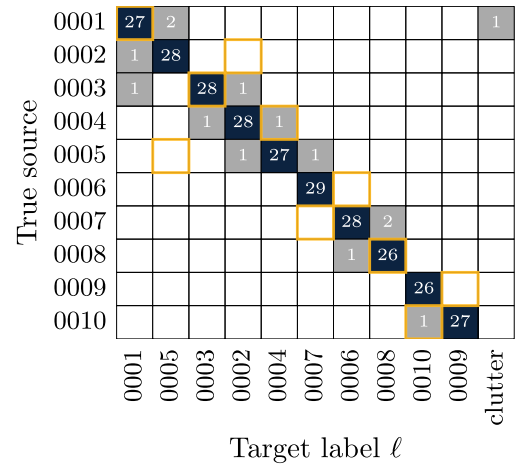



Figure 7.47: Assignment summary for payload separation case with initial state, reordered by final association.

Using the same approach as in the previous section, the association counts are listed in [Table 7.11](#) and the PPV and TPR for the tracklet LMB filter in this scenario become 0.95 and 0.96, respectively.

Table 7.11: Association counts summary for payload separation without birth model

TP	FP	FN	PPV	TPR
261	13	11	0.95	0.96

Furthermore, Figure 7.48 demonstrates that all FPs and FNs are within the first three tracklets per object and thus before the state estimates converged on particular objects. This suggests that one could implement a smoothing method to also refine the first associations. For this reason, this test case can be considered a very successful application of the tracklet LMB filter in LEO.

**Figure 7.48:** Cumulative association counts per object without birth model

Additional note

One of the previous cases (Section 7.2.4) already showed that GM capping can have undesirable effects on the filter performance. It was pointed out that the CAR birth model may produce many components of nearly equal weight, such that removing some of them results in a distorted distribution that reduces the model fidelity. It is good to briefly return to this issue for the current scenario, because it is relevant in a slightly different way.

Notice that the only clutter assignment in Figure 7.47 is for a tracklet from object 0001. In early results for this test case, all 0001 tracklets that came after this one were assigned to the clutter process and the corresponding target ($\ell = 0001$) was *pruned out* of the filter. Figure 7.49 visualises how the existence probability drops after the first measurement gap ($t + 20$ h) and the detection probability is already much lower. Consecutive assignments fail as well and at some point $r^{(0001)}$ drops below the LMB pruning threshold ϑ_P^{LMB} and it is removed from the MO state density.

Apart from all the advantages that the labelled RFS framework provides, the possibility that confirmed targets experience a sudden drop in existence probability or are pruned out entirely is logically inconsistent in a space object cataloguing context. This reflects a mismatch between the physical meaning of object survival (which is nearly certain in the SSA context) and the filtering mechanism that systematically needs to remove hypothesised objects that do not find confirmation in the observations (any more).

Nevertheless, in this specific case, the reason for the poor behaviour was once again the GM capping threshold of 10 components. Since there is no birth model here, the number of components grew because of all the uncertain associations in the early stages of the simulation. The SO density then has an increased number of components because of the various potential associations. Capping the GM resulted in the removal of crucial components and the filter does not recover. Once again, this is a very small but deeply nested configuration problem that is not straightforward to track down. In general, it should be recommended to only rely on GM capping if absolutely necessary for runtime reasons and to avoid LMB pruning in cases without birth model.




Figure 7.49: Detection and existence probabilities for target 0001 with overly aggressive GM capping and LMB pruning

7.4.4 Summary

This test case showed that the constrained admissible region approach yields valid birth densities for LEO targets with space-based observations, albeit with far more variation in orbit geometry than for GEO. Additionally, the GM approximation could benefit from components with non-zero $(\rho, \dot{\rho})$ correlation to better represent the elongated regions that the constraints often produce for a space-based LEO attributable. Using limited prior information, the tracklet LMB filter can extract the tracklet association for a polar payload separation event with a precision of 0.95 and recall of 0.96, where the remaining inaccuracies are mostly due to large uncertainties in the early stages of the simulation and may be resolved by a smoothing approach.

CONCLUSION

This thesis investigated how tracking methods based on labelled random finite sets (LRFs) can be employed to perform tracklet association and orbit determination for resident space objects (RSOs), particularly using space-based angles-only observations. The main contributions resulting from these efforts are first listed below and then framed in the original research questions in [Section 8.1](#). Possible directions for future work are discussed in [Section 8.2](#).

A. Development

The work detailed the development of BASIL, a flexible and extensible estimation library containing the necessary building blocks for RFS-based multi-object tracking and data association, with a focus on the (generalised) labelled multi-Bernoulli (LMB) density and corresponding LMB filter. The core algorithms can work for any multi-object estimation problem, but are implemented and tested in particular for the context of space situational awareness (SSA) through an interface with Orekit [\[85\]](#).

B. Tracklet LMB filter

Since the duration of a single pass through a telescope’s field of view (FOV) is mostly orders of magnitude shorter than one orbital period, the observations are often nearly linear in topocentric right ascension and declination. Therefore, successive measurements can already be combined into *tracklets* without resorting to orbital dynamics. Existing RFS methods generally discard this information and update the state estimates at every distinct observation epoch. This work proposed an adapted LMB filter that can consider groups of tracklets instead of single-epoch measurement scans, where groups are formed dynamically to ensure object uniqueness in every filter step. Computing measurement likelihood based on complete tracklets minimises the contribution of individual outliers and thereby reduces hypothesis explosion in the state density, leading to improved convergence properties. Without measurement-driven target birth, the filter performs as intended in low-clutter environments with LEO or GEO objects, using ground-based and space-based observers.

C. Extracting associations

The LMB filter is shown to allow extracting measurement associations from the update step, so they can be stored for further single-object analysis. It was found that measurements at the edge of the FOV are often missed using the original LMB filter, if the predicted state has not yet entered or already left that FOV. For labelled measurements (based on some external tracklet formation process), accumulating association probabilities helped recover some of these missed measurements, but only when a large portion of the tracklet was within the FOV. The proposed tracklet LMB filter addresses this issue by evaluating the detection probability over the entire tracklet group, significantly reducing sensitivity to sharp FOV boundaries and almost entirely mitigating the problem.

D. Insights for initial orbit determination

While not a core contribution, the work also examined suitable methods for the discovery or *birth* of new RSOs based on single uncorrelated tracklets (UCTs). For target objects in the geostationary (GEO) regime, the constrained admissible region (CAR) provided an excellent

tracklet-driven birth method if the observer was either ground-based or used a fixed pointing strategy, resulting in sufficiently short tracklets for a linear approximation of the angular rates. Observation modes where the observer’s position and attitude changed significantly within the span of a single pass resulted in curved tracklets and poor translation of the constraints to the initial density. For sufficiently long GEO tracklets, a combination of Gooding’s method and batch least squares estimation was shown to be a suitable alternative. Using this method and without prior knowledge, the tracklet LMB filter discovered 86% of a downsampled GEO population, with most objects confirmed after 3 revolutions or less.

Admissible regions for objects in LEO lead to distributions with more variation in orbital geometry and are less accurately represented by the Gaussian mixture approximation that worked well for objects in GEO. A payload separation test case in polar orbit showed that these widespread densities *can* converge based on follow-up measurements. However, if multiple very uncertain birth densities exist, clutter probabilities are implicitly inflated, and the association often fails. This highlights an intrinsic bias introduced by the tracklet clutter model, which should be addressed to make the approach mathematically rigorous.

8.1 Revisiting the research questions

(i) Which available MOT methods are suitable for tracking and data association with space-based optical measurements?

The LRFS framework was identified as an appropriate candidate for RSO tracking, especially through the LMB filter. Unlike its unlabelled counterpart, LRFS has the advantage that it inherently outputs *trajectory* estimates, rather than single-epoch states. Since maintaining object identity is crucial in space object cataloguing, a labelled approach is preferred. This avoids heuristic post-processing steps to associate state estimates with the appropriate catalogue entries. In its most rigorous form, LRFS tracking is based on the GLMB density, but the LMB filter offers an efficient, first-moment approximation that inherits most benefits. Although computationally expensive – scaling at most cubically with the number of measurements per scan – the LMB filter remains significantly more tractable than classical methods such as MHT and JPDA, which are typically NP-hard.

A test case with close-spaced objects in GEO demonstrated that observations from a space-based telescope can be used effectively in the LMB filter, with all estimates converging based on several minutes of observations spread over the first two observation days.

(ii) How can the optimal tracklet association with catalogued RSOs be extracted from MOT methods, independent of the state estimation?

While updating the multi-object state using a scan of measurements, the LMB filter considers a set of the most probable measurement-to-object association hypotheses. The joint likelihood of each hypothesis determines how much it contributes to the corrected state density. This process accounts for possible missed detections and clutter measurements that do not correspond to any real object.

In the LMB/GLMB filter, the associations exist as explicit mappings between the target labels (which may be linked to real objects in a catalogue) and individual measurements. By extracting these and normalising their joint likelihoods, one can construct an association probability distribution for each measurement over all catalogued objects. The mode of this distribution then

yields the final deterministic assignment, allowing the resulting single-object measurement sets to be catalogued and used for further analysis of individual space objects.

If observations are already pre-grouped in tracklets, the notion of *measurement association probability* can be extended to *tracklet association probability* without making any adaptations to the filter itself. A tracklet then simply dictates that all its measurements relate to the same (space) object. To this end, a method to *accumulate* the association probabilities was proposed. It was found that this reduces the negative impact of FOV-based detection models, and improves the overall assignment accuracy.

It is important to note that this approach is only as reliable as the internal representation within the GLMB filter itself. If the filter’s association hypotheses are poor—due to, for example, an inaccurate birth model or high state uncertainty—the extracted associations may also be unreliable. Thus, ensuring well-tuned filter models remains critical for robust tracklet-to-catalogue assignment.

(iii) How can the MOT methods be adapted to use all information of optical tracklets?

To fully incorporate pre-existing tracklets in the MO estimation process, the proposed tracklet LMB filter extends all the relevant building blocks of the original method so that it can interpret *tracklets as measurements*. The main adaptations can be summarised as:

- Instead of measurement scans, the new filter uses *tracklet groups*. A simple grouping method ensures that tracklets within the same group certainly originate from distinct objects.
- The single-object tracklet likelihood is computed as the geometric mean of the likelihoods of the individual measurements.
- The detection model is extended to consider all possible observation times within the duration of the tracklet group.
- The concept of a *clutter tracklet* is ill-defined, since a tracklet might be “corrupted” in various ways. This work assumes that tracklet formation significantly reduces the already low clutter rate in optical observations and uses a preliminary heuristic model to represent tracklets not belonging to any existing target.
- Finally, this version of the LMB filter is directly compatible with single-tracklet birth models, facilitating the discovery of new objects.

Various simulation cases have confirmed the applicability of this method for SSA tasks. The filter reduces the sensitivity of the detection probability to the observer FOV, successfully discovers and maintains a downsampled population of GEO objects, and converges on all new RSOs following a payload separation in LEO. While the filter struggles to confirm targets in LEO when many objects appear simultaneously at the same location, it reliably incorporates all subsequent tracklets once convergence is achieved.

8.2 Limitations and future work

The implementation and results presented in this work indicate several possibilities for improvement and further research.

A. Detection probability model

The classical implementation of the LMB filter was found to be sensitive to wrongly estimated detection probability, particularly because it is modelled constant inside the field of view and near-zero outside. That discontinuity at the edge can result in missed associations if the predicted state is just outside the FOV while the actual object is just inside; the opposite scenario results in

an unjustified reduction of the existence probability. While integrating target position uncertainty over the FOV partially addresses this issue, the results suggest that more realistic visibility models might result in more stable filter behaviour.

B. Tracklet clutter modelling

The Poisson clutter assumptions of the standard multi-object measurement model are not easily enforced for *tracklets as measurements*, because the definition of a false positive tracklet is ambiguous. For example, a tracklet might be considered clutter because it contains clutter measurements, but one could argue that *some* outliers in a tracklet could be acceptable as long as the true positives all belong to the same object. Conversely, a tracklet formed with observations from more than one object should also be considered invalid, even though it contains no actual clutter measurements. For a mathematically rigorous filter, it is necessary to define a consistent model for tracklet clutter intensity that addresses all types of false positives and integrates to the clutter rate over all possible tracklets. The multi-scan filter from Vo and Vo [129][134, Sec. IV-F] might provide the relevant framework to obtain such a model, although some domain-specific knowledge on the tracklet formation process should likely be introduced.

C. Initial orbit determination

In this thesis, two separate birth models were considered; the CAR approach worked well if tracklets were almost linear in right ascension and declination, whereas a batch LS optimisation initialised by Gooding’s method was promising for strongly curved tracklets of sufficient duration. A combined approach that uses the appropriate method based on the tracklet geometry would make the filter more robust and adaptable to scenarios with objects in various orbit regimes and multiple observers in space and on ground.

Furthermore, for objects in LEO observed from space, the admissible regions often result in high uncertainties for the orientation of the orbital plane. A probabilistic admissible region (PAR) approach could help further define the density [110]. Additionally, very short tracklets would benefit from a multi-tracklet birth model. Analogous to the approach used by Cai et al. [16], semi-analytical methods could be used on different combinations of tracklets from two or three consecutive groups to obtain a birth density.

D. Smoothing and multi-scan LMB

In the early stages of object discovery, association confusion often results in *track switching*, before the various uncertain birth components converge on a specific object. After that, the tracklet filter can effectively perform tracklet association, but the early tracklets remain incorrectly assigned. To update the association results for those observations, a smoothing [123] version of the tracklet LMB filter could be considered, using the converged state to go back in time and resolve the previous uncertainties in the tracklet assignment. Stauch et al. [120] successfully applied a similar approach for single measurements association, but an extension to tracklet-level inference could improve the performance of the tracklet LMB filter.

E. Parameter consistency

This work has shown that incorrect choices for certain configuration parameters, particularly the GM and LMB reduction thresholds, can be detrimental to filter accuracy. For instance, if the maximum number of GM components is set too low, essential birth components may be silently discarded (see Section 7.2.4), or ambiguous measurement updates could be inadequately represented (Section 7.4.3). Although these effects are intuitively understood, tracing specific issues back to individual configuration parameters requires detailed knowledge of the methods and their interdependencies. A better understanding of potential failure modes would enable the

implementation of automated checks to detect or warn against problematic configurations, and could even allow for dynamic adaptation mechanisms to make the methods more robust, while reducing the burden of manual parameter tuning.

F. Accumulated probabilities for adaptive birth

In this work, the *accumulated association probability* was used to extract *tracklet associations* from the traditional (single-scan) LMB filter. However, one could also use the accumulated probability that a tracklet is *not* associated to identify UCTs, which may in turn feed back into the adaptive birth process.

G. Tracklets from multiple sensors

The formulation of the traditional LMB filter and its implementation in BASIL do not set any limitations on the variety of measurement types and number of observers used in the estimation process. Specifically for SSA, all measurement models in Orekit [85] are supported out of the box. However, the validity of the new tracklet LMB filter heavily depends on the assumptions that valid *tracklet groups* can be formed that guarantee object uniqueness. The current approach relies on the notion of a minimum time between passes through the FOV, but this is not necessarily relevant in the general multi-sensor case. A more versatile tracklet grouping approach would be necessary to robustly deal with more than one observer.

REFERENCES

- [1] Anselmo, L. and Pardini, P., “Analysis of the consequences in low earth orbit of the collision between COSMOS 2251 and IRIIDIUM 33,” in *21st international symposium on space flight dynamics*, pp. 1–15, tolosa, francia, 28 settembre - 2 ottobre 2009, CNES, Toulouse , Francia, 2009.
- [2] Arasaratnam, I. and Haykin, S., “Cubature Kalman Filters,” *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1254–1269, Jun. 2009, doi: [10.1109/TAC.2009.2019800](https://doi.org/10.1109/TAC.2009.2019800).
- [3] Aristoff, J., Horwood, J., Singh, N., Poore, A., Sheaff, C., and Jah, M., “Multiple hypothesis tracking (MHT) for space surveillance: Theoretical framework,” *Advances in the Astronautical Sciences*, vol. 150, pp. 55–74, Jan. 2014.
- [4] Armellin, R. and Di Lizia, P., “Probabilistic initial orbit determination,” Feb. 2016.
- [5] Bar-Shalom, Y. and Fortmann, T. E., *Tracking and data association*, 3. print., no. 179. in Mathematics in science and engineering. Boston, Mass.: Acad. Press, 1991.
- [6] Bar-Shalom, Y., Daum, F., and Huang, J., “The probabilistic data association filter,” *IEEE Control Systems*, vol. 29, no. 6, pp. 82–100, Dec. 2009, doi: [10.1109/MCS.2009.934469](https://doi.org/10.1109/MCS.2009.934469).
- [7] Bar-Shalom, Y., Willett, P. K., and Tian, X., *Tracking and data fusion: a handbook of algorithms*. Storrs, CT: YBS Publishing, 2011.
- [8] Beard, M., Vo, B.-T., and Vo, B.-N., “A solution for large-scale multi-object tracking,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 2754–2769, 2020, Accessed: Apr. 19, 2024. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9063553/>
- [9] Bertsekas, D. P. and Castanon, D. A., “The auction algorithm for the transportation problem,” *Annals of Operations Research*, vol. 20, no. 1, pp. 67–96, Dec. 1989, doi: [10.1007/BF02216923](https://doi.org/10.1007/BF02216923).
- [10] Blackman, S. S., “Multiple hypothesis tracking for multiple target tracking,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 5–18, Jan. 2004, doi: [10.1109/MAES.2004.1263228](https://doi.org/10.1109/MAES.2004.1263228).
- [11] Blackman, S. S. and Popoli, R., *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.
- [12] Boucher, C., Herring, T., Altamimi, Z., and Pavlis, E. C., “The terrestrial reference frame and the dynamic Earth,” Jan. 2001, doi: [10.1029/EO082I025P00273-01](https://doi.org/10.1029/EO082I025P00273-01).
- [13] Brekke, E. F. and Chitre, M., “The multiple hypothesis tracker derived from finite set statistics,” in *2017 20th International Conference on Information Fusion*, Jul. 2017, pp. 1–8. doi: [10.23919/ICIF.2017.8009708](https://doi.org/10.23919/ICIF.2017.8009708).
- [14] Cai, H., Gehly, S., Yang, Y., and Zhang, K., “Modeling Birth for the Labeled Multi-Bernoulli Filter Using a Boundary-Value Approach,” *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 1, pp. 162–169, Jan. 2020, doi: [10.2514/1.G004112](https://doi.org/10.2514/1.G004112).
- [15] Cai, H., Yang, Y., Gehly, S., He, C., and Jah, M., “Sensor tasking for search and catalog maintenance of geosynchronous space objects,” *Acta Astronautica*, vol. 175, pp. 234–248, Oct. 2020, doi: [10.1016/j.actaastro.2020.05.063](https://doi.org/10.1016/j.actaastro.2020.05.063).
- [16] Cai, H., Yang, Y., Gehly, S., Wu, S., and Zhang, K., “Improved tracklet association for space objects using short-arc optical measurements,” *Acta Astronautica*, vol. 151, pp. 836–847, Oct. 2018, doi: [10.1016/j.actaastro.2018.07.024](https://doi.org/10.1016/j.actaastro.2018.07.024).
- [17] Cament, L., Adams, M., and Barrios, P., “Space Debris Tracking with the Poisson Labeled Multi-Bernoulli Filter,” *Sensors*, vol. 21, no. 11, p. 3684, Jan. 2021, doi: [10.3390/s21113684](https://doi.org/10.3390/s21113684).
- [18] Chamarro, A., “Angles-only Initial Orbit Determination and SST cataloguing,” 2021.
- [19] Chiu, S. N., Stoyan, D., Kendall, W. S., and Mecke, J., *Stochastic Geometry and Its Applications*. John Wiley & Sons, 2013.
- [20] Choi, E.-J., Cho, S., Jo, J. H., Park, J.-H., Chung, T., Park, J., Jeon, H., et al., “Performance Analysis of Sensor Systems for Space Situational Awareness,” *Journal of Astronomy and Space Sciences*, vol. 34, no. 4, pp. 303–313, 2017, doi: [10.5140/JASS.2017.34.4.303](https://doi.org/10.5140/JASS.2017.34.4.303).
- [21] Chui, C. K. and Chen, G., *Kalman Filtering*. Cham: Springer International Publishing, 2017. doi: [10.1007/978-3-319-47612-4](https://doi.org/10.1007/978-3-319-47612-4).

- [22] Clohessy, W. H. and Wiltshire, R. S., "Terminal Guidance System for Satellite Rendezvous," *Journal of the Aerospace Sciences*, vol. 27, no. 9, pp. 653–658, Sep. 1960, doi: [10.2514/8.8704](https://doi.org/10.2514/8.8704).
- [23] Coraluppi, S. and Carthel, C., "Multi-Stage Multiple-Hypothesis Tracking," *J. Adv. Inf. Fusion*, 2011, Accessed: Mar. 18, 2025. [Online]. Available: <https://www.semanticscholar.org/paper/Multi-Stage-Multiple-Hypothesis-Tracking-Coraluppi-Carthel/c8a8c2f572bf55e11f068f684b2dfeaf9958356c>
- [24] Cover, T. M. and Thomas, J. A., *Elements of information theory*, 2nd ed. Hoboken, N.J: Wiley-Interscience, 2006.
- [25] Cox, I. and Hingorani, S., "An efficient implementation of Reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 2, pp. 138–150, Feb. 1996, doi: [10.1109/34.481539](https://doi.org/10.1109/34.481539).
- [26] Curtis, H. D., *Orbital mechanics for engineering students*, Fourth edition. in Elsevier aerospace engineering series. Kidlington, Oxford, United Kingdom: Butterworth-Heinemann, an imprint of Elsevier, 2020.
- [27] Daley, D. J. and Vere-Jones, D., *An Introduction to the Theory of Point Processes*. in Probability and its Applications. New York: Springer-Verlag, 2003. doi: [10.1007/b97277](https://doi.org/10.1007/b97277).
- [28] DeMars, K. J. and Jah, M. K., "Probabilistic Initial Orbit Determination Using Gaussian Mixture Models," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 5, pp. 1324–1335, Sep. 2013, doi: [10.2514/1.59844](https://doi.org/10.2514/1.59844).
- [29] DeMars, K. J., Bishop, R. H., and Jah, M. K., "Entropy-Based Approach for Uncertainty Propagation of Nonlinear Dynamical Systems," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 4, pp. 1047–1057, Jul. 2013, doi: [10.2514/1.58987](https://doi.org/10.2514/1.58987).
- [30] DeMars, K. J., Jah, M. K., and Schumacher, P. W., "Initial Orbit Determination using Short-Arc Angle and Angle Rate Data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 3, pp. 2628–2637, Jul. 2012, doi: [10.1109/TAES.2012.6237613](https://doi.org/10.1109/TAES.2012.6237613).
- [31] Dhole, A., "EnduroSat to build Vyoma's space situational awareness constellation." Accessed: Jun. 01, 2025. [Online]. Available: <https://www.vyoma.space/news-items/endurosat-to-build-vyomas-space-situational-awareness-constellation>
- [32] Eppstein, D., "Finding the k Shortest Paths," *SIAM Journal on Computing*, vol. 28, no. 2, pp. 652–673, Jan. 1998, doi: [10.1137/S0097539795290477](https://doi.org/10.1137/S0097539795290477).
- [33] Erdinc, O., Willett, P., and Bar-Shalom, Y., "Probability hypothesis density filter for multitarget multisensor tracking," in *2005 7th International Conference on Information Fusion*, Jul. 2005, p. 8. doi: [10.1109/ICIF.2005.1591848](https://doi.org/10.1109/ICIF.2005.1591848).
- [34] ESA Space Debris Office, O., "ESA'S Annual Space Environment Report." Accessed: Nov. 18, 2024. [Online]. Available: https://www.sdo.esoc.esa.int/environment_report/Space_Environment_Report_latest.pdf
- [35] Escribano, G., Brandon, J., Sanjurjo-Rivo, M., Siminski, J., Pastor, A., and Escobar, D., "A GLMB Filter for Space Objects with Control Metric Based Maneuvre Detection," *American Astronomical Society*, 2023.
- [36] Fantacci, C., Vo, B.-T., Papi, F., and Vo, B.-N., "The Marginalized δ -GLMB Filter," *IEEE Signal Processing Letters*, vol. 23, no. 6, pp. 863–867, Jun. 2016, doi: [10.1109/LSP.2016.2557078](https://doi.org/10.1109/LSP.2016.2557078).
- [37] Fantacci, C., Vo, B.-N., Vo, B.-T., Battistelli, G., and Chisci, L., "Robust Fusion for Multisensor Multiobject Tracking," *IEEE Signal Processing Letters*, vol. 25, no. 5, pp. 640–644, May 2018, doi: [10.1109/LSP.2018.2811750](https://doi.org/10.1109/LSP.2018.2811750).
- [38] Flohrer, T., Krag, H., Klinkrad, H., and Schildknecht, T., "Feasibility of performing space surveillance tasks with a proposed space-based optical architecture," vol. 38, p. 5, Jan. 2010, Accessed: Mar. 18, 2025. [Online]. Available: <https://ui.adsabs.harvard.edu/abs/2010cosp...38.3913F>
- [39] Fortmann, T., Bar-Shalom, Y., and Scheffe, M., "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE Journal of Oceanic Engineering*, vol. 8, no. 3, pp. 173–184, Jul. 1983, doi: [10.1109/JOE.1983.1145560](https://doi.org/10.1109/JOE.1983.1145560).
- [40] Franken, D., Schmidt, M., and Ulmke, M., "'Spooky Action at a Distance' in the Cardinalized Probability Hypothesis Density Filter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 4, pp. 1657–1664, Oct. 2009, doi: [10.1109/TAES.2009.5310327](https://doi.org/10.1109/TAES.2009.5310327).
- [41] Frueh, C., Fiedler, H., Schildknecht, T., and Herzog, J., "Multi-Target Tracking for Smartnet: Multi-Layer Probability Hypothesis Filter for Near-Earth Object Tracking," 2021.
- [42] Frueh, C. and Jah, M., "Detection probability of earth orbiting objects using optical sensors," *Advances in the Astronautical Sciences*, vol. 150, pp. 3–14, Jan. 2014.

- [43] Früh, C. and Schildknecht, T., “Object image linking of earth orbiting objects in the presence of cosmoics,” *Advances in Space Research*, vol. 49, no. 3, pp. 594–602, Feb. 2012, doi: [10.1016/j.asr.2011.10.021](https://doi.org/10.1016/j.asr.2011.10.021).
- [44] Fujimoto, K. and Scheeres, D. J., “Applications of the admissible region to space-based observations,” *Advances in Space Research*, vol. 52, no. 4, pp. 696–704, Aug. 2013, doi: [10.1016/j.asr.2013.04.020](https://doi.org/10.1016/j.asr.2013.04.020).
- [45] García-Fernández, Á. F. and Svensson, L., “Spooky effect in optimal OSPA estimation and how GOSPA solves it,” 2019, doi: [10.48550/ARXIV.1908.08815](https://doi.org/10.48550/ARXIV.1908.08815).
- [46] García-Fernández, Á. F., Williams, J. L., Granström, K., and Svensson, L., “Poisson Multi-Bernoulli Mixture Filter: Direct Derivation and Implementation,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1883–1901, Aug. 2018, doi: [10.1109/TAES.2018.2805153](https://doi.org/10.1109/TAES.2018.2805153).
- [47] Gehly, S., “Multitarget Estimation, Tracking, and Information Synthesis (METIS).” Accessed: Jan. 24, 2025. [Online]. Available: <https://github.com/gehly/metis>
- [48] Gehly, S., Langbroek, M., and Carter, B., “Optical Tracklet Correlation and Adaptive Birth Modeling for Random Finite Set Filters,” 2023.
- [49] Gehly, S., “Estimation of Geosynchronous Space Objects Using Finite Set Statistics Filtering Methods,” 2016. [Online]. Available: https://www.colorado.edu/ccar/sites/default/files/attached-files/estimation_of_geosynchronous_s.pdf
- [50] Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B., *Bayesian data analysis*, Third edition. in Texts in statistical science series. Boca Raton London New York: CRC Press, Taylor, Francis Group, 2014.
- [51] Genz, A., “Numerical computation of rectangular bivariate and trivariate normal and t probabilities,” *Statistics and Computing*, vol. 14, no. 3, pp. 251–260, Aug. 2004, doi: [10.1023/B:STCO.0000035304.20635.31](https://doi.org/10.1023/B:STCO.0000035304.20635.31).
- [52] Gooding, R., “A New Procedure for Orbit Determination Based on Three Lines of Sight (Angles Only),” Apr. 1993. Accessed: Apr. 05, 2024. [Online]. Available: <https://www.semanticscholar.org/paper/A-New-Procedure-for-Orbit-Determination-Based-on-of-Gooding/49ae3aed1a3d5023487d3c7d8bb641fef88e749c>
- [53] Gooding, R. H., “A new procedure for the solution of the classical problem of minimal orbit determination from three lines of sight,” *Celestial Mechanics and Dynamical Astronomy*, vol. 66, no. 4, pp. 387–423, Dec. 1996, doi: [10.1007/BF00049379](https://doi.org/10.1007/BF00049379).
- [54] Goodman, I. R., Mahler, R., and Nguyen, H. T., *Mathematics of Data Fusion*. Dordrecht: Springer Netherlands, 1997. doi: [10.1007/978-94-015-8929-1](https://doi.org/10.1007/978-94-015-8929-1).
- [55] Gordon, N., Salmond, D., and Smith, A., “Novel approach to nonlinear/non-Gaussian Bayesian state estimation,” *IEE Proceedings F (Radar and Signal Processing)*, vol. 140, no. 2, pp. 107–113, Apr. 1993, doi: [10.1049/ip-f-2.1993.0015](https://doi.org/10.1049/ip-f-2.1993.0015).
- [56] Hoang, H. G., Vo, B.-N., Vo, B.-T., and Mahler, R., “The Cauchy–Schwarz Divergence for Poisson Point Processes,” *IEEE Transactions on Information Theory*, vol. 61, no. 8, pp. 4475–4485, Aug. 2015, doi: [10.1109/TIT.2015.2441709](https://doi.org/10.1109/TIT.2015.2441709).
- [57] Holzinger, M. J., Scheeres, D. J., and Alfriend, K. T., “Object Correlation, Maneuver Detection, and Characterization Using Control Distance Metrics,” *Journal of Guidance, Control, and Dynamics*, vol. 35, no. 4, pp. 1312–1325, Jul. 2012, doi: [10.2514/1.53245](https://doi.org/10.2514/1.53245).
- [58] Horwood, J. T. and Poore, A. B., “Gauss von Mises Distribution for Improved Uncertainty Realism in Space Situational Awareness,” *SIAM/ASA Journal on Uncertainty Quantification*, vol. 2, no. 1, pp. 276–304, Jan. 2014, doi: [10.1137/130917296](https://doi.org/10.1137/130917296).
- [59] Horwood, J. T., Aragon, N. D., and Poore, A. B., “Gaussian Sum Filters for Space Surveillance: Theory and Simulations,” *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1839–1851, Nov. 2011, doi: [10.2514/1.53793](https://doi.org/10.2514/1.53793).
- [60] IERS, “IERS Conventions Centre.” Accessed: Jun. 02, 2025. [Online]. Available: https://iers-conventions.obspm.fr/conventions_material.php
- [61] Ito, K. and Xiong, K., “Gaussian filters for nonlinear filtering problems,” *IEEE Transactions on Automatic Control*, vol. 45, no. 5, pp. 910–927, May 2000, doi: [10.1109/9.855552](https://doi.org/10.1109/9.855552).
- [62] Izzo, D., “Revisiting Lambert’s problem,” *Celestial Mechanics and Dynamical Astronomy*, vol. 121, no. 1, pp. 1–15, Jan. 2015, doi: [10.1007/s10569-014-9587-y](https://doi.org/10.1007/s10569-014-9587-y).
- [63] Jacobsen, M., Ed., *Point Process Theory and Applications: Marked Point and Piecewise Deterministic Processes*. Boston, MA: Birkhäuser, 2006. doi: [10.1007/0-8176-4463-6_2](https://doi.org/10.1007/0-8176-4463-6_2).

- [64] Jones, B. A. and Weisman, R., "Multi-fidelity orbit uncertainty propagation," *Acta Astronautica*, vol. 155, pp. 406–417, Feb. 2019, doi: [10.1016/j.actaastro.2018.10.023](https://doi.org/10.1016/j.actaastro.2018.10.023).
- [65] Jones, B. A., Bryant, D. S., Vo, B.-T., and Vo, B.-N., "Challenges of multi-target tracking for space situational awareness," in *2015 18th International Conference on Information Fusion*, IEEE, 2015, pp. 1278–1285. Accessed: Dec. 10, 2024. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/7266704/>
- [66] Jonker, R. and Volgenant, T., "Improving the Hungarian assignment algorithm," *Operations Research Letters*, vol. 5, no. 4, pp. 171–175, Oct. 1986, doi: [10.1016/0167-6377\(86\)90073-8](https://doi.org/10.1016/0167-6377(86)90073-8).
- [67] Julier, S. and Uhlmann, J., "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, Mar. 2004, doi: [10.1109/JPROC.2003.823141](https://doi.org/10.1109/JPROC.2003.823141).
- [68] Julier, S. J. and Uhlmann, J. K., "New extension of the Kalman filter to nonlinear systems," in *Signal Processing, Sensor Fusion, and Target Recognition VI*, SPIE, Jul. 1997, pp. 182–193. doi: [10.1117/12.280797](https://doi.org/10.1117/12.280797).
- [69] Kay, S. M., *Fundamentals of statistical signal processing. 1: Estimation theory*, 20. pr. Upper Saddle River, NJ: Prentice Hall PTR, 2013.
- [70] Kennewell, J. A. and Vo, B.-N., "An overview of space situational awareness," in *Proceedings of the 16th International Conference on Information Fusion*, Jul. 2013, pp. 1029–1036. Accessed: Dec. 09, 2024. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/6641108>
- [71] Krag, H., Serrano, M., Braun, V., Kuchynka, P., Catania, M., Siminski, J., Schimmerohn, M., et al., "A 1 cm space debris impact onto the Sentinel-1A solar array," *Acta Astronautica*, vol. 137, pp. 434–443, Aug. 2017, doi: [10.1016/j.actaastro.2017.05.010](https://doi.org/10.1016/j.actaastro.2017.05.010).
- [72] Kuhn, H. W., "The Hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1–2, pp. 83–97, 1955, doi: [10.1002/nav.3800020109](https://doi.org/10.1002/nav.3800020109).
- [73] Kullback, S. and Leibler, R. A., "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, Mar. 1951, doi: [10.1214/aoms/1177729694](https://doi.org/10.1214/aoms/1177729694).
- [74] Kurien, T., "Issues in the design of practical multitarget tracking algorithms," *Multitarget-multisensor tracking: advanced applications*, pp. 43–84, 1990, Accessed: Dec. 15, 2024. [Online]. Available: <https://cir.nii.ac.jp/crid/1573387451142656512>
- [75] Lee, B.-s., Kim, W.-G., Lee, J., and Hwang, Y., "Machine Learning Approach to Initial Orbit Determination of Unknown LEO Satellites," *2018 SpaceOps Conference*. in SpaceOps Conferences. American Institute of Aeronautics, Astronautics, May 2018. doi: [10.2514/6.2018-2566](https://doi.org/10.2514/6.2018-2566).
- [76] Leushacke, L., Mehrholz, D., Flury, W., Jehn, R., Klinkrad, K., and Landgraf, M., "Detecting, Tracking and Imaging Space Debris," *ESA Publications*, 2002, [Online]. Available: https://www.esa.int/esapub/bulletin/bullet109/chapter16_bul109.pdf
- [77] Lubey, D. P., "Maneuver Detection and Reconstruction in Data Sparse Systems with an Optimal Control Based Estimator," 2015.
- [78] Lubey, D. P. and Scheeres, D. J., "Combined optimal control and state estimation for the purposes of maneuver detection and reconstruction," in *2014 American Control Conference*, Jun. 2014, pp. 1749–1754. doi: [10.1109/ACC.2014.6859260](https://doi.org/10.1109/ACC.2014.6859260).
- [79] Lubey, D. P. and Scheeres, D. J., "Automated state and dynamics estimation in dynamically mismodeled systems with information from optimal control policies," in *2015 18th International Conference on Information Fusion*, Jul. 2015, pp. 1299–1306. Accessed: Sep. 06, 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/7266707>
- [80] Mahalanobis, P. C., "On test and measures of group divergence : theoretical formulae," 1930, Accessed: Dec. 09, 2024. [Online]. Available: <http://library.isical.ac.in:8080/xmlui/handle/10263/1639>
- [81] Mahler, R., "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003, doi: [10.1109/TAES.2003.1261119](https://doi.org/10.1109/TAES.2003.1261119).
- [82] Mahler, R., *Statistical multisource-multitarget information fusion*. in Artech House information warfare library. Boston London: Artech House, 2007.
- [83] Mahler, R., "PHD filters of higher order in target number," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1523–1543, Oct. 2007, doi: [10.1109/TAES.2007.4441756](https://doi.org/10.1109/TAES.2007.4441756).
- [84] Mahler, R., *Advances in statistical multisource-multitarget information fusion*. in Artech House electronic warfare library. Boston London: Artech House, 2014.

- [85] Maisonobe, L., Pommier, V., and Parraud, P., *OREKIT: AN OPEN SOURCE LIBRARY FOR OPERATIONAL FLIGHT DYNAMICS APPLICATIONS*. 2010.
- [86] Mallick, M., Coraluppi, S., and Carthel, C., “Multitarget Tracking Using Multiple Hypothesis Tracking,” *Integrated Tracking, Classification, and Sensor Management*. Wiley, pp. 163–203, Sep. 2016. doi: [10.1002/9781118450550.ch05](https://doi.org/10.1002/9781118450550.ch05).
- [87] Maruskin, J. M., Scheeres, D. J., and Alfriend, K. T., “Correlation of Optical Observations of Objects in Earth Orbit,” *Journal of Guidance, Control, and Dynamics*, vol. 32, no. 1, pp. 194–209, 2009, doi: [10.2514/1.36398](https://doi.org/10.2514/1.36398).
- [88] Milani, A., Gronchi, G., Vitturi, M., and Knežević, Z., “Orbit determination with very short arcs. I admissible regions,” *Celestial Mechanics and Dynamical Astronomy*, vol. 90, no. 1–2, pp. 59–87, 2004, doi: [10.1007/s10569-004-6593-5](https://doi.org/10.1007/s10569-004-6593-5).
- [89] Mishra, U. R., Faber, W., Chakravorty, S., Hussein, I., Sunderland, B., and Hesar, S., “Geometric Solution to Probabilistic Admissible Region Based Track Initialization,” in *Dynamic Data Driven Applications Systems*, E. Blasch, F. Darema, and A. Aved, Eds., Cham: Springer Nature Switzerland, 2024, pp. 71–80. doi: [10.1007/978-3-031-52670-1_7](https://doi.org/10.1007/978-3-031-52670-1_7).
- [90] Montenbruck, O. and Gill, E., *Satellite Orbits*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000. doi: [10.1007/978-3-642-58351-3](https://doi.org/10.1007/978-3-642-58351-3).
- [91] Murty, K. G., “An Algorithm for Ranking all the Assignments in Order of Increasing Cost,” *Operations Research*, vol. 16, no. 3, pp. 682–687, 1968, Accessed: Dec. 16, 2024. [Online]. Available: <https://www.jstor.org/stable/168595>
- [92] NAIF, “Frames and Coordinate Systems.” Accessed: May 01, 2025. [Online]. Available: https://naif.jpl.nasa.gov/pub/naif/toolkit_docs/Tutorials/pdf/individual_docs/17_frames_and_coordinate_systems.pdf
- [93] NASA, “Spherical Harmonic Representation of the Gravity Field Potential.” Accessed: Jun. 02, 2025. [Online]. Available: https://spsweb.ftps.jpl.nasa.gov/portaldataops/mpg/MPG_Docs/Source%20Docs/gravity-SphericalHarmonics.pdf
- [94] NASA, “Orbital Debris Quarterly News 29-2,” *NASA Orbital Debris Program Office (ODPO)*, vol. 29, no. 2, 2025, [Online]. Available: <https://orbitaldebris.jsc.nasa.gov/quarterly-news/pdfs/ODQNv29i2.pdf>
- [95] Nguyen, T. T. D., Rezatofighi, H., Vo, B.-N., Vo, B.-T., Savarese, S., and Reid, I., “How Trustworthy are Performance Evaluations for Basic Vision Tasks?,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 7, pp. 8538–8552, Jul. 2023, doi: [10.1109/TPAMI.2022.3227571](https://doi.org/10.1109/TPAMI.2022.3227571).
- [96] Papi, F., Vo, B.-N., Vo, B.-T., Fantacci, C., and Beard, M., “Generalized Labeled Multi-Bernoulli Approximation of Multi-Object Densities,” *IEEE Transactions on Signal Processing*, vol. 63, no. 20, pp. 5487–5497, Oct. 2015, doi: [10.1109/TSP.2015.2454478](https://doi.org/10.1109/TSP.2015.2454478).
- [97] Pearce, T., Brintrup, A., and Zhu, J., “Understanding Softmax Confidence and Uncertainty.” Accessed: Jun. 04, 2025. [Online]. Available: <http://arxiv.org/abs/2106.04972>
- [98] Pirovano, L., “Cataloguing space debris: Methods for optical data association,” 2020. Accessed: Mar. 11, 2025. [Online]. Available: <https://openresearch.surrey.ac.uk/esploro/outputs/doctoral/Cataloguing-space-debris-Methods-for-optical/99514583702346>
- [99] Pirovano, L., Armellin, R., Siminski, J., and Flohrer, T., “Differential algebra enabled multi-target tracking for too-short arcs,” *Acta Astronautica*, vol. 182, pp. 310–324, May 2021, doi: [10.1016/j.actaastro.2021.02.023](https://doi.org/10.1016/j.actaastro.2021.02.023).
- [100] Pulford, G., “Taxonomy of multiple target tracking methods,” *IEE Proceedings - Radar, Sonar and Navigation*, vol. 152, no. 5, p. 291, 2005, doi: [10.1049/ip-rsn:20045064](https://doi.org/10.1049/ip-rsn:20045064).
- [101] Rahmathullah, A. S., García-Fernández, Á. F., and Svensson, L., “Generalized optimal sub-pattern assignment metric,” in *2017 20th International Conference on Information Fusion*, Jul. 2017, pp. 1–8. doi: [10.23919/ICIF.2017.8009645](https://doi.org/10.23919/ICIF.2017.8009645).
- [102] Reid, D., “An algorithm for tracking multiple targets,” *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, Dec. 1979, doi: [10.1109/TAC.1979.1102177](https://doi.org/10.1109/TAC.1979.1102177).
- [103] Reid, I. and Term, H., “1 Discrete-time Kalman filter,” 2001.
- [104] Reuter, S., Vo, B.-T., Vo, B.-N., and Dietmayer, K., “The Labeled Multi-Bernoulli Filter,” *IEEE Transactions on Signal Processing*, vol. 62, no. 12, pp. 3246–3260, Jun. 2014, doi: [10.1109/TSP.2014.2323064](https://doi.org/10.1109/TSP.2014.2323064).
- [105] Ristic, B., Clark, D., Vo, B.-N., and Vo, B.-T., “Adaptive Target Birth Intensity for PHD and CPHD Filters,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 2, pp. 1656–1668, Apr. 2012, doi: [10.1109/TAES.2012.6178085](https://doi.org/10.1109/TAES.2012.6178085).

- [106] Ristic, B., Beard, M., and Fantacci, C., “An Overview of Particle Methods for Random Finite Set Models.” Accessed: Nov. 13, 2024. [Online]. Available: <http://arxiv.org/abs/1602.03945>
- [107] Ristic, B., Vo, B.-N., and Clark, D., “Performance evaluation of multi-target tracking using the OSPA metric,” in *2010 13th International Conference on Information Fusion*, Jul. 2010, pp. 1–7. doi: [10.1109/ICIF.2010.5712055](https://doi.org/10.1109/ICIF.2010.5712055).
- [108] Ristic, B., Vo, B.-N., Clark, D., and Vo, B.-T., “A Metric for Performance Evaluation of Multi-Target Tracking Algorithms,” *IEEE Transactions on Signal Processing*, vol. 59, no. 7, pp. 3452–3457, Jul. 2011, doi: [10.1109/TSP.2011.2140111](https://doi.org/10.1109/TSP.2011.2140111).
- [109] Robert, C., *The Bayesian Choice*. in Springer Texts in Statistics. New York, NY: Springer, 2007. doi: [10.1007/0-387-71599-1](https://doi.org/10.1007/0-387-71599-1).
- [110] Roscoe, C. W. T., Hussein, I. I., Wilkins, M. P., and Schumacher, P. W., “The Probabilistic Admissible Region with Additional Constraints,” 2015.
- [111] Ross, S. M., *Introduction to Probability and Statistics for Engineers and Scientists*, 5th ed. San Diego: Elsevier Science & Technology, 2014.
- [112] Rényi, A., “On Measures of Entropy and Information,” *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. University of California Press, pp. 547–562, Jan. 1961.
- [113] Schuhmacher, D. and Xia, A., “A new metric between distributions of point processes.” Accessed: Dec. 09, 2024. [Online]. Available: <http://arxiv.org/abs/0708.2777>
- [114] Schuhmacher, D., Vo, B.-T., and Vo, B.-N., “A Consistent Metric for Performance Evaluation of Multi-Object Filters,” *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3447–3457, Aug. 2008, doi: [10.1109/TSP.2008.920469](https://doi.org/10.1109/TSP.2008.920469).
- [115] Septier, F., Cornebise, J., Godsill, S., and Delignon, Y., “A comparative study of Monte-Carlo methods for multitarget tracking,” in *2011 IEEE Statistical Signal Processing Workshop (SSP)*, Jun. 2011, pp. 205–208. doi: [10.1109/SSP.2011.5967660](https://doi.org/10.1109/SSP.2011.5967660).
- [116] Siminski, J. A., Montenbruck, O., Fiedler, H., and Schildknecht, T., “Short-arc tracklet association for geostationary objects,” *Advances in Space Research*, vol. 53, no. 8, pp. 1184–1194, Apr. 2014, doi: [10.1016/j.asr.2014.01.017](https://doi.org/10.1016/j.asr.2014.01.017).
- [117] Siminski, J., “Object correlation and orbit determination in the geostationary orbit using optical measurements,” 2016. Accessed: Jun. 15, 2025. [Online]. Available: <https://athene-forschung.rz.unibw-muenchen.de/node?id=115158>
- [118] Singh, N., Horwood, J., Aristoff, J., Poore, A., Sheaff, C., and Jah, M., “Multiple hypothesis tracking (MHT) for space surveillance: Theoretical framework,” *Advances in the Astronautical Sciences*, vol. 150, pp. 55–74, Jan. 2014.
- [119] Sorenson, H. W. and Alspach, D. L., “Recursive bayesian estimation using gaussian sums,” *Automatica*, vol. 7, no. 4, pp. 465–479, Jul. 1971, doi: [10.1016/0005-1098\(71\)90097-5](https://doi.org/10.1016/0005-1098(71)90097-5).
- [120] Stauch, J., Bessell, T., Rutten, M., Baldwin, J., Jah, M., and Hill, K., “Joint Probabilistic Data Association and Smoothing Applied to Multiple Space Object Tracking,” *Journal of Guidance, Control, and Dynamics*, Dec. 2017, doi: [10.2514/1.G002230](https://doi.org/10.2514/1.G002230).
- [121] Svensson, L., “Lecture 5: Multi-object conjugate priors.” 2019.
- [122] Svensson, L., “Lecture 4: Random finite sets.” 2021.
- [123] Särkkä, S. and Svensson, L., “Bayesian Filtering and Smoothing,” 2013.
- [124] Tapley, B. D., Schutz, B. E., and Born, G. H., *Statistical Orbit Determination*. Amsterdam: Academic Press, 2004.
- [125] Vallado, D. A., *Fundamentals of astrodynamics and applications*, Fifth edition, first printing., no. 21. in Space technology library. Torrance, CA: Microcosm Press, 2022.
- [126] Vermaak, J., Godsill, S., and Perez, P., “Monte Carlo filtering for multi target tracking and data association,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 1, pp. 309–332, Jan. 2005, doi: [10.1109/TAES.2005.1413764](https://doi.org/10.1109/TAES.2005.1413764).
- [127] Vo, B.-N., Singh, S., and Doucet, A., “Sequential Monte Carlo methods for multitarget filtering with random finite sets,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1224–1245, Oct. 2005, doi: [10.1109/TAES.2005.1561884](https://doi.org/10.1109/TAES.2005.1561884).

- [128] Vo, B.-N. and Ma, W., “The Gaussian Mixture Probability Hypothesis Density Filter,” *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, Nov. 2006, doi: [10.1109/TSP.2006.881190](https://doi.org/10.1109/TSP.2006.881190).
- [129] Vo, B.-N. and Vo, B.-T., “A Multi-Scan Labeled Random Finite Set Model for Multi-Object State Estimation,” *IEEE Transactions on Signal Processing*, vol. 67, no. 19, pp. 4948–4963, Oct. 2019, doi: [10.1109/TSP.2019.2928953](https://doi.org/10.1109/TSP.2019.2928953).
- [130] Vo, B.-N., Dam, N., Phung, D., Tran, Q., and Vo, B.-T., “Model-Based Learning for Point Pattern Data,” *Pattern Recognition*, vol. 84, Jul. 2018, doi: [10.1016/j.patcog.2018.07.008](https://doi.org/10.1016/j.patcog.2018.07.008).
- [131] Vo, B.-N., Mallick, M., Bar-shalom, Y., Coraluppi, S., Osborne III, R., Mahler, R., and Vo, B.-T., “Multitarget Tracking,” *Wiley Encyclopedia of Electrical and Electronics Engineering*. John Wiley & Sons, Ltd, pp. 1–15, 2015. Accessed: Nov. 13, 2024. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/047134608X.W8275>
- [132] Vo, B.-N., Vo, B.-T., and Mahler, R. P. S., “Closed-Form Solutions to Forward-Backward Smoothing,” *IEEE Transactions on Signal Processing*, vol. 60, no. 1, pp. 2–17, Jan. 2012, doi: [10.1109/TSP.2011.2168519](https://doi.org/10.1109/TSP.2011.2168519).
- [133] Vo, B.-N., Vo, B.-T., and Phung, D., “Labeled Random Finite Sets and the Bayes Multi-Target Tracking Filter,” *IEEE Transactions on Signal Processing*, vol. 62, no. 24, pp. 6554–6567, Dec. 2014, doi: [10.1109/TSP.2014.2364014](https://doi.org/10.1109/TSP.2014.2364014).
- [134] Vo, B.-N., Vo, B.-T., Nguyen, T. T. D., and Shim, C., “An Overview of Multi-Object Estimation via Labeled Random Finite Set,” *arXiv.org*, Sep. 2024, Accessed: Oct. 08, 2024. [Online]. Available: <https://arxiv.org/abs/2409.18531v2>
- [135] Vo, B.-T. and Vo, B.-N., “Labeled Random Finite Sets and Multi-Object Conjugate Priors,” *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, Jul. 2013, doi: [10.1109/TSP.2013.2259822](https://doi.org/10.1109/TSP.2013.2259822).
- [136] Vo, B.-T. and Vo, B.-N., “RFS Tracking Toolbox.” Accessed: Oct. 19, 2024. [Online]. Available: <https://batuong.vo-au.com/codes.html>
- [137] Vo, B.-T., Vo, B.-N., and Cantoni, A., “Analytic Implementations of the Cardinalized Probability Hypothesis Density Filter,” *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3553–3567, Jul. 2007, doi: [10.1109/TSP.2007.894241](https://doi.org/10.1109/TSP.2007.894241).
- [138] Vo, B.-T., Vo, B.-N., and Cantoni, A., “The Cardinality Balanced Multi-Target Multi-Bernoulli Filter and Its Implementations,” *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 409–423, Feb. 2009, doi: [10.1109/TSP.2008.2007924](https://doi.org/10.1109/TSP.2008.2007924).
- [139] Wan, E. and Merwe, R. van der, “7 the Unscented Kalman Filter,” 2001. Accessed: Jan. 14, 2025. [Online]. Available: <https://api.semanticscholar.org/CorpusID:14862265>
- [140] Weisstein, E. W., “Chi-Squared Distribution.” Accessed: Mar. 31, 2025. [Online]. Available: <https://mathworld.wolfram.com/Chi-SquaredDistribution.html>
- [141] Williams, J. L., “Hybrid Poisson and multi-Bernoulli filters.” Accessed: Dec. 16, 2024. [Online]. Available: <http://arxiv.org/abs/1203.2992>
- [142] Williams, J. L., “Marginal multi-bernoulli filters: RFS derivation of MHT, JIPDA, and association-based member,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 3, pp. 1664–1687, Jul. 2015, doi: [10.1109/TAES.2015.130550](https://doi.org/10.1109/TAES.2015.130550).
- [143] Wright, J. R., “Orbit Determination Tool Kit,” 2013.
- [144] Wu, Y., Hu, D., Wu, M., and Hu, X., “A Numerical-Integration Perspective on Gaussian Filters,” *IEEE Transactions on Signal Processing*, vol. 54, pp. 2910–2921, Aug. 2006, doi: [10.1109/TSP.2006.875389](https://doi.org/10.1109/TSP.2006.875389).
- [145] Yen, J. Y., “Finding the K Shortest Loopless Paths in a Network,” *Management Science*, vol. 17, no. 11, pp. 712–716, 1971, Accessed: Jan. 24, 2025. [Online]. Available: <https://www.jstor.org/stable/2629312>
- [146] Zardain, L., Farres, A., and Puig, A., “High-fidelity Modeling and Visualizing of Solar Radiation Pressure A Framework for High-fidelity Analysis,” South Lake Tahoe, 2021. Accessed: Jun. 26, 2025. [Online]. Available: <https://ntrs.nasa.gov/citations/20205005240>
- [147] Zhao, G., Liu, L., Li, B., Li, Z., and Sang, J., “An orbit determination approach to associating optical tracklets of space objects,” *Acta Astronautica*, vol. 200, pp. 506–523, Nov. 2022, doi: [10.1016/j.actaastro.2022.08.044](https://doi.org/10.1016/j.actaastro.2022.08.044).
- [148] Zittersteijn, M., Vananti, A., Schildknecht, T., Dolado-Perez, J., and Martinot, V., “A Genetic Algorithm to associate optical measurements and estimate orbits of geocentric objects,” in *ResearchGate*, 2015. Accessed: Jun. 08, 2025. [Online]. Available: https://www.researchgate.net/publication/285769685_A_Genetic_Algorithm_to_associate_optical_measurements_and_estimate_orbits_of_geocentric_objects

LIST OF APPENDICES

Section	Content	
Appendix A	Early literature review	p. 117
Appendix B	Early project planning	p. 124
Appendix C	LMB verification for 2D case by comparison	p. 126
Appendix D	CAR implementation details and verification by comparison	p. 129
Appendix E	The χ^2 distribution and LogSumExp normalisation	p. 135
Appendix F	Complete association probabilities for selected simulation cases	p. 136
Appendix G	Information on simulated objects	p. 140

LITERATURE REVIEW

This appendix contains a slightly reduced version of the literature review from December 2024²⁹.

In general, estimation of any dynamical system refers to the practice of *fine-tuning* the defining components of that system based on noisy measurements. This consists of three main elements:

- (i) *state estimation*, i.e. finding the state of the system/object over time,
- (ii) *system characterisation*, which is adjusting the driving parameters of the system models – including dynamic parameters like coefficients of drag and radiation pressure–, and
- (iii) *control estimation*, referring to the systematic adjustment of state input to match the measurements [134]. The focus of this thesis is on the first element, which is also commonly referred to as *tracking*.

For the context of this work, let *tracking* be the practice of estimating the *trajectory* (or *track*) of one or more *objects* (or *targets*) over time [131].

A.1 Single-object Bayesian filtering

In general, Bayesian filtering refers to the sequential estimation of a time-varying system, based on noisy measurements [123]. Usually, these measurements are indirect observations of the system, and there exists some mapping between the measurements and the estimated state.

In statistical estimation, states and measurements are represented as PDFs, which describe the relative likelihood of each possible value for these quantities [111]. That is, integrating the PDF over some subspace of possible values yields the probability that the true value is indeed part of that subspace.

The Bayesian filter is a statistically optimal estimator that provides the posterior PDF of all estimated states, based on their assumed prior density and a set of measurements. The measurements constitute a partial observation of the states and are distributed according to a likelihood function, conditioned on the prior, i.e. given a prior PDF of the state, the measurement likelihood function can be used to determine the probability of observing something in a part of the measurement space. This update is described in mathematics as an *inversion* problem and is usually formulated using Bayes' theorem [109]. Although the description is theoretically very useful, the presence of multiple uncertain measurements, missed detections and non-linear evolution of general uncertainty causes the computational complexity to increase very quickly and become intractable for any practical application.

For the development of approximation algorithms, it is generally desirable to limit the complexity of the posterior computation by employing a *conjugate* family of distributions for the state PDF [135]. Given a measurement likelihood function, a prior density is called *conjugate prior* if the Bayes posterior always belongs to the same family. A very common choice is the Gaussian distribution, which is conjugate prior for Gaussian likelihood, and particularly attractive because of its analytical convenience and simple formulation [50, p. 40].

²⁹As this version is integrated in the same document as the thesis, abbreviations are not repeated in full. Please refer to the list of acronyms at the start of the document for reference.

The filtering complexity is often reduced by considering measurements one by one and only performing the Bayes update for the estimated state at the time of the latest measurement [124]. The closed-form solution to the linear Gaussian filtering problem is given by the Kalman filter equations. However, it is only valid if the underlying dynamics are linear in the estimated state. If this is not the case, the linear mapping itself can be propagated using Taylor series linearisation, giving rise to the EKF technique [69], at the cost of n^2 additional integrated variables, n being the dimension of the state vector. In highly non-linear cases, where this linearisation is not representative, the UKF can further improve the realism [67][68][139]. This approach directly propagates a representative sample of so-called *sigma points* and reconstructs a Gaussian prior and measurement cross-covariance at the time of interest [123]. As observed by Ito and Xiong [61], the UKF is a special case of what they call general Gaussian filtering, which uses *moment matching* to propagate assumed Gaussian densities and can be approximated using a variety of integration methods like Gauss-Hermite and spherical cubature³⁰ [2].

Asymptotic optimality can be achieved by sequential Monte Carlo (SMC) integration in particle filters (PFs) such as the bootstrap filter [55], but due to their very high computational burden compared to the analytic alternatives, these methods are mostly useful from an academic and benchmarking perspective [144].

One way to preserve non-Gaussian behaviour without resorting to an SMC approach is to represent the PDF as a GM, i.e. a weighted sum of Gaussian densities [119]. A filter such as EKF and UKF can then act independently on each component to compute the Bayesian recursion. This mixture approach also becomes particularly useful in the context of multi-target tracking, where the added complexity of uncertain measurement association can cause highly non-Gaussian posterior PDFs [29][59][64][128].

Within the field of OD, the state consists of 6 orbital parameters (for example, a Cartesian position and velocity) along with any number of estimated dynamic and measurement parameters to tune the motion and observation models [124][143]. The conventional approach in OD is to represent the state's PDF as a single Gaussian using its mean for the estimate and covariance for the uncertainty distribution. However, the exceptional sparsity of measurements in SSA often calls for long propagation times (regularly several orbital periods) so that the true state density quickly loses its Gaussian properties. For example, LEO orbits are influenced by drag effects, which cause rapidly increasing non-Gaussian uncertainties, mainly in the along-track direction [58][125]. Therefore, the orbital PDF propagation can certainly benefit from a GM representation.

The Gaussian assumption remains valid longer when classical orbital elements (COE) or modified equinoctial elements (MEE) are used to describe the object state [59]. This can reduce the required number of components in the GM, although this approach introduces other issues like singularities for specific types of orbits and additional implementation complexity.

A.2 Multiple object tracking

While single-object tracking scenarios already have the potential problem of missed detections, clutter and observation noise, the main challenge posed in the MO case is that of measurement association [5].

³⁰in the context of numerical integration, cubature refers to multi-dimensional integrals, as an extension of quadrature, which is generally reserved for the one-dimensional case.

A.2.1 Tracking after association

Traditional approaches consider the association step separately and then perform single-object updates afterwards. The simplest solutions include the nearest neighbour and strongest neighbour Kalman filters, which only consider the observations closest in statistical distance for each of the single-object filters [100]. Methods like probabilistic data association (PDA) go one step further by including several likely associations in the update and then condensing the GM posterior back to one Gaussian component [6].

A.2.2 Tracking while associating

Note that the single-object approaches have the risk of ambiguous data associations. It is possible for one observation to be a “strong neighbour” of multiple objects. “True” MOT algorithms treat the association problem by considering all measurements and objects simultaneously. The most straightforward example is the global nearest neighbour (GNN) algorithm, which assigns a known number of objects to a set of measurements, by minimising a global association distance metric [11].

The JPDA filter [7][39] instead considers a collection of feasible joint association events, using a validity gate for each object to limit the number of options. The marginal probabilities of the events are then computed jointly, after which uncoupled Bayesian updates can be performed for each object, while merging the update components using a moment matching approach [131]. Although this method has been very popular and effective, one of its drawbacks is the reduced performance in case of densely spaced targets. This is because it merges information related to wrongly associated observations [11][49].

Alternatively, MHT is a method that keeps a separate global hypothesis for any permutation of measurement associations over time [10][86][102]. At every time step, all the valid measurement-to-track associations (M2TAs) are applied to each of the previous hypotheses to update the global hypotheses. In this way, the algorithm delays the association decision, allowing multiple measurement scans to have influence on the result [10]. Since the number of hypotheses can grow exponentially [131], a combination of pruning, merging, clustering and capping the components is usually employed to limit the complexity. The original hypothesis-oriented formulation from Reid [102] has an additional drawback related to the fact that the same association sequences can be present in multiple global hypotheses. For this reason, a track-oriented (TO) formulation as proposed by Kurien [74] can be more desirable, as it stores independent association sequences for each of the objects, together with a lookup table to indicate which combinations of those tracks form the global hypotheses. An additional benefit is that component reduction happens on the track level (a complete sequence of M2TAs) so that the individual hypotheses do not lose any statistical significance³¹ [10]. A variety of other MHT flavours exist, some of which have been successfully applied to SSA [3][70].

A new suite of methods based on FISST has been under rapid development over the last two of decades, after being introduced as a simplification of simple point process theory by Mahler [82][84]. Instead of state and measurement vectors as the primary element of the Bayesian recursion, these methods represent the tracked objects and observations as random finite sets (RFSs) of vectors. Using the notion of multi-object densities and set integrals [54][81] allows the MOT

³¹In the original *measurement-oriented* MHT, the tracks are formed by the superposition of all likely hypotheses at every time step and the association sequence is implicit. Local hypotheses with a single, very unlikely association might be pruned out, thereby reducing the weight of all other (potentially correct) associations in that hypothesis. TO MHT does not have this issue, as association histories are pruned out for every target individually.

problem to be formulated equivalently to the single-object counterpart and the state estimation and measurement association to be performed as one.

Many variations of RFSs exist, but an essential feature is that their MO density contains information on both the distribution of the objects in the state space *and* the cardinality distribution. This means that the number of elements in the set is also a random variable with a certain distribution, and can thus change based on the Bayes measurement update. One of the most common examples is the Poisson RFS or PPP³² [27]. Its elements are i.i.d. based on some single-state density and its cardinality is Poisson distributed. The PPP with uniform PDF is a popular choice for modelling clutter or object birth. Note that the PPP is a special type of i.i.d. Cluster RFS, which allows for any type of cardinality distribution, as long as the components are i.i.d., conditioned on that cardinality [84]. Another useful form is the Bernoulli RFS, which is parametrised by a single-object PDF and an existence probability. The resulting set is therefore either empty or a singleton. By taking the union of independent Bernoulli's, this model can easily be extended to the multi-Bernoulli RFS [134], which can be a very intuitive representation of hypothesised objects in an MOT context.

The RFS framework allows for a statistically optimal formulation of the Bayes' filter for MO systems [106], but since any implementation is intractable for realistic time spans and target counts, many approximations have emerged.

The first-moment approximation of an RFS is its intensity function or PHD, defined such that it integrates to the expected cardinality of the set [81]. Note that this is equivalent to approximating the RFS as a PPP. For practical purposes, the evolution of the PHD can be modelled using SMC filters [115][126][127] or a GM representation [128]. These filters have simple, efficient implementations and work well in straightforward tracking scenarios [42][49], but the PHD's implicit cardinality estimate is particularly erratic in scenarios with dense clutter or frequent missed detections [33]. Mahler [83] addresses this problem with the cardinalised probability hypothesis density (CPHD) filter, which propagates the cardinality distribution alongside the PHD. This keeps the state estimates first-order, but allows for higher-order propagation of the target count and also generalises the PPP assumption to become i.i.d. cluster since the cardinality can now have an arbitrary distribution. Closed-form solutions to this formulation exist and indeed show a drastic reduction in the variance on the number of objects [105][137]. One downside of the CPHD is a complication called the *spooky effect*. If an object is not detected, the local cardinality estimate reduces and the filter compensates this by increasing the PHD near detected objects at arbitrarily large distances from that missed detection [40][45]. This can be very undesirable in sparse measurement cases such as SSA. It is possible to use measurement gating and clustering methods to alleviate this issue [49][84], but more rigorous applications of FISST may be preferable.

Multi-Bernoulli random finite set (MB) filters were suggested by Mahler [84], where each Bernoulli component in an MB represents a hypothesised object and target birth is also modelled as independent Bernoulli components. It was later shown that this filter is biased in the number of targets, mostly in high clutter scenarios. Vo et al. [138] derive the expected bias and present a *cardinality-balanced* version of the filter with SMC and GM implementations, but it only outperforms the CPHD filter in highly non-linear scenarios where GM approximations are invalid.

³²In the context of MOT, the terms *Poisson RFS* and *PPP* are synonymous and can be found interchangeably in the literature.

Multi-Bernoulli mixture (MBM) filters combine the intent of the MB filter with something resembling TO-MHT to model likely association sequences [46][142]. Every component in the mixture then represents a global association hypothesis with a weight that represents the likelihoods and an MB representation of the MO state as described before. Williams [142] suggests a PPP to represent undetected objects and simultaneously model target birth. The resulting Poisson multi-Bernoulli mixture (PMBM) filter is conjugate prior and strongly reduces complexity compared to the equivalent with Bernoulli birth [46]. The general capping, pruning and merging methods from TO-MHT can be applied here as well, and another strong benefit is that Bernoulli components with a low existence probability can be projected onto the Poisson component instead of just deleted, which improves the performance of track initialisation [141] since unassociated measurements will influence the birth intensity at the next time step.

In all of the above-mentioned RFS filters, the tracks of object states are not explicitly represented and one would have to infer this information independently. For many on-line tracking scenarios with frequent object birth and death this is not a problem, since one might only be interested in the current location of the targets [5]. However, cases that call for reconstruction and prediction of the trajectories (like SSA) need a more explicit labelling. For favourable observation conditions (low signal-to-noise ratio (SNR), small sensor noise and few clutter observations), this problem is already largely addressed by conventional approaches like TO-MHT [11][82]. However, for the method to be consistent, it requires that unknown targets are uniformly distributed and can be initialised (*born*) in the filter by some external, heuristic procedure [13][23].

The labelled version of RFS theory (LRFS) solves the track generation problem by explicitly making the label a part of the estimated state representation [84][133][135]. Explicit labelling also gives rise to the notion of joint existence probability, which is very useful from an estimator point of view since it gives more information than the cardinality distribution [134].

Most RFSs can be extended relatively easily to their labelled counterpart by applying the right marking method [134]. One particularly useful LRFS is the LMB [133][135], which is constructed by marking a set of Bernoulli's with a unique label. Note that *unmarking* an LMB only yields an MB if the Bernoulli's are distinct [134]. LRFSs also have an intensity function (PHD), albeit one that is also a function of the label. In the case of the LMB, this labelled PHD provides a complete characterisation of the distribution³³. This property lends itself well for estimators based on the same principles as the MB filter, only without the discussed cardinality bias and with better accuracy performance [104]. It also does not suffer from the spooky effect observed in the CPHD. Although its implementations are more demanding than the (cardinality-balanced) MB [138], the mentioned benefits as well as the fact that it outputs target tracks and performs better in low SNR scenarios make it a very attractive filtering option [96]. Finally, Reuter et al. [104] also introduces a dynamic target grouping method that allows high levels of parallelisation and thus significant runtime reduction.

As it turns out, the LMB is just a special case of the GLMB filter in its δ GLMB form [134]. Vo and Vo [135] and Vo et al. [133] introduced the GLMB RFS family as an analytical, closed-form solution to the MO Bayes' filter for non-linear non-Gaussian dynamic systems. It is characterised by a mixture that accommodates diverse LRFSs such as the LMB and labelled i.i.d.. cluster [134]. Note that δ GLMB refers to the formulation that explicitly separates the various label sets and thereby lends itself better to numerical implementations and approximations [84][96]. GM

³³For this reason, the LMB can be seen as the “mean” (*first moment*) of an LRFS, just as the unlabelled PHD serves this purpose for the unlabelled RFS.

and SMC implementations of the filter exist and show significant improvement in cardinality and state estimation as compared to the corresponding flavours of the CPHD, albeit at much higher computational load [133]. A significant bottleneck occurs in the Bayes correction step, where Murty’s algorithm [91] is used to find the best solutions to the optimal measurement assignment problem.

Many promising variations of the filter have been proposed to reduce the dramatic scaling of components in the posterior, and to improve scalability in specific scenarios [8][36][37].

One thing left undiscussed so far is how to measure the accuracy of an RFS-based MO filter. Since errors can exist in both the estimated cardinality and the estimated states, the usual SO distance metrics such as Euclidean and Mahalanobis distance cannot be easily applied and a unifying metric would be useful to compare performance between implementations. Several classical measures are based on track-to-truth assignment [11, Sec. 13.6], but their definition is ad-hoc and lacks mathematical rigour. The OSPA is a consistent metric proposed by Schuhmacher et al. [114] to compare RFSs, along with a version for labelled sets by Ristic et al. [108]. The metric finds the optimal assignment to match elements from the smaller set to elements from the larger one, and applies a penalty to account for the unmatched elements. Rahmathullah et al. [101] argue that this definition is too focused on the RFS framework, rather than the original MOT problem, and that it does not sufficiently punish false and missed detections. Instead, they suggest the Generalised OSPA (GOSPA), but it is not generally accepted in the literature [95][134].

A.3 MOT for SSA

Applying MOT methods to SSA gives rise to a number of problem-specific challenges, mainly related to the sparsity of measurements, the non-linear orbital dynamics, complications in initialising new objects from UCTs and failed associations because of unannounced manoeuvres by observed satellites.

First of all, the measurements for objects in space surveillance are notoriously sparse. Since most observations are optical angles-only or radar measurements from ground-based observers, an individual observation tracklet might have a duration under a minute, after which it could take many orbital revolutions until the same object is observed again. In the meantime, one must rely on imperfect dynamical models to propagate an uncertain estimated orbit into the future [125], causing wide uncertainty spreads with highly non-linear and often non-Gaussian properties. In particular, small debris objects often have a high area-to-mass ratio, which makes modelling aerodynamic and radiation pressure forces exceedingly challenging [65][90].

The measurement sparsity problem also implies that tracking sensors are scarce and observation time is a valuable resource. Sensor tasking is a whole problem of its own and has solutions that use information theory by maximising the information gain (e.g. [15][49]).

One appealing way to increase the measurement density is by placing an SSA observer in an orbit around the Earth. An observer in LEO will perform a full orbit every 90 minutes and thereby have far more opportunity to collect measurements from objects of interest than any ground-based observer would. In addition, the collected observation geometry has more variation, which is usually good for the convergence and accuracy of estimation algorithms. Contrary to ground-based observers, space-based observers (SBOs) don’t suffer from the influence of atmosphere and weather conditions, nor are there geographical or political restrictions on their position. In

addition, SBOs theoretically have visibility of the complete GEO regime and better observation conditions for small particles in LEO due to their proximity [38].

Another problem that all RFS-based SSA solutions must address is that of target birth. In the context of space objects, this is referred to as IOD and relates to the practice of initialising new components in the set based on well-chosen metrics. Because the state space is in principle unbounded (an optical observation could originate from a satellite or body in any orbital altitude), naive approaches that populate the field of view with Bernoulli components or use a uniform PPP are ineffective or intractable. The fact that target births require a 6-dimensional state initialisation further discourages to randomly populate the search space with birth densities and instead calls for more informed approaches [65].

Because the IOD problem requires information from multiple points across the orbit, it is generally not feasible to obtain an accurate orbit estimate from a single optical measurement. RFS-based filters therefore require the use of a smoothing method to infer IOD from measurements at different epochs [65, p. 1282][132]. Alternatively, approaches based on CAR [88] or PAR [28] have been proposed to represent likely locations for object birth based on tracklets of angular measurements. A mixture density is then initialised over that region so that it can be updated by future measurements. CARs and PARs have been actively used and extended [89][110]. Siminski et al. [116] reformulated IOD for tracklet association with a boundary value problem (BVP) approach using angles and angular rate. The most probable association hypotheses then result in deterministic orbit solutions, meaning that the method does not inherently estimate the uncertainty distribution. This approach has been used as input for adaptive MO birth processes [105] in the LMB filter implementations [14][16] with very promising reported effects. The approach is more efficient since it avoids the need for large Gaussian mixtures and the smooth topography of the BVP formulation allows for quick optimisation [16, p. 840]. Gehly et al. [48] use Gooding IOD [52] to solve the BVP, but several other methods are available. A variety of other IOD approaches exist that haven't been fully integrated with MOT, for example based on differential algebra [4] or machine learning [75].

Note that, apart from modelling object birth, there are equivalent challenges to represent survival and detection probabilities [42], presence of clutter as well as spawning objects [65].

A final major point of interest and active field of research in SSA is that of anomaly and manoeuvre detection. In the field of MOT, this means overcoming all the association challenges related to unknown changes in the dynamics of an RSO. Note that model errors can also cause propagation offsets, so it is very challenging to distinguish those from manoeuvres [65]. Approaches have been suggested based on minimum control effort between UCTs [57] or optimal control problem formulations [77]–[79]. These have reported promising single-object results as compared to process-noise based Kalman filter approaches and have also been applied with in various RFS filters (e.g [35]). However, it remains very challenging to separate model uncertainties from object manoeuvres.

APPENDIX B

PROJECT PLAN

This appendix contains part of the original project plan from December 2024.

A high-level preliminary work breakdown is shown in Figure B.1 to illustrate top-level flow of the thesis efforts. Note that there is some chronological dependency between the phases, but I propose a slightly iterative approach that aims to address the entire problem before increasing the complexity in each of the elements.

For example, a first iteration could be to implement the LMB filter and test it on the case of tracklet formation case (this nearly linear application lends itself well to test the filter in an uncomplicated way). If this is successful, the LMB can be applied to small number of RSOs with simplified dynamics and a naive birth model and some more challenging scenarios can be set up for validation. This can constitute a first iteration as “proof of concept”, after which each of the phases in Figure B.1 can be incrementally expanded to reach more realistic scenarios, more filters to compare and filter adjustments that can reach competitive performance.

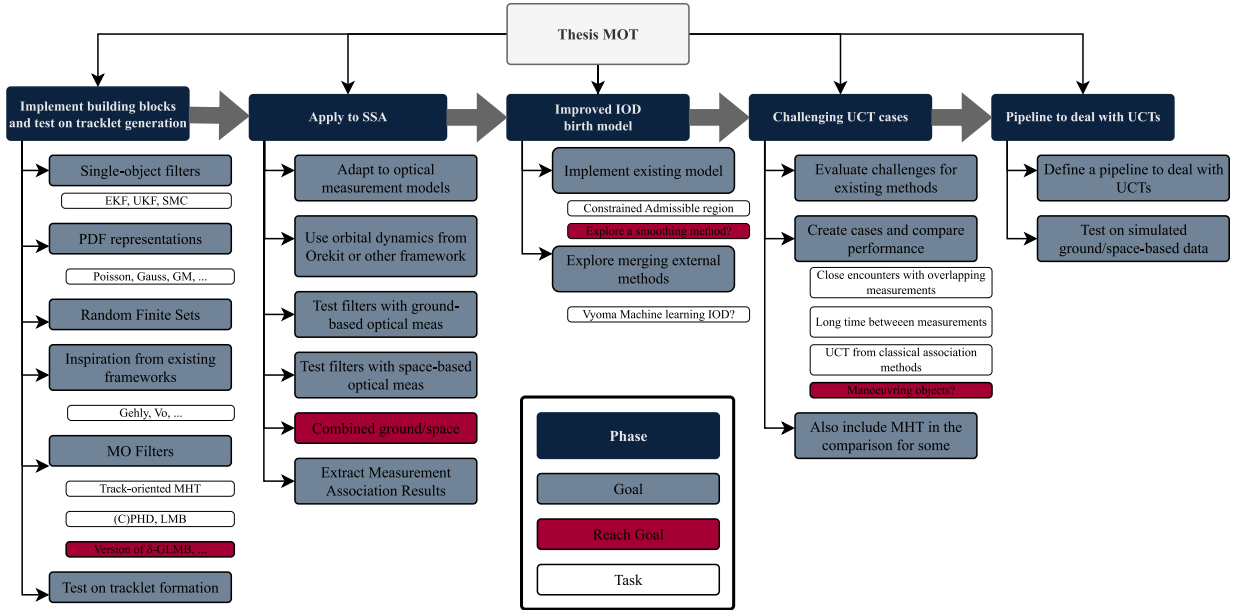


Figure B.1: Work Breakdown Structure

In the interest of planning the main thesis activities over time, a high-level schedule is created, dividing time over admin, reading, writing and coding. This is purposefully kept very general, so that it can serve as a guideline to be refined as the project moves along. Please refer to Figure B.2.




Figure B.2: Preliminary activity time planning

TESTING THE LMB IN TWO DIMENSIONS

To validate the implementation of the conventional LMB filter in BASIL, it is more straightforward to first address a problem of lower complexity. Using a simulation of objects in two dimensions with relatively simple dynamics, the results can be more intuitively interpreted and compared to reference outputs from the literature.

C.1 Reuter and Vo LMB EKF

To validate the high-level working functionality of the LMB filter implementation, the demonstration test case on LMB-EKF/UKF by Reuter et al. [104][136], is reproduced as closely as possible.

A. Setup

The target state is $\mathbf{x} = (x, y, \dot{x}, \dot{y}, \omega)^T$, which represents a two-dimensional position and velocity, along with a turn rate ω . The linearised prediction model then has state transition matrix around ω_k

$$\Phi_{k,k-1} = \begin{pmatrix} 1 & \frac{\sin(\omega_k)}{\omega_k} & 0 & -\frac{1-\cos(\omega_k)}{\omega_k} & 0 \\ 0 & \cos(\omega_k) & 0 & -\sin(\omega_k) & 0 \\ 1 & \frac{1-\cos(\omega_k)}{\omega_k} & 0 & \frac{\sin(\omega_k)}{\omega_k} & 0 \\ 0 & \sin(\omega_k) & 0 & \cos(\omega_k) & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad (\text{C.1})$$

and a discrete-time process noise model is defined as

$$Q_k = \Gamma_{k,k-1} \begin{pmatrix} \sigma_w^2 & 0 & 0 \\ 0 & \sigma_w^2 & 0 \\ 0 & 0 & \sigma_v^2 \end{pmatrix} \Gamma_{k,k-1}^T, \quad \text{with} \quad \Gamma_{k,k-1} = \begin{pmatrix} \frac{1}{2} & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \text{and} \quad \begin{cases} \sigma_w = 15 \text{ m/s}^2 \\ \sigma_v = \frac{\pi}{180} \text{ rad/s} \end{cases}. \quad (\text{C.2})$$

The measurements are $\mathbf{z} = (\theta, \rho)^T$, where $\theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ is the angle to the positive y axis and $\rho \in [0, 2000]$ is the distance to the origin. Clutter is modelled as a PPP with clutter rate λ_c and uniform spacial density over the observable measurement domain. Detection and survival probabilities are constant and listed in Table C.1.

Table C.1: LMB parameters for 2D test case

Parameter	Value
P_S	0.99
P_D	0.98
λ_c	10
ϑ_P^{LMB}	1×10^{-5}

Table C.2: Constant birth model parameters for 2D test case

i	$r_B^{(i)}$	x [m]	y [m/s]	\dot{x} [m/s]	\dot{y} [m/s]	ω [rad/s]
1	0.02	-1500	0	250	0	0
2	0.02	-250	0	1000	0	0
3	0.03	250	0	750	0	0
4	0.03	1000	0	1500	0	0

Finally, the case uses a constant multi-Bernoulli birth model with four Gaussian components such that $\pi_{B,k} = \left\{ \left(r^{(i)}, p_B^{(i)} \right) \right\}_{i=1}^4$ with $p_B^{(i)} = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_B^{(i)}, \mathbf{P}_B)$. The distribution is further defined by a constant covariance matrix $\mathbf{P}_B = \text{diag}[(50, 50, 50, 50, 6\pi/180)^T]$ and the parameters in Table C.2.

A set of objects is simulated along with measurements following the described models and all the information is visualised in Figure C.1. Note that the filter does not have any way to distinguish between real measurements and false positives from the input itself.




Figure C.1: 2D test case ground truth and measurements (Based on [136])

B. Results

The LMB filter behaves as expected and is able to detect the birth of all simulated targets. Figure C.2 shows that all 10 objects are found by the filter and the estimates accurately approach the ground truth.




Figure C.2: 2D test case LMB estimates

The cardinality estimate shows that there is some delay to confirm newborn targets and remove those that no longer exist (see Figure C.3), which is reflected as the temporary spikes in OSPA

distance as shown in Figure C.4. Again, this is expected in MO filtering framework, to avoid that clutter measurements or missed detections too easily affect the filtering estimates.




Figure C.3: 2D test case LMB cardinality




Figure C.4: 2D test case LMB OSPA metric

It should be noted that the cardinality is at times underestimated incorrectly, even when no objects are dying or being born. The likely reason is that two missed detections occurred in immediate succession and the existence probability therefore dropped significantly for one of the targets. Because the cardinality estimate \hat{n} is obtained as the MAP of its distribution (Section 3.3.4-F), and only \hat{n} are then estimated, the filter might be slightly too sensitive to these types of missed detection incidents.

All these observations are consistent with the results in the RFS-MOT toolbox from Vo and Vo [136] and the details reported by Reuter et al. [104]. For that reason, the basic verification of the core LMB algorithms is considered successful.

APPENDIX D

CONSTRAINED ADMISSIBLE REGION IOD

This appendix gives the relevant details to derive the constrained admissible region (CAR) approach as proposed by [Milani et al. \[88\]](#) and discussed in [Section 5.3.2–A](#). The development follows that of [DeMars and Jah \[28\]](#), as repeated by e.g. [Gehly \[47\]\[49\]](#) and [Cament et al. \[17\]](#).

Using a tracklet of topocentric right ascension/declination measurements $(\alpha, \delta)^T$ or some other measurement source for an RSO, the *attributable* observation is obtained of the form

$$\mathbf{a} = (\alpha \ \delta \ \dot{\alpha} \ \dot{\delta}) \in [-\pi, \pi) \times \left(-\frac{\pi}{2}, \frac{\pi}{2}\right) \times \mathbb{R}^2. \quad (\text{D.1})$$

Denote the ECI 6-dimensional Cartesian object state as $\mathbf{x} = (\mathbf{r}^T, \dot{\mathbf{r}}^T)^T$, where \mathbf{r} is its position vector and $\dot{\mathbf{r}} = \mathbf{v}$ is the velocity vector. The two-body *specific orbital energy* is then given by

$$\mathcal{E} = \frac{\|\dot{\mathbf{r}}\|^2}{2} - \frac{\mu}{\|\mathbf{r}\|} = -\frac{\mu}{2a}, \quad (\text{D.2})$$

where a is the orbit's semi-major axis. Let an observer with state \mathbf{x}_o observe the object such that $\mathbf{x}_\rho = (\boldsymbol{\rho}^T, \dot{\boldsymbol{\rho}}^T)^T$ is the state vector relative to the observer, i.e. $\mathbf{x} = \mathbf{x}_o + \mathbf{x}_\rho$. Let this relative state vector be expressed in spherical coordinates,

$$\boldsymbol{\rho} = \rho \mathbf{u}_\rho, \quad \text{and} \quad \dot{\boldsymbol{\rho}} = \dot{\rho} \mathbf{u}_\rho + \rho \dot{\alpha} \mathbf{u}_\alpha + \rho \dot{\delta} \mathbf{u}_\delta, \quad (\text{D.3})$$

where the unit vectors are defined by

$$\mathbf{u}_\rho = \begin{pmatrix} \cos(\alpha) \cos(\delta) \\ \sin(\alpha) \cos(\delta) \\ \sin(\delta) \end{pmatrix}, \quad \mathbf{u}_\alpha = \begin{pmatrix} -\sin(\alpha) \cos(\delta) \\ \cos(\alpha) \cos(\delta) \\ 0 \end{pmatrix}, \quad (\text{D.4a})$$

$$\text{and} \quad \mathbf{u}_\delta = \begin{pmatrix} -\cos(\alpha) \sin(\delta) \\ -\sin(\alpha) \sin(\delta) \\ \cos(\delta) \end{pmatrix}. \quad (\text{D.4b})$$

Then, define the scalars

$$w_0 = \|\mathbf{r}_o\|^2, \quad w_1 = 2\langle \dot{\mathbf{r}}_o, \mathbf{u}_\rho \rangle, \quad w_2 = \dot{\alpha}^2 \cos^2(\delta) + \dot{\delta}^2, \quad (\text{D.5a})$$

$$w_3 = 2\dot{\alpha} \langle \dot{\mathbf{r}}_o, \mathbf{u}_\delta \rangle + 2\dot{\delta} \langle \dot{\mathbf{r}}_o, \mathbf{u}_\alpha \rangle, \quad w_4 = \|\dot{\mathbf{r}}_o\|^2, \quad (\text{D.5b})$$

$$\text{and} \quad w_5 = 2\langle \mathbf{r}_o, \mathbf{u}_\rho \rangle. \quad (\text{D.5c})$$

The squared norm of the observed object's position and velocity vector in ECI are then

$$\|\mathbf{r}\|^2 = \rho^2 + w_5 \rho + w_0, \quad (\text{D.6a})$$

$$\|\dot{\mathbf{r}}\|^2 = \dot{\rho}^2 + w_1 \dot{\rho} + w_2 \rho^2 + w_3 \rho + w_4. \quad (\text{D.6b})$$

These expressions can be substituted into [Eq. \(D.2\)](#) to obtain

$$\dot{\rho}^2 + w_1 \dot{\rho} + F(\rho) - 2\mathcal{E} = 0, \quad (\text{D.7})$$

$$\text{where } F(\rho) = w_2\rho^2 + w_3\rho + w_4 - \frac{2\mu}{\sqrt{\rho^2 + w_5\rho + w_0}}. \quad (\text{D.8})$$

Solving Eq. (D.7), an expression can be obtained for $\dot{\rho}$ as a function of ρ :

$$\dot{\rho} = -\frac{w_1}{2} \pm \sqrt{\left(\frac{w_1}{2}\right)^2 - F(\rho) + 2\mathcal{E}}. \quad (\text{D.9})$$

If one now requires that the observed object is in orbit around Earth, the region of *admissible* combinations of ρ and $\dot{\rho}$ is bounded by setting $\mathcal{E} < 0$.

D.1 Constraining the admissible region

Instead of only requiring the orbital energy to be negative, it can be useful in many cases to further constrain the allowable orbits for IOD. This is referred to as a constrained admissible region (CAR).

A. Semi-major axis

Define an admissible range of semi-major axis $a \in [a_{\min}, a_{\max}]$. This translates directly to an energy constraint by Eq. (D.2), so two bounding curves can be constructed to constrain the semi-major axis.

B. Eccentricity

Further, it is often desirable to limit the search to (near-circular) orbits or just put a limit on how eccentric the target orbits can be. To achieve such an eccentricity constraint, first recall the definition of specific angular momentum:

$$\mathbf{h} = \mathbf{r} \times \dot{\mathbf{r}}. \quad (\text{D.10})$$

It can further be shown that

$$\mathbf{h} = \dot{\rho}\mathbf{h}_1 + \rho^2\mathbf{h}_2 + \rho\mathbf{h}_3 + \mathbf{h}_4, \quad (\text{D.11a})$$

$$\text{and } \|\mathbf{h}\|^2 = c_0\dot{\rho}^2 + P(\rho)\dot{\rho} + U(\rho), \quad (\text{D.11b})$$

based on the definition of four vector parameters

$$\mathbf{h}_1 = \mathbf{r}_o \times \mathbf{u}_\rho, \quad \mathbf{h}_2 = \mathbf{u}_\rho \times (\dot{\alpha}\mathbf{u}_\alpha + \dot{\delta}\mathbf{u}_\delta), \quad (\text{D.12a})$$

$$\mathbf{h}_4 = \mathbf{r}_o \times \dot{\mathbf{r}}_o, \text{ and } \mathbf{h}_3 = \mathbf{u}_\rho \times \dot{\mathbf{r}}_o + \mathbf{r}_o \times (\dot{\alpha}\mathbf{u}_\alpha + \dot{\delta}\mathbf{u}_\delta), \quad (\text{D.12b})$$

nine scalar parameters

$$c_0 = \|\mathbf{h}_1\|^2, \quad c_1 = 2\langle \mathbf{h}_1, \mathbf{h}_2 \rangle, \quad c_2 = 2\langle \mathbf{h}_1, \mathbf{h}_3 \rangle, \quad (\text{D.13a})$$

$$c_3 = 2\langle \mathbf{h}_1, \mathbf{h}_4 \rangle, \quad c_4 = \|\mathbf{h}_2\|^2, \quad c_5 = 2\langle \mathbf{h}_2, \mathbf{h}_3 \rangle, \quad (\text{D.13b})$$

$$c_6 = 2\langle \mathbf{h}_2, \mathbf{h}_4 \rangle + \|\mathbf{h}_3\|^2, \quad c_7 = 2\langle \mathbf{h}_3, \mathbf{h}_4 \rangle, \text{ and } c_8 = \|\mathbf{h}_4\|^2, \quad (\text{D.13c})$$

and two scalar functions

$$P(\rho) = c_1\rho^2 + c_2\rho + c_3, \quad (\text{D.14a})$$

$$U(\rho) = c_4\rho^4 + c_5\rho^3 + c_6\rho^2 + c_7\rho + c_8. \quad (\text{D.14b})$$

The relation between eccentricity, angular momentum and specific orbital energy can be formulated as [125, p. 97]

$$2\mathcal{E}\|\mathbf{h}\|^2 = -\mu^2(1 - e^2). \quad (\text{D.15})$$

which can be expanded, by substitution of Eq. (D.7) for $2\mathcal{E}$ and Eq. (D.11) for $\|\mathbf{h}\|^2$, to

$$[\dot{\rho}^2 + w_1\dot{\rho} + F(\rho)][c_0\dot{\rho}^2 + P(\rho)\dot{\rho} + U(\rho)] = -\mu^2(1 - e^2) \quad (\text{D.16})$$

and otherwise rewritten as a fourth order polynomial equation in $\dot{\rho}$

$$a_4\dot{\rho}^4 + a_3\dot{\rho}^3 + a_2\dot{\rho}^2 + a_1\dot{\rho} + a_0 = 0, \quad (\text{D.17})$$

where

$$a_4 = c_0, \quad a_3 = P(\rho) + c_0w_1, \quad a_2 = U(\rho) + c_0F(\rho) + w_1P(\rho), \quad (\text{D.18a})$$

$$a_1 = F(\rho)P(\rho) + w_1U(\rho), \text{ and } a_0 = F(\rho)U(\rho) + \mu^2(1 - e^2). \quad (\text{D.18b})$$

The constraint $e < e_{\max}$ may be enforced by finding the real-valued roots of Eq. (D.7) for all values of ρ with $e = e_{\max}$.

C. Approximation by Gaussian mixture

The procedure to convert the uniform admissible region in $(\rho, \dot{\rho})$ to a Gaussian mixture in the Cartesian orbit space was already summarised in Section 5.3.2-B and more extensive details can be found in the paper by DeMars and Jah [28].

D.2 Attributable prediction

Given the observer state \mathbf{x}_o and a target state \mathbf{x} (guess/estimate), one can compute the expected attributable. Note that the slant range vector $\boldsymbol{\rho}$ and its derivative $\dot{\boldsymbol{\rho}}$ are simply the difference between the target and the observer state

$$\begin{pmatrix} \boldsymbol{\rho} \\ \dot{\boldsymbol{\rho}} \end{pmatrix} = \mathbf{x} - \mathbf{x}_o. \quad (\text{D.19})$$

Define the basis vectors of the topocentric reference frame parallel to ECI and centred at the observer as

$$(\mathbf{u}_i \ \mathbf{u}_j \ \mathbf{u}_k) = \mathbf{I}_{3 \times 3} \quad (\text{D.20})$$

Then, compute the slant range and its direction by

$$\rho = \|\boldsymbol{\rho}\|, \quad \text{and} \quad \mathbf{u}_\rho = \frac{\boldsymbol{\rho}}{\rho}. \quad (\text{D.21})$$

Further solve the unit vectors in Eq. (D.4) to find α and δ

$$\delta = \arcsin(\langle \mathbf{u}_\rho, \mathbf{u}_k \rangle), \quad \cos(\delta) = \sqrt{1 - \langle \mathbf{u}_\rho, \mathbf{u}_k \rangle^2} \quad (\text{D.22a})$$

$$\alpha = \text{atan2}(\langle \mathbf{u}_\rho, \mathbf{u}_j \rangle, \langle \mathbf{u}_\rho, \mathbf{u}_i \rangle) \quad (\text{D.22b})$$

where $\text{atan2}(y, x)$ is defined as $\arctan(\frac{y}{x})$, but adjusted to be in the correct quadrant such that the result is in $[-\pi, \pi)$:

$$\text{atan2}(y, x) = \begin{cases} \arctan\left(\frac{y}{x}\right) & \text{if } x > 0 \\ \arctan\left(\frac{y}{x}\right) + \pi & \text{if } x < 0 \wedge y \geq 0 \\ \arctan\left(\frac{y}{x}\right) - \pi & \text{if } x < 0 \wedge y < 0 \\ +\frac{\pi}{2} & \text{if } x = 0 \wedge y > 0 \\ -\frac{\pi}{2} & \text{if } x = 0 \wedge y < 0 \\ \text{undefined} & \text{if } x = 0 \wedge y = 0 \end{cases} \quad (\text{D.23})$$

Finally, use [Eq. \(D.3\)](#) to solve for the angular rates

$$\dot{\rho} = \|\dot{\boldsymbol{\rho}}\|, \quad \dot{\alpha} = \frac{1}{\rho} \langle \dot{\boldsymbol{\rho}}, \mathbf{u}_\alpha \rangle, \quad \text{and} \quad \dot{\delta} = \frac{1}{\rho} \langle \dot{\boldsymbol{\rho}}, \mathbf{u}_\delta \rangle. \quad (\text{D.24})$$

D.3 CAR Examples

For the purpose of validating the implementation, it is useful to reproduce several examples of CARs from the literature. Additionally, an example using space-based observations of a GEO object is briefly introduced for comparison. A LEO region with space-based observer is addressed in [Section 7.4](#), related to one of the simulation cases.

D.3.1 Ground-based observer

A. Reproducing DeMars result

The scenario from [DeMars and Jah \[28\]](#) was used as example in [Section 5.3.2–A](#) and the relevant illustrations and plots can also be found there. By visual inspection, it is confirmed that [Figure 5.9](#) and [Figure 5.14](#) match the expected region and approximation from the original paper. One discrepancy is that the design values of σ_ρ and $\sigma_{\dot{\rho}}$ are much higher in this case, leading to fewer GM components. It was deliberately chosen to limit the number of components to 15x15 for the present work, to limit computation time. However allowing increasingly dense approximations would require only minor modifications in the implementation.

B. Reproducing Gehly result

For further validation purposes, the second example follows the settings from a unit test in the METIS library [\[47\]](#), which was also discussed by [Gehly \[49\]](#). The observer is located on the Earth’s surface at 30° latitude, and the attributable observation is $\mathbf{a} = (10 \text{ deg}, -2 \text{ deg}, 15 \text{ deg/h}, 3 \text{ deg/h})$. The $(\rho, \dot{\rho})$ space is constrained in semi-major axis $a \in [42565, 41764] \text{ km}$ and in eccentricity $e < 0.1$.

[Figure D.1](#) shows the CAR as computed using the BASIL implementation, which closely resembles the results from METIS in [Figure D.2](#). The corresponding GM approximation of the uniform region in $(\rho, \dot{\rho})$ is plotted in [Figure D.4](#) and [Figure D.6](#) for BASIL and METIS, respectively. The number of components and locations of the means are indistinguishable between both implementations.

The only noticeable difference is in the range-marginal PDF. METIS does not handle the concavity of the CAR correctly, resulting in a discontinuity in the range-marginal PDF ([Figure D.5](#)). The BASIL solution in [Figure D.3](#) does not have this issue and corresponds well with expectations.

In summary, the results from both ground-based test cases match the literature and reference code, so the verification of BASIL’s CAR implementation is considered successful.




Figure D.1: Constrained admissible region with ground-based observer for Gehly test case, produced with BASIL




Figure D.2: GM approximation for Gehly test case, produced with METIS [47]




Figure D.3: Range-marginal PDF for Gehly test case, produced with BASIL




Figure D.4: GM approximation for the Gehly test case produced with BASIL




Figure D.5: Range-marginal PDF for Gehly test case, produced with METIS [47]




Figure D.6: GM approximation for Gehly test case, produced with METIS [47]

D.3.2 Space-based observer

For comparison, [Figure D.7](#) shows a CAR based on a space-based observer in a circular LEO orbit. There are now two disjoint eccentricity regions, where one includes the trivial solution $(0, 0)$, representing the observer orbit. Space-based CARs are further discussed in [Section 7.4](#).




Figure D.7: Example constrained admissible region with space-based observer and near-circular, near-GEO constraints. The observer is at 500 km altitude and the attributable observation is $\mathbf{a} = (74.6 \text{ deg}, 0.0 \text{ deg}, -13.2 \text{ deg/h}, 0.0 \text{ deg/h})$.




Figure D.8: GM approximation with $\sigma_\rho = 300 \text{ km}$ and $\sigma_{\dot{\rho}} = 80 \text{ m/s}$.




Figure D.9: Range-marginal probability density function and GM for $\sigma_\rho = 300 \text{ km}$.

AUXILIARY COMPUTATIONS

E.1 Covariance ellipsoids and the χ^2 distribution

The uncertainty spread in a multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \mathbf{P})$ is often represented by a so-called *covariance ellipsoid*. This is a locus of limited Mahalanobis distance $d_M(\cdot; \boldsymbol{\mu}, \mathbf{P}) < d_M^{\max}$ around the mean $\boldsymbol{\mu}$, and represents a confidence interval for the given distribution. Since the squared Mahalanobis distance d_M^2 follows a χ^2 distribution, the confidence level for a particular d_M ellipsoid depends on the number of dimensions it represents. Table E.1 summarises the confidence levels for various choices of d_M^{\max} .

Note that the $d_M^{\max} = 1$ ellipsoid in 1 dimension corresponds to the $\pm 1\sigma$ uncertainty bound, representing a 68% confidence interval. However, in two and three dimensions, $d_M^{\max} = 1$ bounds only 39% and 20%, respectively. Covariance ellipsoids in higher dimensions are not easily visualised and not really used, but they are relevant in MOT for the purpose of e.g. GM merging (Section 2.1.2–C) and measurement gating (Section 3.3.3–B).

Table E.1: Cumulative densities of the χ distribution. This is the probability that the Mahalanobis distance is smaller than d_M^{\max}

Dimensions	$d_M^{\max} = 1$	$d_M^{\max} = 2$	$d_M^{\max} = 3$	$d_M^{\max} = 4$	$d_M^{\max} = 5$	$d_M^{\max} = 6$
1	0.683	0.954	0.997	0.999937	0.999999	1.000000
2	0.393	0.865	0.989	0.999665	0.999996	1.000000
3	0.199	0.739	0.971	0.998866	0.999985	1.000000
4	0.090	0.594	0.939	0.996981	0.999950	1.000000
5	0.037	0.451	0.891	0.993156	0.999861	0.999999
6	0.014	0.323	0.826	0.986246	0.999659	0.999997

E.2 LogSumExp normalisation

In MOT, there are many cases where a set of weights need to be normalised (e.g. after GM reduction or in managing GLMB hypotheses). Since it is very common that some weights are many orders of magnitude smaller than others, simply dividing all weights by the sum can lead to underflow or overflow complications. These can be avoided by solving the normalisation in log-space [97].

Let $\mathbf{x} \in \mathbb{R}^N$ be a vector of weights to be normalised and $\mathbf{1} \triangleq \mathbf{1}_N \in \mathbb{R}^N$ the vector of ones. Then

$$\tilde{\mathbf{x}} = \frac{\mathbf{x}}{\mathbf{1}^T \mathbf{x}} \Rightarrow \mathbf{x} = \log(\tilde{\mathbf{x}}) + \log(\mathbf{1}^T \mathbf{x}) \Rightarrow \tilde{\mathbf{x}} = \exp \left(\boldsymbol{\eta} - \underbrace{\log[\mathbf{1}^T \exp(\boldsymbol{\eta})]}_{\text{LogSumExp}(\boldsymbol{\eta})} \right), \quad (\text{E.1})$$

where $\boldsymbol{\eta} = \log(\mathbf{x})$ and $\text{LogSumExp}(\boldsymbol{\eta})$ is computed using this stable formulation:

$$\exp(\boldsymbol{\eta}) = \exp(\eta_{\max}) \exp(\boldsymbol{\eta} - \eta_{\max} \mathbf{1}) \Rightarrow \mathbf{1}^T \exp(\boldsymbol{\eta}) = \exp(\eta_{\max}) [\mathbf{1}^T \exp(\boldsymbol{\eta} - \eta_{\max} \mathbf{1})] \quad (\text{E.2a})$$

$$\Rightarrow \text{LogSumExp}(\boldsymbol{\eta}) = \eta_{\max} + \text{LogSumExp}(\boldsymbol{\eta} - \eta_{\max} \mathbf{1}) \quad (\text{E.2b})$$

APPENDIX F

ASSOCIATION MATRICES

This appendix contains the association probabilities for some of the cases in [Chapter 7](#), as computed from the GLMB update. Every row represents a tracklet and the columns correspond to all targets in the filter with at least 1 associated tracklet. Since all these results use the tracklet LMB filter from [Chapter 6](#), no accumulation of the probabilities is required and the values in these table are the *time-marginal association probabilities* from [Eq. \(4.2\)](#).

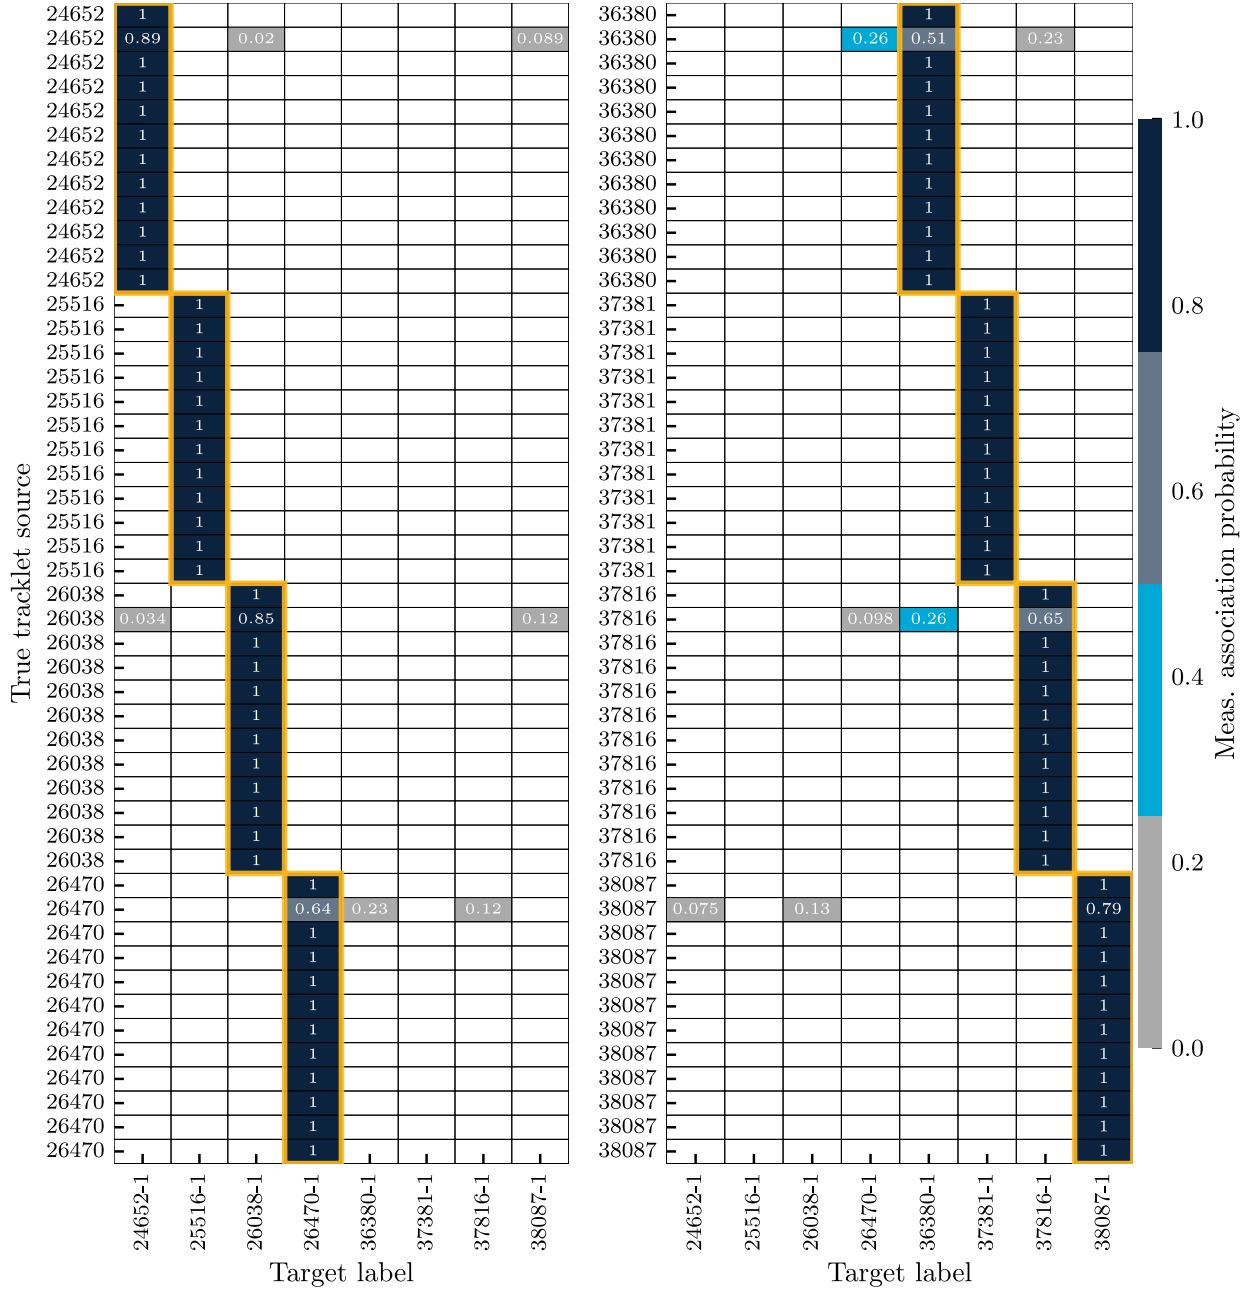


Figure F.1: Measurement association probabilities for GEO objects and ground-based observer without initial information (see [Section 7.1](#))

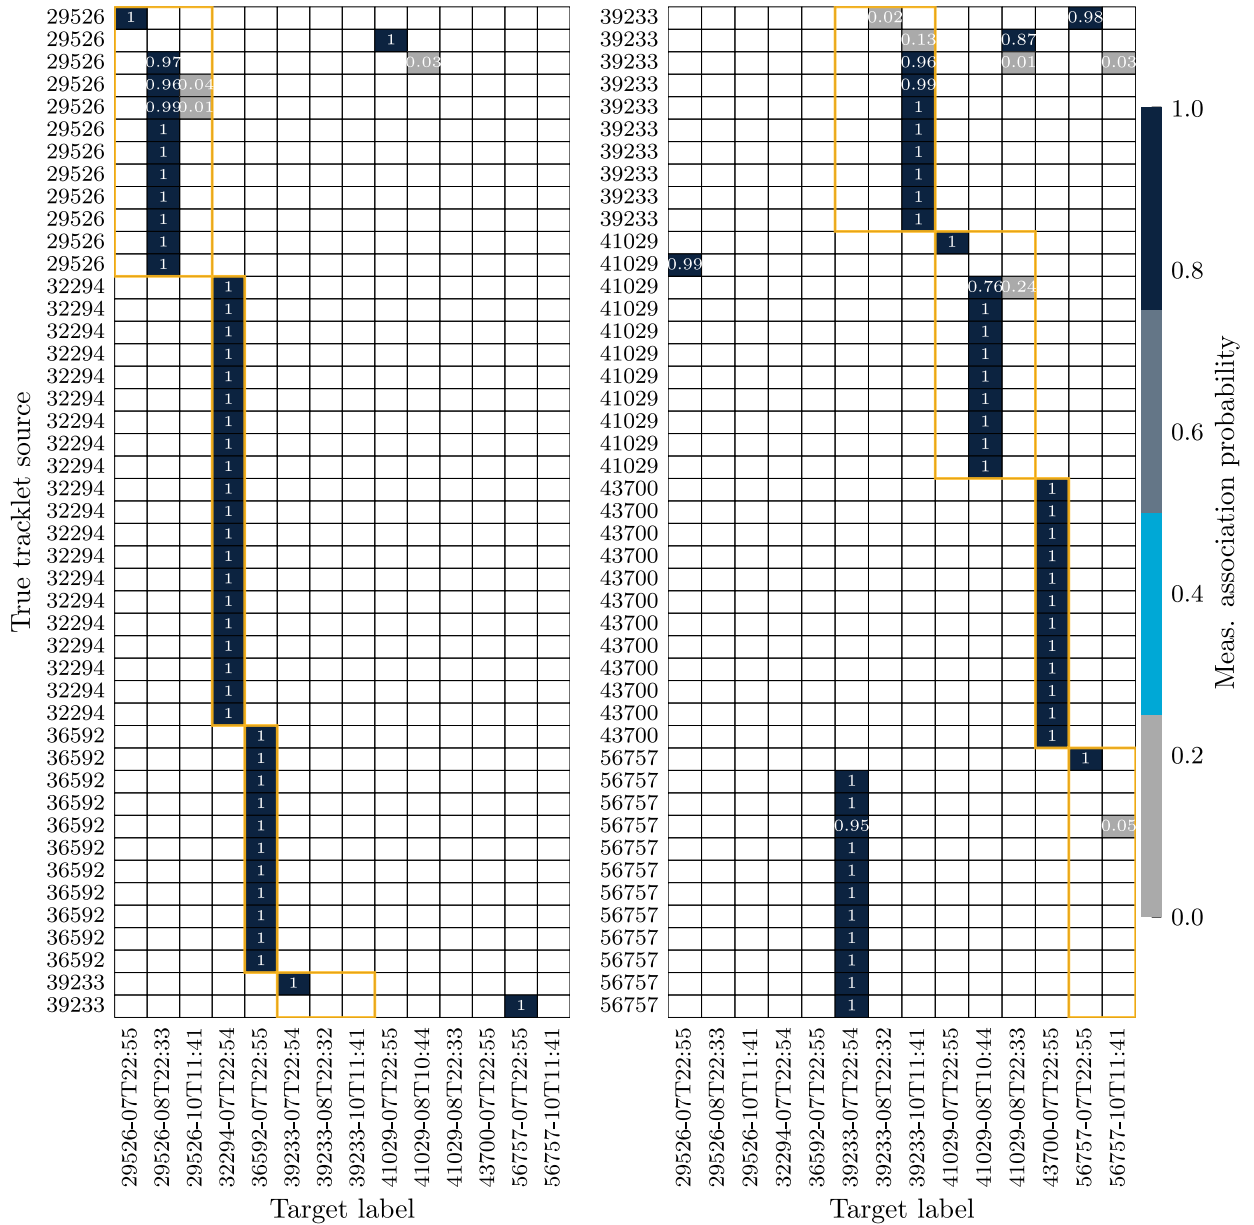



Figure F.2: Measurement association probabilities for close-spaced GEO objects and space-based observer without initial information (Section 7.2.4). Here, Gaussian mixture were capped at a maximum of 10 components, far too few in comparison to the CAR birth mixtures.






Figure F.4: measurement association probabilities for payload separation with space-based observer. Here, the CAR birth model was omitted and the only birth components were single Gaussians from the rocket upper stage. The association matrix is truncated after 10 passes (first three days); all remaining tracklets had 100% association probability with the correct target.

APPENDIX G

SIMULATED OBJECT DATA

Table G.1: Orbital elements for test case objects at epoch 2016-01-14T12:00:00 UTC, based on [Pirovano \[98, Appendix A\]](#)

Norad	a [km]	e [-]	i [deg]	Ω [deg]	ω [deg]	θ [deg]
26470	42165.70	0.00058	0.714	205.356	84.527	330.293
36380	42165.97	0.00031	0.085	179.045	111.543	329.532
37381	43674.63	0.02663	7.002	222.357	42.568	43.194
37816	42166.11	0.00047	0.142	222.187	58.793	338.872
38087	42166.49	0.00024	0.112	211.319	73.127	346.808
26038	42106.22	0.00059	0.128	265.581	91.955	272.955
24652	42407.16	0.00062	2.806	199.149	73.446	359.985
25516	42430.15	0.00040	4.538	216.120	63.980	348.544

Table G.2: TLE data for closely spaced GEO objects, obtained from SpaceTrack

Name	Norad	TLE
Arabsat 4B	29526	1 29526U 06051A 25159.72868902 .00000164 00000-0 00000-0 0 9991
		2 29526 1.0405 85.0357 0003686 330.7572 130.1283 1.00270415 47520
Arabsat 6B	41029	1 41029U 15065B 25159.72860213 .00000164 00000-0 00000-0 0 9995
		2 41029 0.0374 235.5605 0002925 329.8417 340.2334 1.00270611 35027
Arabsat 5B	36592	1 36592U 10025A 25159.72860213 .00000164 00000-0 00000-0 0 9995
		2 36592 0.0591 122.6279 0004899 347.5433 75.4588 1.00273063 49473
Arabsat 7B	56757	1 56757U 23075A 25159.27405853 .00000163 00000-0 00000-0 0 9993
		2 56757 0.0772 28.6718 0004707 356.2770 356.6138 1.00271624 7370
ES'HAIL 2	43700	1 43700U 18090A 25159.72860213 .00000164 00000-0 00000-0 0 9994
		2 43700 0.0212 59.6695 0002403 51.1694 74.5997 1.00273860 23931
ES'HAIL 1	39233	1 39233U 13044A 25159.72851525 .00000162 00000-0 00000-0 0 9998
		2 39233 0.0481 354.7340 0002473 82.6728 107.6878 1.00271738 42888
SKYNET 5B	32294	1 32294U 07056B 25159.51422862 .00000157 00000-0 00000-0 0 9993
		2 32294 4.3017 77.9940 0003726 351.0512 38.3310 1.00268975 64453

Table G.3: Orbital elements for payload separation test case objects at epoch 2025-05-25T20:00:00 UTC,

ID	a [km]	e [-]	i [deg]	Ω [deg]	ω [deg]	θ [deg]
0000	7170.30	0.00115	0.028	-7.606	-184.370	176.765
0001	7172.23	0.00088	0.028	-7.606	-183.387	175.782
0002	7174.15	0.00062	0.028	-7.606	-181.550	173.945
0003	7176.07	0.00035	0.028	-7.606	-176.914	169.308
0004	7178.00	0.00010	0.028	-7.606	-147.042	139.437
0005	7179.93	0.00020	0.028	-7.606	-26.281	18.676
0006	7181.86	0.00047	0.028	-7.606	-15.641	8.036
0007	7183.78	0.00073	0.028	-7.606	-12.705	5.099
0008	7185.72	0.00100	0.028	-7.606	-11.338	3.733
0009	7187.65	0.00127	0.028	-7.606	-10.549	2.943