Bayes Risk-Based Radar Resource Management

for Joint Multi-Target Track Maintenance and Classification

Joost Henricus Frederick Jaspers

Delft University of Technology

Bayes Risk-Based Radar Resource Management

for Joint Multi-Target Track Maintenance and Classification

by

Joost Henricus Frederick Jaspers

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on the 7th of July, 2026 at 13:00.

Student number:	5372925		
Project duration:	September 8, 2025 – July 7, 2026		
Thesis committee:	Prof. DSc. O. Yarovyi,	TU Delft MS3,	Chair
	dr. ir. R. Heusdens,	TU Delft SPS,	External committee member
	dr. ir. J.N. Driessen,	TU Delft MS3,	Daily and responsible supervisor
Committee advisor:	S. Chioccarello MSc,	TU Delft MS3,	Daily co-supervisor

Style: TU Delft Report Style, with modifications by Daan Zwaneveld
Cover Page made with AI

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft

Preface

This thesis marks the completion of my Master's degree at Delft University of Technology. The research was conducted at the Microwave Sensing, Signals and Systems group (MS3) and focused on radar resource management. More specifically, the aim of my project was to develop a new resource allocation approach that accounts for both target tracking and classification, while allowing the operational consequences of different radar decisions to be incorporated into the allocation process.

I would first like to thank Professor Olexander Yarovyi for connecting me with Hans Driessen. This connection introduced me to radar resource management and led me to a thesis subject that closely matched my interests.

I am especially grateful to my supervisor, Hans Driessen, and my co-supervisor, Stefano Chioccarello, for their guidance, flexibility, and the many thoughtful discussions during the project. During the first part of my thesis, I was also completing an exchange semester in Lausanne and could therefore only work on the project parttime. Their willingness to accommodate this situation allowed me to combine both experiences and still work towards graduating on time. I greatly appreciated their continued support and availability, including the occasional thesis discussion during one of their holidays.

Throughout this project, I did not only learn about radar systems, Bayesian decision-making, and resource management, but also about the process of conducting research itself. I learned how important it is to translate complicated ideas into explanations that are understandable to others. Attempting to formulate the methodology clearly often helped me develop a better understanding of the underlying concepts myself.

I would also like to thank my friends, my girlfriend, and my family for their support throughout my studies and this thesis project. I am especially grateful to my parents, who continuously motivated me to keep going and challenged me to reconsider the boundaries I placed on myself. Their support helped me believe that more was possible than I sometimes thought and, most importantly, that the sky is not the limit.

During my project, I used ChatGPT 5.5 by OpenAI to support me with grammar correction, language editing, textual refinement, and coding assistance for material originally written or developed by me. This coding assistance was limited to code structuring, debugging, and implementation suggestions. No AI-generated content was used as a source of scientific results, and all final text, equations, code, figures, interpretations, and conclusions were reviewed, tested, and approved by me.

Finally, I hope that the work presented in this thesis will be useful for others working on radar resource management. I also hope that it provides a meaningful contribution to the field. It would be particularly rewarding if my ideas developed here could support future implementations in real world radar systems. Ultimately, I hope that this research, even in a small way, can contribute to the development of more effective radar systems.

Joost Henricus Frederick Jaspers
Delft, June 2026

Abstract

Modern multifunction radar systems have to assign limited sensing resources across multiple targets and radar tasks. For mission specific scenarios, the allocation should not only improve simple performance measures such as estimation accuracy or information gain. The allocation should be focusing on the operational consequences of the decisions supported by the radar. Therefore, in this thesis a novel Bayes risk-based radar resource management framework for joint multi-target track maintenance and classification is proposed. The framework assigns resources by minimizing the expected cost of incorrect target existence and classification decisions under a shared radar time budget.

The proposed approach represents target classes and absence within a novel joint Bayesian hypothesis space. Each radar update consists of a kinematic range-azimuth measurement and a scalar class related feature measurement. These measurements are used in a recursive Bayesian update that accounts for target existence uncertainty, class uncertainty, missed detections, and clutter. The resulting belief state is then used in a Bayes risk decision formulation with a user-defined cost matrix. Mission-specific objectives can then be encoded directly, by choosing the right costs in this matrix. These could for example include different penalties for: missed targets, false tracks, class-confusion errors, high-priority target classes, or mission specific important regions.

Because the resulting expected Bayes risk cannot generally be evaluated in a closed analytical form, the conditional decision probabilities have to be approximated using Monte Carlo sampling. The sampled risk values are then converted into deterministic surrogate risk functions through smoothing, monotonicity enforcement, and finally interpolation. After this processing, the resulting surrogate functions are used in a constrained resource optimization.

The simulation results show that the proposed framework produces interpretable adaptive allocation behavior. This behavior can be related directly to the specified mission costs and the estimated belief state. The ability to encode mission-specific costs in the Bayes risk objective therefore becomes one of the key features of the framework. By changing the cost matrix, the user can determine which decision errors, target classes, or operational regions should have the strongest influence on the resource allocation.

More resources are assigned to targets for which additional sensing is expected to reduce the decision risk most strongly. This can occur because of poor measurement quality, high uncertainty, target appearance or disappearance, or increased operational cost. The framework performed particularly well in scenarios with mission-specific objectives, with cumulative risk improvements of up to approximately 30% compared with a fixed equal-budget allocation. Overall, the optimized allocation reduced the cumulative realized Bayes risk in all considered scenarios.

These results show that the proposed Bayes risk objective is suitable for joint track maintenance and classification RRM. The framework shows a direct link between sensing resource allocation and mission dependent decision costs. The observed fluctuations in the realized risk also show limitations of the myopic one step ahead implementation. This limitation motivates future work on longer horizon optimization and the joint optimization of multiple resources at the same time.

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
EKF	Extended Kalman filter
JTC	Joint tracking and classification
MAP	Maximum a posteriori
PDA	Probabilistic data association
PDA-EKF	Probabilistic data association extended Kalman filter
PCHIP	Piecewise cubic Hermite interpolating polynomial
POMDP	Partially observable Markov decision process
RRM	Radar resource management
SNR	Signal-to-noise ratio

Symbols

Symbol	Definition	Unit
Latin symbols		
$\mathbf{a}_k$	radar sensing action at time step k	–
$B_{\max}$	maximum available normalized radar time budget	–
$C_{ij,k}$	cost of selecting decision $D_{i,k}$ when hypothesis $H_{j,k}$ is true	–
$\mathbf{C}_k$	Bayesian cost matrix at time step k	–
$D_{i,k}$	decision corresponding to hypothesis $H_{i,k}$ at time step k	–
$\hat{D}_k$	selected minimum-risk decision at time step k	–
d_r	paired cumulative-risk difference for Monte Carlo run r	–
d_{ref}	reference range for the range-dependent cost function	m
F	state-transition matrix	–
$F^{-1}_{\chi^2_3}(\cdot)$	inverse cumulative distribution function of a chi-square distribution with three degrees of freedom	–
$\mathcal{G}_k$	effective validation region at time step k	–
$\mathbf{g}^k$	subgradient or finite-difference gradient vector at optimization iteration k	–
$H_{i,k}$	joint hypothesis i at time step k	–
$H_k^{(i)}$	measurement Jacobian under class hypothesis $H_{i,k}$	–
$h(\cdot)$	nonlinear radar measurement function	–
$\mathbb{I}(\cdot)$	indicator function	–
J_{eq}	cumulative realized Bayes risk under equal-budget allocation	–
J_{opt}	cumulative realized Bayes risk under optimized allocation	–
K	number of discrete simulation time steps	–
$K_k^{(i)}$	EKF gain under class hypothesis $H_{i,k}$	–
k	discrete time-step index	–
$L_{i,k}(Z_k)$	likelihood of the realized measurement set under hypothesis $H_{i,k}$	–
$L_{i,k}^{(m)}$	target-generated likelihood of validated measurement m under hypothesis $H_{i,k}$	–
M_k	number of validated measurements at time step k	–
$M_{\text{FA},k}$	number of false-alarm measurements inside the validation region	–
$M_{\max}$	maximum measurement-set cardinality used for truncation	–

Symbol	Definition	Unit
N	number of targets	–
N_c	number of target classes	–
N_{init}	number of optimizer initializations	–
N_{MC}	number of Monte Carlo samples	–
n	target index	–
$P_{D,k}^{(i)}$	probability of detection under present-target hypothesis $H_{i,k}$ at time step k	–
P_F	constant false-alarm probability	–
P_G	validation-gate probability	–
$P_{k\|k-1}^{(i)}$	predicted state covariance under class hypothesis $H_{i,k}$	–
$P_{k\|k}^{(i)}$	posterior state covariance under class hypothesis $H_{i,k}$	–
$p_F(m)$	probability of receiving m false-alarm measurements inside the validation region	–
$p_0(\mathbf{z}_k \mid Z_{1:k-1})$	single clutter generated density	–
$p_i(\mathbf{z}_k \mid Z_{1:k-1})$	single target-generated measurement density under hypothesis $H_{i,k}$	–
$Q^{(i)}$	process-noise covariance under class hypothesis $H_{i,k}$	–
r	target range	m
$\hat{r}_{k\|k-1}^{(i)}$	predicted target range under class hypothesis $H_{i,k}$	m
$r(D_{i,k} \mid Z_{1:k})$	conditional Bayes risk of decision $D_{i,k}$ after observing measurements up to time k	–
$r_{\min}$	inner boundary of the active classification region	m
$r_{\max}$	outer boundary of the active classification region	m
$R_k^{(i)}$	kinematic measurement-noise covariance under class hypothesis $H_{i,k}$	–
R_k	expected Bayes risk at time step k	–
$R_k^{(n)}$	expected Bayes risk associated with target n at time step k	–
$\tilde{R}_n(u_n)$	surrogate Bayes-risk function for target n	–
$\mathbf{s}_k$	target kinematic state vector at time step k	–
$\hat{\mathbf{s}}_{k\|k-1}^{(i)}$	predicted state estimate under class hypothesis $H_{i,k}$	–
$\hat{\mathbf{s}}_{k\|k}^{(i)}$	posterior state estimate under class hypothesis $H_{i,k}$	–
$S_k^{(i)}$	innovation covariance under class hypothesis $H_{i,k}$	–
$\mathrm{SNR}_k^{(i)}$	signal-to-noise ratio under class hypothesis $H_{i,k}$	–
SNR_0	reference signal-to-noise ratio	–
T	revisit interval	s
T_n	revisit interval assigned to target n	s
T_{opt}	optimization interval	s
$\mathbf{T}$	Markov hypothesis transition matrix	–
T_{ij}	transition probability from hypothesis $H_{j,k-1}$ to hypothesis $H_{i,k}$	–
u_n	selected scalar resource variable for target n	–
$V_k^{(0)}$	effective validation volume at time step k	–
$\mathbf{v}_{y,k}^{(i)}$	kinematic measurement-noise vector under class hypothesis $H_{i,k}$	–
$v_{\xi,k}^{(i)}$	feature measurement noise under class hypothesis $H_{i,k}$	–
$\mathbf{w}_k^{(i)}$	process-noise vector under class hypothesis $H_{i,k}$	–
$w(r)$	range-dependent cost-weighting function	–
$w_{i,k}$	normalized predicted class weight for class hypothesis $H_{i,k}$	–
x_k	target position in the x-direction at time step k	m
$\dot{x}_k$	target velocity in the x-direction at time step k	m/s
y_k	target position in the y-direction at time step k	m
$\dot{y}_k$	target velocity in the y-direction at time step k	m/s
$\mathbf{y}_k$	kinematic radar measurement vector containing range and azimuth at time step k	–

Symbol	Definition	Unit
$\mathbf{y}_k^{(m)}$	kinematic component of the m-th validated measurement at time step k	–
$\hat{\mathbf{y}}_{k\|k-1}^{(i)}$	predicted kinematic radar measurement under class hypothesis $H_{i,k}$	–
$\tilde{\mathbf{y}}_k^{(i)}$	kinematic measurement innovation under hypothesis $H_{i,k}$	–
Z_k	validated measurement set at time step k	–
$Z_{1:k}$	sequence of measurement sets up to time step k	–
$\mathbf{z}_k$	combined kinematic-feature measurement vector	–
$\mathbf{z}_k^{(m)}$	m-th validated measurement at time step k	–
$\hat{\mathbf{z}}_{k\|k-1}^{(i)}$	predicted joint kinematic-feature measurement under class hypothesis $H_{i,k}$	–
Greek symbols		
α	strength of the range-dependent cost increase	–
β_{FA}	false-alarm density in the validation region	–
$\beta_{i,k}^{(m)}$	PDA association probability of validated measurement m under hypothesis $H_{i,k}$	–
$\beta_{i,k}^{(0)}$	PDA no-association probability under hypothesis $H_{i,k}$	–
Γ_i	Bayesian decision region associated with decision $D_{i,k}$	–
$\Gamma_i^{(m)}$	Bayesian decision region for decision $D_{i,k}$ and measurement-set cardinality m	–
γ	probability-mass threshold used for cardinality truncation	–
$\delta_k(Z_k, m)$	Bayesian decision rule applied to measurement set Z_k with cardinality m	–
ΔJ	difference in cumulative realized Bayes risk between equal-budget and optimized allocation	–
ϵ	finite-difference perturbation used for gradient approximation	–
$\lambda_{\mathrm{FA},k}$	expected number of false alarms inside the validation region at time step k	–
μ_i	class-dependent feature mean for class i	–
μ_ξ	vector of class-dependent feature means	–
$\boldsymbol{\nu}_{i,k}^{(m)}$	kinematic innovation vector of validated measurement m under hypothesis $H_{i,k}$	–
$\bar{\boldsymbol{\nu}}_{i,k}$	PDA-weighted combined innovation vector under hypothesis $H_{i,k}$	–
$C_{\nu,i,k}$	innovation-spread covariance due to association uncertainty under hypothesis $H_{i,k}$	–
ξ_k	scalar class-related feature measurement at time step k	–
$\pi_{i,k}$	posterior probability of hypothesis $H_{i,k}$	–
$\pi_{i,k\|k-1}$	predicted prior probability of hypothesis $H_{i,k}$ before receiving Z_k	–
$\boldsymbol{\pi}_k$	joint target-existence and classification belief vector	–
$\boldsymbol{\pi}_{\mathrm{birth}}$	initial belief vector assigned to a newly appearing target	–
ρ_k	step size at optimization iteration k	–
ρ_{base}	base step-size parameter	–
$\Sigma_k^{(0)}$	effective validation covariance obtained by moment matching	–
$\Sigma_k^{(i)}$	joint measurement covariance under class hypothesis $H_{i,k}$	–
$\sigma_{r,k}^{(i)}$	range-measurement standard deviation under class hypothesis $H_{i,k}$	m
$\sigma_{r,0}^{(i)}$	reference range-measurement standard deviation under class hypothesis $H_{i,k}$	m
$\sigma_{\theta,k}^{(i)}$	azimuth-measurement standard deviation under class hypothesis $H_{i,k}$	rad

Symbol	Definition	Unit
$\sigma_{\theta,0}^{(i)}$	reference azimuth-measurement standard deviation under class hypothesis $H_{i,k}$	rad
$\sigma_{\xi,k}^{(i)}$	feature-measurement standard deviation under class hypothesis $H_{i,k}$	–
$\sigma_{\xi,0}^{(i)}$	reference feature-measurement standard deviation under class hypothesis $H_{i,k}$	–
$\sigma_{\omega}^{(i)}$	maneuverability standard deviation under class hypothesis $H_{i,k}$	m/s^2
τ_n	dwell time allocated to target n	s
τ_0	reference dwell time	s
$\chi_{\mathrm{border}}(r)$	indicator function for the active classification region	–
θ_k	azimuth measurement at time step k	rad

1 Introduction

1.1. Multifunction Radar and Radar Resource Management

Modern radar systems are increasingly required to operate in complex and dynamic environments, where they should handle multiple objectives and tasks under limited resources [1], [2]. A radar system may need to perform detection, target tracking, and classification, while adapting its sensing actions to the current operational situation. This requires flexibility, since different targets and sensing tasks may become important at different times.

Older mechanically rotating radars offered only limited flexibility, since their updates were determined by the mechanical scan pattern. Modern electronically steered phased array antennas remove this dependence on mechanical antenna pointing [3]. By steering the beam electronically, a phased array radar can rapidly switch between different directions and targets without physically repositioning the antenna. This enables a single multifunction radar to perform multiple radar functions, such as detection, target tracking, classification, and other mission-dependent sensing tasks [4], [5].

The flexibility of such radar hardware can only be fully used when the radar actively manages its sensing resources. For example, the radar may allocate more sensing time to a sector with a high expected threat level, revisit uncertain tracks more often, or spend more time observing a target when improved classification confidence is required. An example of the resulting competing sensing tasks is shown in Figure 1.1.

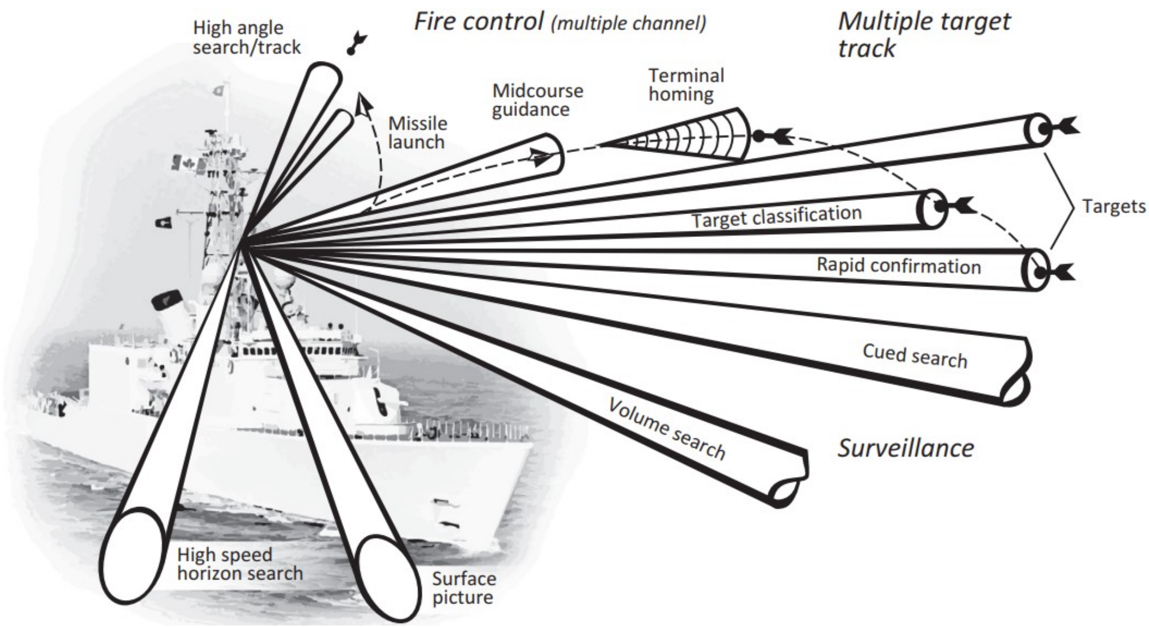

Figure 1.1: Multiple functions of a ship borne radar system, adapted from Moo and Ding [6].

Despite the radar's flexibility, the available radar resources remain limited. Resources have to be shared among multiple targets and competing radar functions. Increasing the resources assigned to one target may improve its probability of detection, measurement accuracy, and classification performance, but simultaneously reduces the resources available for other targets or sensing tasks [7]. This creates a resource allocation problem. The radar must continuously decide how to distribute its limited sensing capacity over it's tasks.

RRM addresses this allocation problem by determining how the available radar resources should be divided over targets and sensing tasks. Candidate resource allocations are often evaluated using an objective function that assigns a numerical value to the expected radar performance resulting from each allocation. The algorithm then selects the allocation that minimizes this objective function, according to the available radar-resource constraints. In this way, the chosen objective function translates the operational goal into an objective measure, so that minimizing the objective corresponds to achieving the goal as well as possible within the available resources.

For the joint track maintenance and classification problem considered in this thesis, the objective should thus not only reward accurate state estimation, but also account for the consequences of target classification decisions. More generally, RRM objectives may represent tracking uncertainty, estimation error, information gain, target or class priority, or the expected consequences of radar decision errors.

The selected objective function therefore defines what the radar management system considers most valuable. It determines which aspects of radar performance are improved, how competing targets and sensing tasks are compared, and which targets receive priority when the available resources are insufficient to optimize all tasks at the same time. We can thus conclude that the design of the objective function plays a central role in determining whether the resulting resource allocation reflects the operational goals of the radar user.

1.2. Conventional Radar Resource Management Objectives

Simple resource allocation strategies, such as equal allocation, fixed scheduling, and priority-based scheduling, do not continuously evaluate the expected benefit of spending additional resources. Although these methods are sometimes straightforward to implement and provide useful reference strategies, they cannot fully adapt it's allocation to changes in target uncertainty, measurement conditions, or sensing requirements.

More adaptive RRM approaches often use tracking based objective functions such as tracking accuracy or estimation uncertainty. Possible measures include individual entries, the trace, or the determinant of the predicted state-error covariance. These objectives assign more radar resources to targets for which additional measurements are expected to reduce the state estimation uncertainty most.

Schöpe [8] applied covariance-based objectives to multi-target tracking RRM under a shared radar budget. Scalar measures of the predicted Kalman-filter covariance were used to evaluate the expected tracking performance associated with different resource allocations. As a result, the radar assigned more resources to targets for which additional sensing was expected to reduce the tracking uncertainty most, for example due to increased maneuverability, reduced measurement quality or targets at further ranges.

Theoretical estimation bounds may similarly be used as the RRM objective. For example, Yi *et al.* [9] formulated radar resource allocation by minimizing Bayesian estimation-error bounds. Such objectives provide a mathematically well-defined relationship between the allocated radar resources and the expected estimation performance.

Information-theoretic quantities provide another class of performance measures. Instead of directly evaluating the tracking covariance, these objectives quantify the expected information obtained from a future measurement. Information gain and Rényi divergence have been used to evaluate candidate sensing actions in detection and tracking problems [10], [11]. Kullback-Leibler information has similarly been applied to waveform selection for target classification [12]. Such measures can provide a common scalar representation of uncertainty reduction across different sensing actions and tasks [8].

Although these objectives support adaptive resource allocation, they primarily quantify tracking-based estimation quality or uncertainty reduction. They do not directly represent the operational consequences of the radar decisions made using the resulting information. A reduction in estimation uncertainty is not necessarily equally valuable for every target, target class, or operational situation.

For example, a covariance-based objective may allocate more resources to a distant target because its measurements are less accurate and its predicted state uncertainty is larger. This behavior is consistent with the selected objective, since additional sensing may produce a relatively large reduction in tracking uncertainty. However, a nearby target may be more operationally relevant, even when its state is already estimated more accurately. An allocation that is optimal in terms of estimation accuracy may therefore not reflect the priorities of the radar user.

Operational preferences can be introduced into conventional objectives using weighting factors, target priorities, or additional correction terms. However, the relationship between these parameters and the consequences of the resulting radar decisions is generally indirect. Changing the operational objective may therefore require the performance measure to be redesigned or retuned, while the resulting allocation behavior may remain difficult to interpret.

This motivates the development of RRM objectives that represent not only sensing or estimation quality, but also the operational relevance of the targets and the consequences of the decisions supported by the radar measurements.

1.3. Threat-Based Radar Resource Management

Threat-based RRM provides a means of incorporating mission relevance into the resource-allocation objective. Instead of evaluating a target only through its estimation uncertainty, threat-based methods map target information to a measure of operational importance.

Early threat-based sensor-management approaches used target-state and classification information to estimate the operational threat associated with a target [13], [14]. In this way, the selected sensing action could depend on the operational context rather than on tracking or classification accuracy alone.

Building on this multi-target resource management framework, Schöpe extended the approach to joint target tracking and classification using a threat-variance-based objective [8]. In this formulation, the target state and class probabilities are mapped to a common threat quantity, and the resource allocation is determined by the expected reduction in uncertainty of that threat estimate. Tracking and classification performance are therefore combined within a single scalar objective, while the same constrained multi-target budget-allocation structure is retained.

This established that tracking and classification do not need to be treated as independent sensing tasks with separate objectives. Instead, their contributions can be evaluated through a shared mission-related quantity, allowing the resource allocation to account jointly for the target state, class belief, and operational context.

However, the threat value remains an intermediate representation of mission relevance. It indicates how important or uncertain a target is considered to be, but it does not directly specify the operational consequences of the possible radar decisions or distinguish explicitly between different types of decision errors.

Threat-based objectives therefore provide an important connection between radar resource allocation and mission relevance. A more direct decision-oriented formulation can build on this idea by representing the consequences of the radar decisions themselves, rather than expressing operational relevance only through an intermediate threat quantity.

1.4. Bayes Risk-Based Radar Resource Management

Following the constrained multi-target RRM settings considered in the preceding work, Chioccarello *et al.* [15] introduced a Bayes Risk-based objective for radar track maintenance. Instead of using a resource allocation through estimation covariance, information gain, or an intermediate threat quantity, their formulation connected the allocation to the expected consequences of target existence decisions.

In the considered track maintenance problem, the radar must determine if a maintained track corresponds to a real target or to a non-existent object. Incorrect decisions include continuing a false track and terminating a track while the target is still present.

The proposed objective function is based on Bayes risk. Each possible track maintenance outcome is assigned a decision cost. Then the probability of that outcome depends on the available radar information and the allocated sensing resources. The objective function finally represents the expected operational cost of the possible track maintenance decisions.

This formulation has two important advantages for RRM. First, it provides an interpretable connection between the resource allocation and the decisions that the radar system is expected to make. Resources are assigned according to their expected contribution to reduce costly decision errors, rather than according to an abstract performance measure alone.

Second, the decision costs provide a direct mean of incorporating the preferences of the radar user. Different errors can be assigned different levels of importance according to their operational consequences. The resource allocation behavior can therefore be adapted to different missions by changing the costs associated with the decision outcomes, without redefining the underlying tracking or resource management framework.

Compared with the threat-based objectives from Schöpe *et al.*[16], Bayes risk avoids the need to first transform target information into an intermediate threat score. The operational relevance is instead represented directly through the consequences assigned to the possible radar decisions. This makes the relationship between the user-defined objective and the resulting allocation behavior more interpretable.

The formulation introduced by Chioccarello *et al.*[17], however, considers a binary track maintenance problem that distinguishes only between target absence and target presence. It does not distinguish between different target classes. Therefore it cannot directly represent classification errors or differences in the operational consequences of different classes.

This limitation motivates the extension considered in this thesis. The objective is to retain the interpretability and adaptability of Bayes Risk-based RRM while extending the decision problem to jointly account for track maintenance as well as the classification of targets.

1.5. Research Gap and Thesis Objective

Joint tracking and classification (JTC) methods have shown that kinematic state information and target-class information can support one another. Tracking information may provide evidence about target class, while class dependent motion and measurement models may improve state estimation [18], [19], [20]. However, these methods primarily aim to improve inference performance and do not generally use the combined operational consequences of target existence and classification decisions as the objective of multi-target radar resource allocation.

To the best of the author's knowledge, a unified Bayes Risk-based RRM formulation in which track maintenance and classification are represented within one decision objective has not yet been established for multi-target radar resource allocation.

The objective of this thesis is therefore to extend the Bayes Risk-based track maintenance RRM from Chioccarello *et al.*[15] to a joint track maintenance and classification framework. The resulting objective function should allow limited radar resources to be allocated according to both the probabilities and the user-defined operational costs of possible target existence and classification decisions. This leads to the following research question and subquestions:

Research question How can a Bayes Risk-based objective function be designed for the joint track maintenance and classification of multiple targets under limited radar resources while accounting for the differing operational costs of decision errors?

- **SQ1:** How can a Bayes Risk-based objective function be formulated for joint track maintenance and classification in radar resource management?
- **SQ2:** How can operationally different decision errors be incorporated into this objective function through user defined costs?

- **SQ3:** How can the coupling between tracking and classification be reflected in the risk-based resource allocation framework?
- **SQ4:** How does the proposed objective influence the realized Bayes Risk and resource allocation behavior under different multi-target scenarios and decision-cost settings compared with a fixed equal budget strategy?

1.6. Contributions

The main contributions of this thesis are:

- the extension of binary Bayes Risk-based track maintenance RRM to a novel joint multi-hypothesis formulation containing target absence and multiple target classes, integrating target existence beliefs, classification beliefs, and class dependent tracking information within one shared recursive estimation and resource management framework;
- the formulation of a measurement-set-based Bayes Risk objective that accounts for multiple validated measurements, missed detections, clutter measurements, and uncertain target-measurement association;
- the derivation and numerical evaluation of the conditional decision probabilities for all combinations of target absence, target presence, and target class decisions under this multiple measurement formulation;
- the incorporation of user defined decision cost matrices in which each matrix entry is directly associated with a specific class dependent decision error, allowing different operational consequences to be weighted explicitly;
- the development of deterministic surrogate risk functions that enable the numerically evaluated Bayes Risk objective to be used in a constrained dwell time allocation;
- the evaluation of the proposed framework under varying target scenarios and operational cost structures.

1.7. Thesis Structure

The remainder of this thesis is structured as follows.

The RRM problem considered in this work is formulated in chapter 2. This chapter introduces the target hypotheses and decisions, the user defined decision costs, the available radar resources, and the constrained multi-target resource allocation problem.

The proposed Bayes Risk-based methodology for joint track maintenance and classification is presented in chapter 3. This chapter introduces the joint target existence and classification belief representation. Then it derives the corresponding Bayesian decision objective, and explains the numerical evaluation of the conditional decision probabilities and expected Bayes risk. It also describes the construction of surrogate risk functions and their use in the constrained resource allocation problem.

The assumed radar and target scenario in this work is defined in chapter 4. This chapter presents the state evolution and radar measurement models, the dependence of the detection probability and measurement quality on dwell time and target range, and the recursive tracking and probabilistic data-association framework used to update the target state estimates.

The simulation setup and numerical evaluation are presented in chapter 5. This chapter evaluates the proposed framework under different multi-target scenarios and operational cost configurations. The results are used to investigate the realized Bayes risk, the resulting resource allocation behavior, and the influence of the user defined decision costs.

Finally, the main findings, limitations, and recommendations for future work are summarized in chapter 6.

2

Problem Formulation

This chapter introduces the problem formulation for the joint track maintenance and classification RRM framework considered in this work. The goal is to dynamically allocate the limited radar sensing resources over multiple targets such that the expected cost of tracking and classification errors, as defined in the user specified objective function, is reduced. The framework therefore does not aim to minimize all errors equally, but to prioritize the errors that are most costly for the considered operational objective.

2.1. Modeling

2.1.1. State space motion model

To describe the target motion over time, a discrete time state space model is used. Each target is represented by a kinematic state vector containing position and velocity components in two spatial dimensions,

$$\mathbf{s}_k = \begin{bmatrix} x_k \\ y_k \\ \dot{x}_k \\ \dot{y}_k \end{bmatrix}, \tag{2.1}$$

where x_k and y_k denote the target position at timestep k, and $\dot{x}_k$ and $\dot{y}_k$ denote the corresponding velocity components.

The target dynamics are described by a general discrete time state transition model,

$$\mathbf{s}_{k+1} = f(\mathbf{s}_k, \mathbf{w}_k, \mathbf{a}_k), \tag{2.2}$$

where $f(\cdot)$ denotes the state evolution model. The term $\mathbf{w}_k$ represents process noise, which accounts for uncertainty in the target dynamics and unmodeled target maneuvers.

The sensing action $\mathbf{a}_k$ contains the radar resource variables selected by the RRM framework. These variables determine how the available sensing resources are distributed over the targets and therefore influence the information obtained from it's future measurements.

In addition to the kinematic state, each target is assumed to belong to one of a finite set of target classes. The class c_n is assumed constant during the considered observation interval.

2.1.2. Radar measurement model

The radar measurements are assumed to consist of two components: kinematic measurements used for target tracking and class dependent feature measurements used for target classification.

The kinematic radar measurement vector is modeled as

$$\mathbf{y}_k = h(\mathbf{s}_k, \mathbf{v}_{y,k}, \mathbf{a}_k), \tag{2.3}$$

where $\mathbf{y}_k$ denotes the kinematic radar measurement vector, $h(\cdot)$ denotes the radar measurement function, $\mathbf{s}_k$ represents the target kinematic state, and $\mathbf{v}_{y,k}$ represents the kinematic measurement noise. The dependence on the sensing action $\mathbf{a}_k$ indicates that the selected radar resource allocation affects the measurement noise statistics.

In addition to the kinematic measurement vector, the radar is assumed to provide a scalar class dependent feature measurement,

$$\xi_k = g(c_n, v_{\xi,k}, \mathbf{a}_k), \tag{2.4}$$

where ξ_k denotes the feature measurement, c_n denotes the discrete target class from a finite set of classes, $g(\cdot)$ denotes the class dependent feature measurement function, and $v_{\xi,k}$ represents the feature measurement noise.

The feature measurement could represent target specific properties that can support target classification. Possible examples could include the radar cross section or micro-Doppler signatures generated by rotating or vibrating target components.

The radar resource allocation thus does not only influence the tracking accuracy, but also the target classification performance, by changing the reliability of both the kinematic and feature measurements.

2.2. Algorithms

For the tracking and classification algorithms, many possible approaches exist. The goal in this work is to estimate both the target track and the target class recursively using information from previous measurements. The objective is not necessarily to identify the optimal algorithm, but rather to provide a flexible RRM framework that is compatible with a wide range of these recursive estimation methods. An overview of possible tracking filters is already provided in the work by Chioccarello *et al.* [17].

For the classification problem considered in this thesis, the objective is to infer the target class from noisy radar measurements. Several classification approaches can be used depending on the assumed measurement model and decision framework.

For statistical classification problems, Maximum Likelihood (ML) and Maximum A Posteriori (MAP) classifiers provide common solutions by selecting the class with the highest likelihood or posterior probability, respectively. In radar applications, Likelihood Ratio Tests (LRTs) are frequently employed for binary classification and detection problems due to their analytical tractability and their natural connection to Bayesian decision theory[21].

For more complex feature spaces, machine learning approaches such as Support Vector Machines (SVMs) [22] and Neural Networks [23] may also be considered. Recent developments in radar target classification using micro-Doppler signatures have demonstrated strong performance using data driven learning methods [24].

Although these machine learning approaches are not explicitly integrated into the recursive Bayesian framework considered in this work, they could later provide valuable class related information that can be incorporated into the recursive estimation structure in future work. For example, the output of a trained classifier could be interpreted as a probabilistic feature measurement or class likelihood and subsequently combined with the tracking information within the Bayesian belief update. In this way, the proposed framework can stay compatible with future extensions involving learned feature representations or data driven classification modules.

For JTC algorithms, tracking and classification are performed within a shared recursive estimation framework to estimate the kinematic target state and obtain the target class from sequential measurements at the same time. Since both processes can be coupled, the estimated target motion can provide information about the target class, while class information can be used to select more appropriate motion models and thereby improve the tracking performance.

Multiple Bayesian JTC approaches have been proposed in the literature. Challa and Pulford [19] formulated the problem using a joint state class probability density, in which class dependent motion models and flight envelope information couple the tracking and classification processes. Particle filter implementations were proposed by Gordon *et al.* [18] and Vercauteren *et al.* [25], allowing nonlinear and non Gaussian JTC problems to be treated using class conditioned particles and sequential Monte Carlo filtering. Multiple-model approaches have also been developed, including the particle filter and mixture Kalman filter algorithms proposed by Angelova and Mihaylova [26], in which separate class conditioned filters provide both state estimates and likelihoods for the recursive update of the class probabilities.

The framework developed in this thesis is therefore not tied to one specific tracking or classification algorithm, but can be combined with any recursive estimation approach that provides the state estimates, class related likelihoods, and belief information required for Bayes risk-based resource allocation.

2.3. Radar Resource Management Problem for Multi-Target Joint Track Maintenance and Classification

2.3.1. Resource variables: dwell time and revisit interval

The limited radar resources in this work are formulated in terms of time allocation per target. It is assumed that the radar can point to only one target at a time, such that the available sensing time must be distributed over all targets under consideration.

For each target n, two resource variables are assigned: the dwell time τ_n, representing the duration for which the radar continuously observes the target during one visit, and the revisit interval T_n, representing the time between two consecutive observations of the same target.

The fraction of radar time occupied by target n within one scheduling interval is therefore given by

$$\frac{\tau_n}{T_n}. \tag{2.5}$$

This quantity expresses how much of the available radar time budget is used for a single target and forms the basic resource term in the global budget constraint.

2.3.2. Global budget constraint

As introduced in the previous section, each target consumes a fraction of the available radar time through the ratio between the dwell time and the revisit interval. Since the radar can point at only one target at a time, the sum of all allocated target fractions must remain within the available radar budget during one scheduling interval.

This leads to the global time budget constraint

$$\sum_{n=1}^{N} \frac{\tau_n}{T_n} \leq B_{\max}, \tag{2.6}$$

where N denotes the number of targets under consideration and $B_{\max} \in [0, 1]$denotes the maximum fraction of the scheduling interval that may be occupied by radar sensing. A value of $B_{\max} = 1$ corresponds to full utilization of the available radar time, while lower values reserve part of the interval for other radar functions or operational limitations.

2.4. Final optimization problem

The objective of the RRM framework is to allocate the dwell times and revisit intervals over all targets such that the expected decision cost is minimized while satisfying the available radar time budget.

The resulting optimization problem is formulated as

$$
\begin{aligned}
\min_{\mathbf{T},\boldsymbol{\tau}} \quad & \sum_{n=1}^{N_k} \mathcal{C}(\mathbf{s}_k^n, \tau_n, T_n) \\
\text{s.t.} \quad & \sum_{n=1}^{N_k} \frac{\tau_n}{T_n} \leq B_{\max}, \\
& \boldsymbol{\tau} \geq 0, \\
& \mathbf{T} \geq 0.
\end{aligned}
\tag{2.7}
$$

where $\mathcal{C}(\cdot)$ is a generic target dependent objective function used to evaluate the quality of a resource allocation.

2.5. Bayes Risk-Based Track Maintenance as Starting Point

The optimization problem in Equation 2.7 is formulated using the generic objective function $\mathcal{C}(\cdot)$. In this thesis, the objective function $\mathcal{C}(\cdot)$ is based on the Bayes risk-based RRM framework for multi-target track maintenance introduced by Chioccarello *et al.* [17].

In this framework, RRM is formulated around the decisions made during track maintenance, rather than around tracking accuracy or estimation uncertainty alone. For each maintained track, the target existence problem is represented by the binary hypothesis set

$$
\mathcal{H}_k = \{H_{0,k}, H_{1,k}\}, \tag{2.8}
$$

where $H_{0,k}$ denotes the hypothesis that the target is absent at time step k, and $H_{1,k}$ denotes the hypothesis that the target is present.

The uncertainty about target existence is represented by a binary Bayesian belief state,

$$
\boldsymbol{\pi}_{k|k} = \begin{bmatrix} \pi_{0,k|k} \\ \pi_{1,k|k} \end{bmatrix}, \tag{2.9}
$$

where $\pi_{0,k|k} = P(H_{0,k} \mid Z_{1:k})$ and $\pi_{1,k|k} = P(H_{1,k} \mid Z_{1:k})$ denote the posterior probabilities for target absence and target presence, respectively, after processing the measurement history $Z_{1:k}$.

Before processing the measurement set at time step k, the belief is first predicted from the previous posterior belief using a Markov transition model for target existence,

$$
\begin{bmatrix} \pi_{0,k|k-1} \\ \pi_{1,k|k-1} \end{bmatrix} = \begin{bmatrix} P(H_{0,k} \mid H_{0,k-1}) & P(H_{0,k} \mid H_{1,k-1}) \\ P(H_{1,k} \mid H_{0,k-1}) & P(H_{1,k} \mid H_{1,k-1}) \end{bmatrix} \begin{bmatrix} \pi_{0,k-1|k-1} \\ \pi_{1,k-1|k-1} \end{bmatrix}, \tag{2.10}
$$

The predicted belief is therefore given by

$$
\boldsymbol{\pi}_{k|k-1} = \begin{bmatrix} \pi_{0,k|k-1} \\ \pi_{1,k|k-1} \end{bmatrix}, \tag{2.11}
$$

where $\pi_{0,k|k-1}$ and $\pi_{1,k|k-1}$ denote the predicted probabilities of target absence and target presence before receiving the measurement set Z_k. The transition probabilities account for the possibility that a target remains absent, appears, disappears, or remains present between two consecutive time steps.

After receiving the measurement set Z_k, the predicted belief is corrected using Bayes' rule,

$$
\pi_{i,k|k} = P(H_{i,k} \mid Z_{1:k}) = \frac{p(Z_k \mid H_{i,k}, Z_{1:k-1})\pi_{i,k|k-1}}{\sum_{\ell=0}^{1} p(Z_k \mid H_{\ell,k}, Z_{1:k-1})\pi_{\ell,k|k-1}}, \qquad i \in \{0,1\}. \tag{2.12}
$$

Here, $p(Z_k \mid H_{i,k}, Z_{1:k-1})$ denotes the likelihood of the received measurement set under hypothesis $H_{i,k}$, conditioned on the previous measurement history. This prediction-correction structure allows the binary track maintenance framework to recursively update the probability that a maintained track corresponds to an existing target.

Based on the updated binary belief state, the track maintenance framework selects one of two corresponding decisions,

$$\mathcal{D}_k = \{D_{0,k}, D_{1,k}\}, \tag{2.13}$$

where $D_{0,k}$ represents the decision to terminate the track and $D_{1,k}$ represents the decision to maintain the track. Consequently, two types of incorrect track maintenance decisions can occur: a false track may be maintained while the target is absent, or a true track may be terminated while the target is still present.

The operational consequences of the possible decision outcomes are represented through the binary cost matrix

$$\mathbf{C}_k = \begin{bmatrix} C_{00,k} & C_{01,k} \\ C_{10,k} & C_{11,k} \end{bmatrix}, \tag{2.14}$$

where $C_{ij,k}$ denotes the cost of selecting decision $D_{i,k}$ when hypothesis $H_{j,k}$ is true. The diagonal entries represent the costs associated with correct decisions, while the off-diagonal entries represent the costs of incorrect track maintenance decisions. In particular, $C_{01,k}$ represents the cost of terminating a track while the target is present, whereas $C_{10,k}$ represents the cost of maintaining a track while the target is absent.

The expected Bayes risk associated with the track maintenance problem can be expressed at a high level as

$$R_k = \sum_{i=0}^{1}\sum_{j=0}^{1} C_{ij,k} P(D_{i,k} \mid H_{j,k}) \pi_{j,k|k-1}, \tag{2.15}$$

where $\pi_{j,k|k-1}$ denotes the predicted probability of hypothesis $H_{j,k}$ before receiving the measurement set at time step k, and $P(D_{i,k} \mid H_{j,k})$ denotes the probability of selecting decision $D_{i,k}$ when hypothesis $H_{j,k}$ is true. The Bayes Risk therefore combines the probability of each decision outcome with the operational cost assigned to that outcome.

The decision probabilities depend on the information provided by the radar measurements. Since the allocated dwell time and revisit interval affect the probability of detection and the measurement quality, the Bayes Risk is indirectly a function of the radar resource allocation. The generic target dependent objective function in Equation 2.7 can therefore be instantiated using the expected Bayes Risk,

$$\mathcal{C}\left(\mathbf{s}_k^n, \boldsymbol{\pi}_{k|k-1}^n, \tau_n, T_n\right) = R_k^n\left(\mathbf{s}_k^n, \boldsymbol{\pi}_{k|k-1}^n, \tau_n, T_n\right). \tag{2.16}$$

The individual target risks can be summed and minimized under the shared radar time constraint. In this way, the allocated resources are selected according to their expected contribution to reducing costly track maintenance decision errors.

The formulation of Chioccarello *et al.*[17] is constrained to the binary target existence problem and therefore considers only track maintenance decisions. Target classification is not yet included as part of the resource management objective. In particular, the belief state only distinguishes between target absence and target presence, and the cost matrix only represents the consequences of maintaining or terminating a track. The formulation therefore does not account for class related information, for the operational consequences of incorrect classification decisions, or for mission dependent priorities in which errors involving one target class may be more important than errors involving another.

The problem addressed in this thesis is therefore to investigate how classification can be incorporated into this Bayes Risk-based RRM framework. The objective is to determine how the operational consequences of both track maintenance and classification errors can be represented through user defined costs.

2.6. System overview

The proposed framework consists of two interacting components: a joint track maintenance and classification module, and a RRM module. Their interaction is shown schematically in Figure 2.1.

The joint track maintenance and classification module processes the radar measurements and recursively estimates the kinematic target states, target existence probabilities, and target class beliefs. These estimates form the information state used by the resource management module.

The RRM module then determines how the available radar time budget should be distributed among the targets. Based on the current target state estimates, classification beliefs, sensing conditions, and user defined decision costs, the module evaluates the expected Bayes Risk associated with the possible resource allocations.

The resulting resource allocation is fed back into the next radar scan, thereby affecting the future measurement quality, detection probability, belief update, and subsequent resource allocation decision. The framework therefore forms a closed loop adaptive radar management system in which sensing, estimation, decision making, and resource allocation are coupled over time.

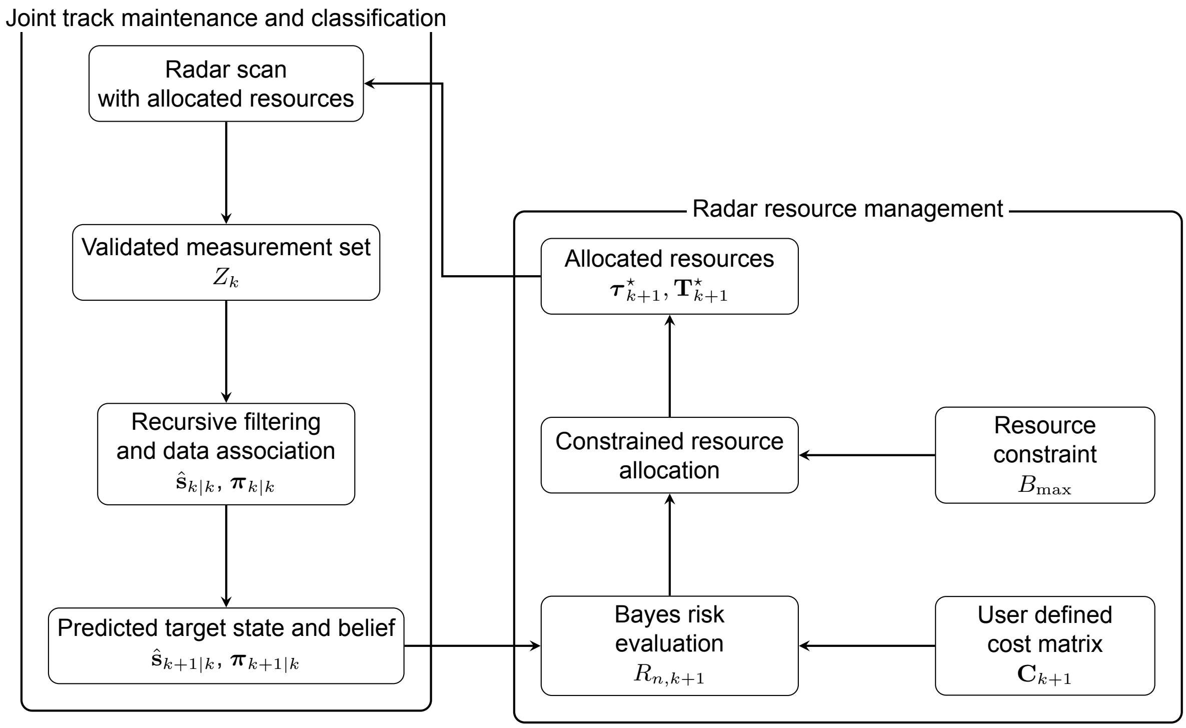

Figure 2.1: System level overview of the proposed Bayes Risk-based RRM framework. The joint track maintenance and classification module processes the validated radar measurements and estimates the target states and joint beliefs. The RRM module evaluates the expected Bayes Risk and allocates the available radar resources for the next scan under the imposed resource constraint.

3 Methodology

This chapter develops the methodology proposed in this thesis for solving the joint track maintenance and classification RRM problem formulated in Equation 2.7. Building on the Bayes risk-based track maintenance framework introduced in section 2.5, the proposed methodology extends the binary target existence decision problem to a joint target existence and classification setting. The central objective is to connect the operational consequences of both track maintenance and classification decision errors to the allocation of the available sensing resources.

The chapter first introduces the proposed joint belief representation used to describe target existence and classification uncertainty. It then presents the measurement-set model required for Bayesian belief updating and Bayes Risk evaluation. Based on these beliefs, a Bayes Risk decision framework is introduced, after which the expected Bayes risk is evaluated numerically. Finally, the resulting risk objective is used into a surrogate-based constrained radar resource allocation method. The general formulation allows both dwell time and revisit interval to be considered as radar resource variables. The implemented optimization method in this thesis focuses on a one dimensional surrogate risk allocation.

3.1. Joint Belief Representation

In the joint track maintenance and classification problem, the radar must infer both whether a target is present and, if present, to which class it belongs. To represent this coupled uncertainty, this thesis proposes a joint hypothesis space in which target absence and the possible target classes are modeled as mutually exclusive hypotheses.

This joint formulation is important because these two sources of information can be coupled. Kinematic evidence that supports the presence of a target, should also affect the credibility of the class hypotheses, while class dependent feature measurements can provide extra evidence that a target is indeed present. For example, a measurement with a feature value that is highly consistent with one target class should not only increase the probability of that class, but also increase the probability that the measurement originated from a real target rather than clutter. On the other hand, if the target existence probability is low, high conditional class likelihoods should not by themselves be interpreted as reliable evidence that a real target of that class is present.

The hypotheses are therefore formulated as mutually exclusive joint events: either no target is present, or a target is present and belongs to one of the possible classes. This allows the posterior belief to capture the interaction between track maintenance and classification in a single probability vector. As a result, the following Bayesian decision rule and Bayes Risk objective can evaluate the operational consequences of absence, missed detection, and class confusion within one consistent framework.

The joint hypothesis space is defined as

$$\mathcal{H} = \{H_{0,k}, H_{1,k}, \ldots, H_{N_c,k}\}, \tag{3.1}$$

where $H_{0,k}$ denotes the hypothesis that the target is absent at time k, while $H_{i,k}$, for $i = 1, \ldots, N_c$, denotes the hypothesis that the target is present and belongs to class i from a set of total N_c classes.

The belief over the hypotheses can then be represented by the posterior probability vector

$$\boldsymbol{\pi}_k = \begin{bmatrix} \pi_{0,k} \\ \pi_{1,k} \\ \vdots \\ \pi_{N_c,k} \end{bmatrix}, \qquad \pi_{i,k} = P(H_{i,k} \mid Z_{1:k}), \tag{3.2}$$

where $Z_{1:k} = \{Z_1, \ldots, Z_k\}$ denotes the sequence of measurements up to time k.

3.2. Prior Prediction of Hypothesis Probabilities

Between consecutive radar scans, the joint hypothesis probabilities may change even before a new measurement is received. For example, a target that was present at the previous time step may disappear from the surveillance region because it was neutralized. Conversely, a previously absent target may enter the surveillance region. The belief state should therefore not only be updated by measurements, but also predicted forward in time to account for possible changes in target existence and class state.

To describe this temporal evolution, the hypotheses are assumed to evolve according to a first-order Markov process. The predicted prior probabilities are therefore given by

$$\boldsymbol{\pi}_{k|k-1} = \mathbf{T}\boldsymbol{\pi}_{k-1}, \tag{3.3}$$

where $\mathbf{T}$ is the hypothesis transition matrix with entries

$$T_{ij} = P(H_{i,k} \mid H_{j,k-1}). \tag{3.4}$$

The transition matrix therefore models the probability that the joint target existence and class hypothesis changes between two consecutive scans. In this way, the predicted prior $\boldsymbol{\pi}_{k|k-1}$ represents the belief before the new measurement set Z_k is incorporated.

3.3. Measurement-Set Likelihood Model

This section develops the probabilistic measurement-set model required to evaluate the Bayesian decision framework introduced in section 3.1. The formulation accounts for multiple measurements, missed detections, clutter generated measurements, and uncertain target measurement associations.

The resulting measurement-set densities are used both for the recursive Bayesian belief update and for the numerical evaluation of the conditional decision probabilities required in the Bayes risk.

3.3.1. Joint Kinematic and Feature Measurements

The radar measurement process is represented through conditional measurement-set densities under the joint hypotheses defined in Equation 3.1. These densities describe the possible measurement-set realizations under each hypothesis and provide the likelihood model used for both Bayesian belief updating and Bayes Risk evaluation.

The complete single target radar measurement $\mathbf{z}_k$ is written as

$$\mathbf{z}_k = \begin{bmatrix} \mathbf{y}_k \\ \xi_k \end{bmatrix}. \tag{3.5}$$

where $\mathbf{y}_k$ contains the kinematic measurement , and ξ_k denotes the class related feature measurement.

3.3.2. Single-Measurement Target Likelihoods

For each present state hypothesis $H_{i,k}$, with $i = 1, \ldots, N_c$, the likelihood of a target generated measurement is evaluated using both kinematic and class related information. The kinematic component describes whether the measurements are consistent with the predicted target position, while the feature component describes whether the measured class related feature is consistent with the assumed target class. Combining these two sources of information provides a more discriminative likelihood than using either source individually, since a measurement must be both spatially plausible and class consistent to receive a high target originated likelihood.

The complete measurement vector therefore contains both the kinematic measurement $\mathbf{y}_k$ and the class related feature measurement ξ_k,

$$p_i(\mathbf{z}_k \mid Z_{1:k-1}) = p(\mathbf{y}_k, \xi_k \mid H_{i,k}, Z_{1:k-1}). \tag{3.6}$$

For the absent hypothesis $H_{0,k}$, the measurements are assumed to originate from clutter. The clutter density is shown as,

$$p_0(\mathbf{z}_k \mid Z_{1:k-1}) = p(\mathbf{y}_k, \xi_k \mid H_{0,k}, Z_{1:k-1}).. \tag{3.7}$$

3.3.3. False Alarm Cardinality Model

At scan k, the radar may receive a random number of measurements. Therefore, the complete radar measurement is represented as the random measurement set

$$Z_k = \{\mathbf{z}_k^{(1)}, \dots, \mathbf{z}_k^{(M_k)}\}, \tag{3.8}$$

where M_k denotes the number of measurements in the set.

Let M_{FA} denote the number of false alarm measurements inside the measurement region. The probability of receiving exactly m false alarm measurements is defined by

$$p_F(m) \triangleq P(M_{\mathrm{FA}} = m \mid Z_{1:k-1}), \qquad m = 0, 1, 2, \dots. \tag{3.9}$$

The function $p_F(m)$ describes the false alarm count distribution inside the measurement region.

Under the absent hypothesis, the total number of measurements equals the number of false alarms. Hence,

$$P(M_k = m \mid H_{0,k}, Z_{1:k-1}) = p_F(m). \tag{3.10}$$

Under a present target hypothesis, the total number of measurements depends on if the target is detected. If the target is missed, all m measurements are false alarms. If the target is detected, we assume one measurement is target generated and the remaining $m-1$ measurements are false alarms. Therefore,

$$P(M_k = m \mid H_{i,k}, Z_{1:k-1}) = (1 - P_D^{(i)}) p_F(m) + P_D^{(i)} p_F(m-1), \qquad i = 1, \dots, N_c, \tag{3.11}$$

with the convention that $p_F(-1) = 0$, as a negative amount of measurements is assumed impossible.

3.3.4. Count Weighted Measurement-Set Likelihoods

The cardinality model above only describes the probability of obtaining a measurement set with $M_k = m$ elements. To evaluate a realized measurement set, the densities of the measurement values must also be included. The resulting count weighted measurement-set likelihoods combine the probability of the required false alarm count with the density of the observed measurement values.

Under the absent hypothesis $H_{0,k}$, all measurements are clutter originated. For a realized set with $M_k = m$, this gives

$$p(Z_k, M_k = m \mid H_{0,k}, Z_{1:k-1}) = p_F(m) \prod_{\ell=1}^{m} p_0\left(\mathbf{z}_k^{(\ell)} \mid Z_{1:k-1}\right). \tag{3.12}$$

Under a present-class hypothesis $H_{i,k}$, the same two cardinality cases from Equation 3.11 are weighted by the corresponding measurement-value densities. This gives

$$\begin{aligned} &p(Z_k, M_k = m \mid H_{i,k}, Z_{1:k-1}) \\ &= (1 - P_D^{(i)}) p_F(m) \prod_{\ell=1}^{m} p_0\left(\mathbf{z}_k^{(\ell)} \mid Z_{1:k-1}\right) \\ &\quad + P_D^{(i)} p_F(m-1) \sum_{\ell=1}^{m} p_i\left(\mathbf{z}_k^{(\ell)} \mid Z_{1:k-1}\right) \prod_{\substack{r=1 \\ r \neq \ell}}^{m} p_0\left(\mathbf{z}_k^{(r)} \mid Z_{1:k-1}\right). \end{aligned} \tag{3.13}$$

for $i = 1, \dots, N_c$. The first term corresponds to a missed detection, while the second term sums over all possible associations of the single target-originated measurement.

After Z_k has been observed, its realized cardinality $M_k = m$ is part of the observation. Therefore, the count weighted densities in Equation 3.12 and Equation 3.13 can be used directly as the likelihoods in the posterior belief update.

3.4. Recursive Bayesian Belief Update

The measurement-set densities derived above are next used to evaluate the likelihoods required in the Bayesian posterior update.

During recursive filtering, the radar receives one realized measurement set,

$$Z_k = \{\mathbf{z}_k^{(1)}, \dots, \mathbf{z}_k^{(M_k)}\}, \tag{3.14}$$

for which the number of measurements is known. Let this realized cardinality be denoted by $M_k = m$.

After the measurement set has been observed, the realized cardinality is part of the observation. Therefore, the likelihood used in the recursive Bayesian update is the count weighted measurement-set density

$$L_{i,k}(Z_k) \triangleq p(Z_k, M_k = m \mid H_{i,k}, Z_{1:k-1}), \qquad i = 0, \dots, N_c. \tag{3.15}$$

This likelihood includes both the density of the received measurement values and the probability of obtaining the observed number of measurements.

The evaluated likelihood values are then inserted into the Bayesian posterior update equation to compute the posterior hypothesis probabilities. After receiving the measurement set Z_k, with realized cardinality $M_k = m$, the posterior probabilities are updated using Bayes' rule:

$$\pi_{i,k} = \frac{L_{i,k}(Z_k)\,\pi_{i,k|k-1}}{\sum_{j=0}^{N_c} L_{j,k}(Z_k)\,\pi_{j,k|k-1}}, \qquad i = 0, \dots, N_c. \tag{3.16}$$

Here, $L_{i,k}(Z_k)$ denotes the likelihood of the realized measurement set and its realized cardinality under hypothesis $H_{i,k}$, conditioned on the measurement history up to time step $k-1$.

Altogether, the first-order Markov hypothesis transition model, the Bayesian measurement update, and the radar resource allocation actions through the dwell time τ and revisit interval T can be interpreted as a myopic Partially Observable Markov Decision Process (POMDP) framework [27], [28].

3.5. Bayes Risk Decision Framework

3.5.1. Decision Set and Cost Matrix

Based on the posterior hypothesis probabilities $P(H_{i,k} \mid Z_{1:k})$, a decision is made from the decision set

$$\mathcal{D} = \{D_{0,k}, D_{1,k}, \dots, D_{N_c,k}\}, \tag{3.17}$$

where $D_{i,k}$ denotes the decision to select hypothesis $H_{i,k}$.

To account for the operational consequences of incorrect decisions, a joint cost matrix is proposed:

$$\mathbf{C}_k = \begin{bmatrix} C_{00,k} & C_{01,k} & \cdots & C_{0N_c,k} \\ C_{10,k} & C_{11,k} & \cdots & C_{1N_c,k} \\ \vdots & \vdots & \ddots & \vdots \\ C_{N_c0,k} & C_{N_c1,k} & \cdots & C_{N_cN_c,k} \end{bmatrix}, \tag{3.18}$$

where $C_{ij,k}$ denotes the cost associated with selecting decision $D_{i,k}$ when hypothesis $H_{j,k}$ is true.

Introducing this joint cost matrix allows different decision errors to be weighted according to their operational significance very directly. In practical radar systems, not all decision errors have equal consequences. For example, failing to detect a target or incorrectly classifying a hostile target may be significantly more critical than confusing two similar non-threatening target classes. By explicitly incorporating these relative penalties into the decision process, the Bayesian framework can prioritize decisions that minimize the expected operational impact rather than only maximizing classification accuracy. This provides a very direct relation between the assigned costs and the decision errors.

3.5.2. Conditional Bayes Risk and Minimum-Risk Decision Rule

The conditional Bayes risk associated with selecting decision $D_{i,k}$ is defined as the expected posterior cost [21], [29]:

$$r(D_{i,k} \mid Z_{1:k}) = \sum_{j=0}^{N_c} C_{ij,k} P(H_{j,k} \mid Z_{1:k}). \tag{3.19}$$

The minimum Bayes risk decision rule selects the decision that minimizes the conditional risk:

$$\hat{D}_k = \arg \min_{D_{i,k} \in \mathcal{D}} r(D_{i,k} \mid Z_{1:k}). \tag{3.20}$$

Although these decisions are currently not used in this framework to delete tracks, they are used to define the Bayesian decision regions.

3.5.3. Bayesian Decision Regions

The Bayesian decision rule partitions the measurement space into decision regions corresponding to the minimum risk decision. The decision region associated with decision $D_{i,k}$ is therefore defined as

$$\Gamma_i = \{Z_k : r(D_{i,k} \mid Z_{1:k}) \leq r(D_{\ell,k} \mid Z_{1:k}),\ \forall \ell \in \{0, \ldots, N_c\}\}, \tag{3.21}$$

where Γ_i denotes the region of the measurement space for which decision $D_{i,k}$ yields the minimum conditional Bayes risk.

Equivalently, the decision regions can be expressed through the pairwise inequalities

$$r(D_{i,k} \mid Z_{1:k}) \leq r(D_{\ell,k} \mid Z_{1:k}), \qquad \ell \neq i. \tag{3.22}$$

For measurement sets with variable cardinality, these regions are later indexed by the realized cardinality m, since each value of m corresponds to a different measurement value space.

3.6. Expected Bayes Risk Evaluation

3.6.1. Expected Bayes Risk Objective

The expected Bayes risk is obtained by taking the expectation of the decision cost over all possible true hypotheses and decisions. Therefore, for a single target at time k, the expected risk can be written as

$$R_k = \sum_{i=0}^{N_c} \sum_{j=0}^{N_c} C_{ij,k} P(D_{i,k} \mid H_{j,k}, Z_{1:k-1}) \pi_{j,k|k-1}. \tag{3.23}$$

where $P(D_{i,k} \mid H_{j,k}, Z_{1:k-1})$ denotes the probability of selecting decision $D_{i,k}$ when hypothesis $H_{j,k}$ is true, conditioned on the measurement history available up to time step $k-1$.

Building on the Bayes risk-based objective formulation introduced in chapter 2, the objective function $\mathcal{C}(\cdot)$ in Equation 2.7 is now specified using the total risk derived in Equation 3.23. Since this risk is evaluated before the next measurement set is observed, the resulting allocation is myopic and optimizes the expected decision cost one step ahead. This total risk combines the expected operational consequences of the joint track maintenance and classification decisions into a single objective value. Assuming that the targets are sufficiently separated and can be treated as conditionally independent, the overall system objective is modeled as the sum of the individual target risks,

$$\mathcal{C}\left(\mathbf{s}_k^n, \boldsymbol{\pi}_{k|k-1}^n, \tau_n, T_n\right) = R_k^n\left(\mathbf{s}_k^n, \boldsymbol{\pi}_{k|k-1}^n, \tau_n, T_n\right). \tag{3.24}$$

where R_k^n denotes the expected Bayes risk associated with target n.

3.6.2. Conditional Decision Probabilities

Evaluating the Bayes risk in Equation 3.23 requires the conditional decision probabilities

$$P(D_{i,k} \mid H_{j,k}, Z_{1:k-1}). \tag{3.25}$$

These probabilities describe how likely it is that the Bayesian decision rule selects decision $D_{i,k}$ when hypothesis $H_{j,k}$ is true, conditioned on the measurement history available before scan k.

In the multiple measurement case, these probabilities must account for all possible numbers of measurements and all possible target measurement associations. Let $\delta_k(Z_k, m)$ denote the Bayesian decision rule from Equation 3.20 applied to a measurement set Z_k with cardinality $M_k = m$. For each possible cardinality, the corresponding Bayesian decision region is defined as

$$\Gamma_i^{(m)} = \{Z_k : \delta_k(Z_k, m) = D_{i,k}\}. \tag{3.26}$$

Hence, $\Gamma_i^{(m)}$ contains all m-measurement sets for which the Bayesian decision rule selects decision $D_{i,k}$.

Using the measurement-set formulation from Equation 3.8, the conditional decision probability can then be written directly in terms of the count weighted measurement-set density as

$$P(D_{i,k} \mid H_{j,k}, Z_{1:k-1}) = \sum_{m=0}^{\infty} \int_{\Gamma_i^{(m)}} p(Z_k, M_k = m \mid H_{j,k}, Z_{1:k-1})\, dZ_k. \tag{3.27}$$

This expression integrates the probability density of all measurement sets that lead to decision $D_{i,k}$, while summing over all possible measurement-set cardinalities.

The count weighted form in Equation 3.27 is equivalent to first conditioning on the measurement-set cardinality. Indeed, for fixed $M_k = m$, the joint density of the measurement values and the cardinality can be factorized as

$$p(Z_k, M_k = m \mid H_{j,k}, Z_{1:k-1}) = P(M_k = m \mid H_{j,k}, Z_{1:k-1}) p(Z_k \mid H_{j,k}, M_k = m, Z_{1:k-1}). \tag{3.28}$$

Substituting Equation 3.28 into Equation 3.27 gives

$$P(D_{i,k} \mid H_{j,k}, Z_{1:k-1}) = \sum_{m=0}^{\infty} P(M_k = m \mid H_{j,k}, Z_{1:k-1}) \int_{\Gamma_i^{(m)}} p(Z_k \mid H_{j,k}, M_k = m, Z_{1:k-1})\, dZ_k. \tag{3.29}$$

Here, $P(M_k = m \mid H_{j,k}, Z_{1:k-1})$ can be taken outside the integral because it does not depend on the measurement values in Z_k. The integral itself is the probability that the Bayesian decision rule selects $D_{i,k}$, conditioned on both the true hypothesis $H_{j,k}$ and the measurement-set cardinality $M_k = m$. Hence, the count weighted expression in Equation 3.27 is equivalent to averaging the cardinality-conditioned decision probabilities over the probability mass of M_k.

In general, the decision regions $\Gamma_i^{(m)}$ are defined implicitly by the Bayesian decision rule. These regions may have complex nonlinear boundaries and therefore do not generally admit closed-form analytical integration. Numerical approximation is therefore used in this work.

3.6.3. Motivation for Numerical Approximation

In principle, the decision region integrals in Equation 3.29 can be evaluated analytically once the Bayesian decision regions are explicitly characterized. However, these regions are not defined directly, but implicitly through pairwise comparisons of conditional Bayes risks. As a result, their boundaries are generally nonlinear functions of the measurement variables.

For low dimensional problems with Gaussian likelihoods and equal covariance matrices, such decision boundaries may reduce to analytically tractable hypersurfaces. In the present joint tracking and classification formulation, however, the measurement model contains nonlinear radar measurements, class dependent likelihoods, missed detections, clutter, and uncertain data association. Consequently, the corresponding Bayesian decision regions become highly complex and generally do not admit a closed form analytical expression.

The first analytical route considered in this work was therefore to simplify or approximate the Bayesian decision regions themselves. Analytical formulations for cost sensitive Bayes risk have been investigated in the literature, but the directly applicable reductions identified in this study are primarily restricted to binary classification problems [30], [31], [32]. In the binary case, the Bayes risk can be expressed using two competing conditional-risk functions, and the corresponding decision boundary can often be characterized by a single likelihood-ratio threshold, as also shown in [17]. This structure does not extend straightforwardly to the multiclass setting considered here, where each decision region is determined by the minimum over multiple cost weighted conditional risks, as shown in Equation 3.20.

Several approximate analytical formulations were also considered to avoid direct integration over the full measurement space. In particular, hierarchical and gate based approximations were investigated, in which the joint decision process is decomposed into sequential detection and classification stages. In such a formulation, the detection decision can first be approximated by neglecting class confusion errors, leading to likelihood-ratio thresholds determined by the missed-detection and false alarm costs, such as C_{0i} and C_{i0}. Classification can then be performed only within the region where target presence is selected. This decomposition reduces the complexity of the decision problem by replacing the global Bayesian decision regions with locally defined detection and classification subregions.

However, this sequential decomposition breaks the equivalence with the global Bayes decision rule in Equation 3.20. In particular, the resulting subregions do not directly preserve the decision regions induced by the full cost matrix, since target existence and class confusion costs are treated in separate stages rather than jointly. This neglects the coupling between target existence, classification, missed detections, clutter, and data association in the complete joint measurement-set formulation. Therefore, no sufficiently accurate and general analytical approximation of the Bayesian decision regions was found.

A second alternative considered in this work was to replace the exact conditional decision probabilities by analytical bounds. A substantial body of literature provides bounds on the probability of error of the MAP decision rule. Classical examples include the Bhattacharyya and Chernoff bounds, as well as bounds based on other statistical divergence measures [30], [33]. Although many of these methods were originally developed for binary hypothesis testing problems, extensions to multiclass MAP classification have also been proposed. These extensions commonly express the multiclass error probability in terms of pairwise binary risks, pairwise class affinities, or generalized multiclass separability measures [34], [35], [36], [37].

However, these multiclass bounds generally assume a zero-one loss function, under which every incorrect decision incurs the same cost. In that case, the Bayes risk reduces to the total probability of misclassification, and the Bayes optimal decision rule becomes the MAP rule. This differs from the decision problem considered in this work, where an arbitrary Bayesian cost matrix assigns different consequences to different combinations of true and selected hypotheses. The resource management objective is therefore explicitly cost sensitive, and cannot be evaluated from the total MAP error probability alone.

Consequently, although analytical approximations are available for certain binary cost sensitive problems and analytical bounds are available for multiclass MAP error probabilities under zero-one loss, no sufficiently general and computationally applicable analytical method was found for the cost sensitive multihypothesis problem considered in this thesis. More importantly, the available methods do not directly provide the individual conditional decision probabilities required to evaluate the complete Bayesian cost matrix in Equation 3.19.

Multiclass Bayes Risk formulations with arbitrary loss functions have been presented in the literature [38]. When the resulting decision region integrals do not admit closed form expressions, these probabilities are commonly approximated numerically by sampling from the conditional distributions and evaluating the resulting classification decisions. This motivates the Monte Carlo approximation adopted in this work, which directly estimates the required conditional decision probabilities without requiring an explicit analytical characterization of the Bayesian decision regions.

3.6.4. Monte Carlo Approximation of Conditional Decision Probabilities

The decision probabilities in Equation 3.29 contain two sources of randomness. First, the number of measurements M_k is random. Second, for a fixed measurement-set cardinality $M_k = m$, the measurement values in Z_k are random. The computation is therefore separated into two steps. In this subsection, the decision probability is estimated for a fixed hypothesis $H_{j,k}$ and a fixed cardinality m. In subsection 3.6.5, these conditional estimates are then averaged over the possible values of m.

For a fixed hypothesis $H_{j,k}$ and measurement-set cardinality $M_k = m$, measurement sets are sampled from the corresponding conditional measurement-value density,

$$Z_k^{(q,j,m)} \sim p(Z_k \mid H_{j,k}, M_k = m, Z_{1:k-1}), \qquad q = 1, \ldots, N_{\mathrm{MC}}. \tag{3.30}$$

Each sampled measurement set is then passed through the same Bayesian decision procedure that would be used for a realized measurement set. For every Monte Carlo realization, the count weighted likelihood values

$$p(Z_k^{(q,j,m)}, M_k = m \mid H_{\ell,k}, Z_{1:k-1}), \qquad \ell = 0, \ldots, N_c, \tag{3.31}$$

are evaluated and inserted into the Bayesian posterior update from Equation 3.16. The resulting posterior probabilities are then used in the Bayesian decision rule from Equation 3.20. This gives the selected decision

$$\hat{D}_k^{(q,j,m)} = \delta_k\left(Z_k^{(q,j,m)}, m\right). \tag{3.32}$$

The conditional decision probability for fixed $H_{j,k}$ and $M_k = m$ is then estimated by the relative frequency with which decision $D_{i,k}$ is selected:

$$\widehat{P}(D_{i,k} \mid H_{j,k}, M_k = m, Z_{1:k-1}) = \frac{1}{N_{\mathrm{MC}}} \sum_{q=1}^{N_{\mathrm{MC}}} \mathbb{I}\Big(\delta_k(Z_k^{(q,j,m)}, m) = D_{i,k}\Big). \tag{3.33}$$

Here, N_{MC} denotes the number of Monte Carlo samples and $\mathbb{I}(\cdot)$ represents the indicator function.

This estimator can equivalently be interpreted as a Monte Carlo approximation of the decision region integral. The decision region $\Gamma_i^{(m)}$ contains all measurement sets with cardinality m for which the decision rule selects $D_{i,k}$. Hence,

$$Z_k^{(q,j,m)} \in \Gamma_i^{(m)} \qquad \Longleftrightarrow \qquad \delta_k(Z_k^{(q,j,m)}, m) = D_{i,k}. \tag{3.34}$$

Therefore,

$$\int_{\Gamma_i^{(m)}} p(Z_k \mid H_{j,k}, M_k = m, Z_{1:k-1})\, dZ_k \approx \frac{1}{N_{\mathrm{MC}}} \sum_{q=1}^{N_{\mathrm{MC}}} \mathbb{I}\Big(\delta_k(Z_k^{(q,j,m)}, m) = D_{i,k}\Big). \tag{3.35}$$

3.6.5. Cardinality Averaging and Truncation

The conditional Monte Carlo estimate in Equation 3.33 only accounts for the randomness of the measurement values given a fixed cardinality m. To obtain the full conditional decision probability given only the true hypothesis $H_{j,k}$, the result must also be averaged over the possible values of the random measurement-set cardinality M_k.

Substituting the fixed cardinality Monte Carlo estimate into the count weighted decision probability in Equation 3.29 gives

$$P(D_{i,k} \mid H_{j,k}, Z_{1:k-1}) \approx \sum_{m=0}^{\infty} P(M_k = m \mid H_{j,k}, Z_{1:k-1}) \widehat{P}\left(D_{i,k} \mid H_{j,k}, M_k = m, Z_{1:k-1}\right). \tag{3.36}$$

In practice, the infinite summation over m must be truncated at a maximum measurement-set cardinality $M_{\max}$. Using the false-alarm count probability introduced in Equation 3.9, $M_{\max}$ is selected such that

$$P(M_{\mathrm{FA}} \leq M_{\max}) \geq \gamma, \tag{3.37}$$

where γ is a probability mass threshold selected close to one. This ensures that the omitted false alarm cardinalities contain at most $1 - \gamma$ of the total false alarm probability mass.

The resulting truncated estimator combines the two steps: Monte Carlo averaging over measurement values for each fixed m, followed by probability weighted averaging over the possible cardinalities:

$$\widehat{P}(D_{i,k} \mid H_{j,k}, Z_{1:k-1}) = \sum_{m=0}^{M_{\max}} P(M_k = m \mid H_{j,k}, Z_{1:k-1}) \frac{1}{N_{\mathrm{MC}}} \sum_{q=1}^{N_{\mathrm{MC}}} \mathbb{I}\Big(\delta_k(Z_k^{(q,j,m)}, m) = D_{i,k}\Big). \quad (3.38)$$

The outer summation accounts for the probability of receiving different numbers of measurements, while the inner Monte Carlo summation estimates the probability that the Bayesian decision rule selects $D_{i,k}$ for measurement sets with that fixed cardinality. In this way, the approximation preserves the statistical effects of missed detections, multiple clutter measurements, and probabilistic measurement association while keeping the Bayes Risk evaluation tractable for adaptive RRM.

3.7. Surrogate-Based Resource Allocation

3.7.1. Solution Strategy

The Bayes Risk objective derived in the preceding sections is incorporated into the constrained resource allocation problem from Equation 2.7. The selected solution method must therefore support constrained minimization while maintaining feasibility with respect to the available radar time budget.

Lagrangian relaxation provided an elegant framework in previous works[15], [39], offering both per target decomposition and an intuitive interpretation of the dual variable . However in the current non-convex settings, the associated dual subgradient method may exhibit slow convergence and sensitivity to step size selection [40]. Maximizing the dual function yields the best lower bound on the primal objective, but does not guarantee recovery of a feasible or optimal primal solution in non-convex problems such as the objective function considered in this thesis.

For this reason, a primal method with explicit feasibility enforcement is used. This ensures that the radar time budget constraint is satisfied at every iteration.

3.7.2. Original Resource Allocation Problem

For completeness, the resource allocation problem from Equation 2.7 is restated here in terms of the individual target risks. In the general formulation, the decision variables are the dwell times τ_n and revisit intervals T_n:

$$\min_{\{\tau_n, T_n\}} \sum_{n=1}^{N} R_{n,k}(\tau_n, T_n) \quad (3.39)$$

subject to

$$\sum_{n=1}^{N} \frac{\tau_n}{T_n} \leq B_{\max}. \quad (3.40)$$

The individual risk $R_{n,k}$ also depends on the predicted target state and predicted joint belief state. These quantities are not shown explicitly in this subsection because they are not decision variables of the resource allocation problem. At a given resource allocation step, the predicted state and belief are treated as fixed inputs, while the optimization is performed only over the selected radar resource variables.

This formulation allows the Bayes Risk objective to be evaluated for different radar resource variables. In the surrogate based method developed below, the optimization is formulated for one selected scalar resource variable u_n. Depending on the assumed radar implementation, this variable may represent dwell time or revisit interval.

3.7.3. Computational Challenge of Direct Bayes Risk Optimization

Direct optimization of the Bayes Risk objective is computationally demanding because its evaluation relies on the Monte Carlo procedure from subsection 3.6.4. Consequently, each evaluation of the objective function would require a new set of sampled measurement realizations and repeated Bayesian decision evaluations.

This introduces two major difficulties. First, the estimated risk values contain sampling noise, such that the resulting risk curve is no longer perfectly smooth. Small fluctuations may appear even when the true underlying risk function is monotonic. Second, during iterative optimization, the objective function must be evaluated repeatedly at different candidate allocations. If a new Monte Carlo simulation is required at every iteration, the computational cost becomes very big.

Thus, the optimization problem becomes both expensive and numerically unstable. The combination of noisy objective values and finite difference gradient approximations leads to unreliable descent directions and slow or no convergence. To overcome these issues, a deterministic surrogate approximation of the Bayes risk function is proposed.

3.7.4. Surrogate Risk Construction

To avoid repeated Monte Carlo evaluation during the iterative optimization, a deterministic surrogate risk function is constructed for each target. The Bayes risk is evaluated on a predefined grid of the selected resource variable u_n, after which the sampled values are processed and interpolated to obtain a continuous approximation of the curve.

For every target n, a resource grid is defined as

$$u_n \in \{u_1, u_2, \dots, u_{N_u}\}. \tag{3.41}$$

At each grid point u_j, the corresponding Bayes risk is evaluated while the remaining resource variables are kept fixed:

$$R_{n,k}(u_j). \tag{3.42}$$

Since these values are obtained through Monte Carlo simulation, the resulting sampled curve consists of discrete noisy estimates of the true underlying risk function with respect to the selected resource variable.

The sampled Monte Carlo risk values are postprocessed in order to obtain a numerically stable approximation:

First, a Savitzky-Golay smoothing filter [41] is applied to suppress high frequency sampling fluctuations, while preserving the overall curvature of the risk curve.

Second, a monotonicity constraint is enforced using isotonic regression [42]. The monotonicity direction is chosen according to the physical interpretation of the selected resource variable. If the resource variable is expected to improve sensing performance when increased, such as dwell time, the surrogate risk is constrained to be non increasing. If increasing the variable is expected to reduce information availability, such as for the revisit interval, the surrogate risk is constrained to be non decreasing. This regularization suppresses Monte Carlo fluctuations that violate the expected physical trend of the Bayes Risk function.

After the regularization, the processed samples are converted into a continuous surrogate function using piecewise cubic Hermite interpolation (PCHIP) [43]. This interpolation method preserves the local shape of the data and avoids the oscillatory behavior that may happen with standard cubic splines, while maintaining the monotonic trend of the processed samples [44].

The final surrogate risk function for target n is denoted by

$$\tilde{R}_{n,k}(u_n), \tag{3.43}$$

which provides a smooth and computationally efficient approximation of the original Monte Carlo Bayes risk.

3.7.5. Scalar Resource Optimization

After the surrogate risk functions have been constructed, the stochastic Bayes Risk optimization problem is replaced by a deterministic scalar resource optimization problem. Restricting the optimization to one scalar resource variable per target avoids the additional complexity of constructing and optimizing a multidimensional surrogate risk function. The resulting optimization can then be performed directly on the one dimensional surrogate risk functions.

Using the surrogate risk approximation, the scalar optimization problem becomes

$$\min_{\{u_n\}} \sum_{n=1}^{N} \tilde{R}_{n,k}(u_n), \tag{3.44}$$

where $\tilde{R}_{n,k}(u_n)$ denotes the surrogate Bayes Risk function for target n as a function of the selected resource variable.

The corresponding feasible set depends on the interpretation of u_n. If u_n is chosen as the dwell time, the revisit intervals are fixed and the radar time budget constraint becomes

$$\sum_{n=1}^{N} \frac{u_n}{T} \leq B_{\max}. \tag{3.45}$$

If u_n is chosen as the revisit interval, the dwell times are fixed and the budget constraint becomes

$$\sum_{n=1}^{N} \frac{\tau}{u_n} \leq B_{\max}. \tag{3.46}$$

The projected optimization method below is written for a generic feasible set $\mathcal{U}$. Therefore, the same optimization structure can be used for any selected scalar resource variable, provided that projection onto $\mathcal{U}$ can be performed efficiently.

The constrained scalar resource problem is solved using a projected subgradient method[40]. At iteration ℓ, the resource allocation vector is written as

$$\mathbf{u}^{\ell} = \begin{bmatrix} u_1^{\ell} & \cdots & u_N^{\ell} \end{bmatrix}^{\top}. \tag{3.47}$$

The vector is first updated in the negative subgradient direction:

$$\mathbf{u}^{\ell+\frac{1}{2}} = \mathbf{u}^{\ell} - \rho_{\ell}\mathbf{g}^{\ell}, \tag{3.48}$$

where $\mathbf{g}^{\ell}$ is a subgradient of the surrogate objective function and ρ_{ℓ} is the step size.

To ensure that the radar time budget constraint remains satisfied after each update, the intermediate iterate is subsequently projected back onto the feasible set:

$$\mathbf{u}^{\ell+1} = \Pi_{\mathcal{U}}\left(\mathbf{u}^{\ell+\frac{1}{2}}\right), \tag{3.49}$$

where $\mathcal{U}$ denotes the feasible set from the selected scalar resource variable.

This projection enforces the maximum budget constraint and the imposed lower and upper bounds on the selected resource variable at every iteration. Due to the projection, every accepted iterate will result in a feasible radar resource allocation.

Since the surrogate objective does not explicitly provide analytical derivatives, the subgradient is approximated numerically using finite differences. For each target n, the corresponding component of the subgradient is approximated as

$$g_n^{\ell} \approx \frac{\tilde{R}_{n,k}(u_n^{\ell} + \epsilon) - \tilde{R}_{n,k}(u_n^{\ell})}{\epsilon}, \tag{3.50}$$

where ϵ is a small perturbation chosen such that $u_n^{\ell} + \epsilon$ remains within the predefined resource interval.

Although the feasible set of the scalar resource optimization problem may be convex for a suitable choice of u_n, the surrogate Bayes Risk functions are known to be non convex. Consequently, the projected subgradient method may converge to different stationary solutions depending on the initial resource allocation. To reduce the sensitivity of the solution to the initialization, the solver is executed from multiple deterministic starting points.

The first initialization is the previously selected resource allocation. This provides a warm start based on the allocation obtained during the preceding resource allocation step and is expected to be particularly suitable when the target states and belief probabilities change gradually over time.

The second initialization distributes the available radar time budget equally among all targets. In addition, a set of target prioritized initializations is considered. For each target, one initialization assigns as much of the available resource budget as permitted to that target, while the remaining budget is distributed among the other targets. These initializations are referred to as single heavy initializations.

Each initialization is projected onto the feasible set before the iterative optimization begins. The projected subgradient solver is then run independently from every starting point. After convergence, the total surrogate Bayes risk is evaluated for each resulting feasible allocation. The final resource allocation is selected according to

$$\mathbf{u}^{\star} = \underset{\mathbf{u}^{(j)} \in \left\{\mathbf{u}^{(1)},\ldots,\mathbf{u}^{(N_{\text{init}})}\right\}}{\arg\min} \sum_{n=1}^{N} \tilde{R}_{n,k}\left(u_n^{(j)}\right). \tag{3.51}$$

Because the surrogate Bayes Risk functions are non convex, this multi start procedure does not guarantee recovery of the globally optimal allocation. It only selects the best feasible allocation among the stationary solutions found from the considered deterministic initializations. The resulting solution should therefore be interpreted as a practical locally optimized allocation, rather than as a global optimum.

3.7.6. Remark on Multi Variable Resource Optimization

The scalar resource formulation optimizes one resource variable per target while keeping the remaining resource variables fixed. This means that the method can be applied to dwell time optimization or revisit time optimization, but it does not optimize both simultaneously.

A full joint dwell and revisit formulation would require a surrogate risk function of the form

$$\tilde{R}_{n,k}(\tau_n, T_n), \tag{3.52}$$

and the resulting optimization problem would be

$$\min_{\{\tau_n, T_n\}} \sum_{n=1}^{N} \tilde{R}_{n,k}(\tau_n, T_n) \tag{3.53}$$

subject to

$$\sum_{n=1}^{N} \frac{\tau_n}{T_n} \leq B_{\max}. \tag{3.54}$$

Such a formulation would allow the radar to adapt both how long each target is observed and how often each target is revisited. However, it would require constructing a multidimensional surrogate risk function and solving an optimization problem with nonlinear coupling between the decision variables through the ratio τ_n/T_n. The reliable combination of such joint optimization with the Monte Carlo-based surrogate risk construction used in this thesis is left as future work.

4 Assumed Scenario

This chapter presents the modeling assumptions used in the final simulations. The goal is to model a radar system that performs both target tracking and target classification under limited sensing resources.

The radar is assumed to dynamically allocate dwell time to multiple targets, while the revisit time is kept fixed for all targets and class hypotheses. This choice keeps all targets on a common discrete time update grid and avoids the need for an explicit scheduling model that determines when each target is revisited. As a result, the adaptive resource allocation can be studied independently from radar timeline scheduling constraints. The allocated dwell time affects the signal to noise ratio (SNR), which in turn influences the probability of detection and the quality of both the kinematic and feature measurements. The resulting measurements are subsequently used for recursive kinematic state estimation, classification, and Bayes Risk-based resource allocation.

Throughout this chapter, superscript (i) shows dependence on the present state class hypothesis $H_{i,k}$, with $i = 1, \ldots, N_c$. The absent hypothesis is denoted by $H_{0,k}$. The time index k is included when a quantity changes over time, for example through the range and dwell time dependent SNR. For readability, the target index n is omitted unless it is needed explicitly.

4.1. Target Motion Model

The target motion is modeled using a discrete time constant velocity linear state space model. Conditioned on present state class hypothesis $H_{i,k}$, the target dynamics evolve according to

$$\mathbf{s}_{k+1} = F\mathbf{s}_k + \mathbf{w}_k^{(i)}, \tag{4.1}$$

where the state vector is defined as

$$\mathbf{s}_k = \begin{bmatrix} x_k & y_k & \dot{x}_k & \dot{y}_k \end{bmatrix}^\top. \tag{4.2}$$

The state-transition matrix is given by

$$F = \begin{bmatrix} 1 & 0 & T & 0 \\ 0 & 1 & 0 & T \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{4.3}$$

where T denotes the revisit time. With a fixed revisit time, all targets are updated on a common discrete time grid and the state transition matrix F can be evaluated using the same value of T.

The process noise is assumed Gaussian,

$$\mathbf{w}_k^{(i)} \sim \mathcal{N}\left(\mathbf{0}, Q^{(i)}\right), \tag{4.4}$$

with class dependent process noise covariance

$$Q^{(i)} = \left(\sigma_{\omega}^{(i)}\right)^2 \begin{bmatrix} \frac{T^4}{4} & 0 & \frac{T^3}{2} & 0 \\ 0 & \frac{T^4}{4} & 0 & \frac{T^3}{2} \\ \frac{T^3}{2} & 0 & T^2 & 0 \\ 0 & \frac{T^3}{2} & 0 & T^2 \end{bmatrix}. \tag{4.5}$$

Here, $\sigma_{\omega}^{(i)}$ denotes the class dependent maneuverability standard deviation. The process noise accounts for unmodeled target accelerations and deviations from ideal constant-velocity motion.

4.2. Radar Measurement Model

The considered radar system is assumed to provide both kinematic measurements and target dependent feature measurements. The kinematic measurements are used for target tracking, while the feature measurements provide additional information for target classification.

The radar directly measures the target range and azimuth angle. The nonlinear kinematic measurement function is

$$h(\mathbf{s}_k) = \begin{bmatrix} \sqrt{x_k^2 + y_k^2} \\ \arctan 2(y_k, x_k) \end{bmatrix}. \tag{4.6}$$

The first component corresponds to the target range, while the second component corresponds to the azimuth angle.

The kinematic measurement under hypothesis $H_{i,k}$ is modeled as

$$\mathbf{y}_k = h(\mathbf{s}_k) + \mathbf{v}_{y,k}^{(i)}, \tag{4.7}$$

where $\mathbf{y}_k$ denotes the range-azimuth measurement vector and the kinematic measurement noise is additive zero mean Gaussian,

$$\mathbf{v}_{y,k}^{(i)} \sim \mathcal{N}\left(\mathbf{0}, R_k^{(i)}\right). \tag{4.8}$$

The corresponding kinematic measurement noise covariance is

$$R_k^{(i)} = \begin{bmatrix} \left(\sigma_{r,k}^{(i)}\right)^2 & 0 \\ 0 & \left(\sigma_{\theta,k}^{(i)}\right)^2 \end{bmatrix}. \tag{4.9}$$

The range and azimuth measurement errors are assumed statistically independent.

Under class hypothesis $H_{i,k}$, the feature measurement is modeled as

$$\xi_k = \mu_i + v_{\xi,k}^{(i)}, \tag{4.10}$$

where μ_i denotes the nominal feature value associated with class i, and $v_{\xi,k}^{(i)}$ represents the feature measurement noise. The feature noise is assumed additive zero mean Gaussian,

$$v_{\xi,k}^{(i)} \sim \mathcal{N}\left(0, \left(\sigma_{\xi,k}^{(i)}\right)^2\right). \tag{4.11}$$

Each validated radar measurement is written as the combined kinematic feature measurement

$$\mathbf{z}_k = \begin{bmatrix} \mathbf{y}_k \\ \xi_k \end{bmatrix}, \qquad m = 1, \dots, M_k, \tag{4.12}$$

where M_k denotes the number of validated measurements at time step k.

4.3. Measurement Quality

The quality of the radar measurements is assumed to depend on the allocated dwell time and the target range. This dependency is modeled through the SNR. For class hypothesis $H_{i,k}$, the predicted target range is obtained from the class conditioned predicted state estimate as

$$\hat{r}^{(i)}_{k|k-1} = \sqrt{\left(\hat{x}^{(i)}_{k|k-1}\right)^2 + \left(\hat{y}^{(i)}_{k|k-1}\right)^2}. \tag{4.13}$$

The SNR is modeled as [45]

$$\mathrm{SNR}^{(i)}_k = \mathrm{SNR}_0 \left(\frac{\tau}{\tau_0}\right) \left(\frac{r_0}{\hat{r}^{(i)}_{k|k-1}}\right)^4, \tag{4.14}$$

where τ_k denotes the allocated dwell time, τ_0 denotes the reference dwell time, r_0 denotes the reference range, and SNR_0 is the reference SNR.

The probability of detection is modeled using a Swerling-I target fluctuation model [46],

$$P^{(i)}_{D,k} = P_F^{\frac{1}{1+\mathrm{SNR}^{(i)}_k}}, \tag{4.15}$$

where P_F denotes the constant false-alarm probability.

An increase in SNR is assumed to improve the radar measurement accuracy. The baseline noise parameters are allowed to depend on the target class. The resulting time and class dependent measurement standard deviations are modeled as [45]

$$\sigma^{(i)}_{r,k} = \frac{\sigma^{(i)}_{r,0}}{\sqrt{\mathrm{SNR}^{(i)}_k}}, \qquad \sigma^{(i)}_{\theta,k} = \frac{\sigma^{(i)}_{\theta,0}}{\sqrt{\mathrm{SNR}^{(i)}_k}}, \qquad \sigma^{(i)}_{\xi,k} = \frac{\sigma^{(i)}_{\xi,0}}{\sqrt{\mathrm{SNR}^{(i)}_k}}. \tag{4.16}$$

Here, $\sigma^{(i)}_{r,0}$, $\sigma^{(i)}_{\theta,0}$, and $\sigma^{(i)}_{\xi,0}$ denote the class dependent reference standard deviations for range, azimuth, and feature measurements, respectively. The quantities $\sigma^{(i)}_{r,k}$, $\sigma^{(i)}_{\theta,k}$, and $\sigma^{(i)}_{\xi,k}$ denote standard deviations; their squared values are used as variances in the corresponding covariance matrices.

Consequently, larger dwell times generally lead to improved detection performance and more accurate tracking and classification measurements.

4.4. Class Conditioned Predicted Measurement

The class conditioned predicted state is assumed Gaussian, as maintained by the class conditioned EKF prediction step. To be able to define the single measurement likelihoods, the corresponding predicted measurement quantities have to be defined. For each present state class hypothesis $H_{i,k}$, the predicted state distribution is written as

$$p(\mathbf{s}_k \mid H_{i,k}, Z_{1:k-1}) = \mathcal{N}\left(\hat{\mathbf{s}}^{(i)}_{k|k-1}, P^{(i)}_{k|k-1}\right). \tag{4.17}$$

The prediction step is performed using the constant velocity motion model,

$$\hat{\mathbf{s}}^{(i)}_{k|k-1} = F\hat{\mathbf{s}}^{(i)}_{k-1|k-1}, \tag{4.18}$$

and

$$P^{(i)}_{k|k-1} = FP^{(i)}_{k-1|k-1}F^\top + Q^{(i)}. \tag{4.19}$$

For each class hypothesis, the predicted kinematic radar measurement is then obtained as

$$\hat{\mathbf{y}}^{(i)}_{k|k-1} = h\left(\hat{\mathbf{s}}^{(i)}_{k|k-1}\right). \tag{4.20}$$

Because the range-azimuth measurement model is nonlinear, the predicted kinematic measurement distribution is approximated by linearizing $h(\cdot)$ around the predicted state estimate.

The class conditioned measurement Jacobian is denoted by $H_k^{(i)}$ and is defined as

$$H_k^{(i)} = \left. \frac{\partial h(\mathbf{s})}{\partial \mathbf{s}} \right|_{\mathbf{s}=\hat{\mathbf{s}}_{k|k-1}^{(i)}} . \tag{4.21}$$

For the state vector $\mathbf{s}_k = [x_k,\, y_k,\, \dot{x}_k,\, \dot{y}_k]^\top$, this gives

$$H_k^{(i)} = \begin{bmatrix} \frac{x_k}{\sqrt{x_k^2+y_k^2}} & \frac{y_k}{\sqrt{x_k^2+y_k^2}} & 0 & 0 \\ -\frac{y_k}{x_k^2+y_k^2} & \frac{x_k}{x_k^2+y_k^2} & 0 & 0 \end{bmatrix}_{\mathbf{s}_k=\hat{\mathbf{s}}_{k|k-1}^{(i)}} . \tag{4.22}$$

This matrix maps uncertainty in the Cartesian state estimate to uncertainty in the range-azimuth measurement space.

The class conditioned innovation covariance is given by

$$S_k^{(i)} = H_k^{(i)} P_{k|k-1}^{(i)} \left(H_k^{(i)} \right)^\top + R_k^{(i)} . \tag{4.23}$$

The covariance $S_k^{(i)}$ defines the predicted uncertainty of the kinematic measurement under class hypothesis $H_{i,k}$.

These class conditioned predicted measurement quantities are used both to define the measurement likelihoods in section 4.5 and to perform the PDA-EKF update in section 4.6.

4.5. Measurement Likelihoods and Validation Region

This section specifies the single measurement likelihoods under the present target and absent target hypotheses. The present target likelihoods are first defined using the Gaussian measurement assumptions. These likelihoods are then used to construct a common effective validation region. Finally, the absent hypothesis clutter density and clutter-count model are defined using the resulting validation volume.

4.5.1. Present Target Measurement Likelihood

For each present target hypothesis $H_{i,k}$, with $i = 1, \ldots, N_c$, the target generated likelihood of validated measurement $\mathbf{z}_k$ is denoted by

$$p_i(\mathbf{z}_k \mid Z_{1:k-1}). \tag{4.24}$$

In the assumed scenario, the kinematic measurement and feature measurement are conditionally independent given the class hypothesis and the prior information $Z_{1:k-1}$. Therefore,

$$p_i(\mathbf{z}_k \mid Z_{1:k-1}) = p(\mathbf{y}_k \mid H_{i,k}, Z_{1:k-1}) p(\xi_k \mid H_{i,k}, Z_{1:k-1}). \tag{4.25}$$

Using the Gaussian measurement noise assumptions, the kinematic likelihood is given by

$$p(\mathbf{y}_k \mid H_{i,k}, Z_{1:k-1}) = \mathcal{N}\left(\mathbf{y}_k; \hat{\mathbf{y}}_{k|k-1}^{(i)}, S_k^{(i)} \right), \tag{4.26}$$

and the feature likelihood is given by

$$p(\xi_k \mid H_{i,k}, Z_{1:k-1}) = \mathcal{N}\left(\xi_k; \mu_i, \left(\sigma_{\xi,k}^{(i)} \right)^2 \right). \tag{4.27}$$

Consequently, the target generated single measurement likelihood becomes

$$p_i(\mathbf{z}_k \mid Z_{1:k-1}) = \mathcal{N}\left(\mathbf{y}_k; \hat{\mathbf{y}}_{k|k-1}^{(i)}, S_k^{(i)} \right) \mathcal{N}\left(\xi_k; \mu_i, \left(\sigma_{\xi,k}^{(i)} \right)^2 \right). \tag{4.28}$$

Equivalently, this likelihood can be written as a Gaussian density in the joint measurement space,

$$p_i(\mathbf{z}_k \mid Z_{1:k-1}) = \mathcal{N}\left(\mathbf{z}_k; \hat{\mathbf{z}}_{k|k-1}^{(i)}, \Sigma_k^{(i)} \right), \tag{4.29}$$

with

$$\hat{\mathbf{z}}_{k|k-1}^{(i)} = \begin{bmatrix} \hat{\mathbf{y}}_{k|k-1}^{(i)} \\ \mu_i \end{bmatrix}, \qquad \Sigma_k^{(i)} = \begin{bmatrix} S_k^{(i)} & 0 \\ 0 & \left(\sigma_{\xi,k}^{(i)} \right)^2 \end{bmatrix}. \tag{4.30}$$

The joint Gaussian form in Equation 4.29 is used to construct the common validation region.

4.5.2. Moment-Matched Validation Region

A common effective validation region is introduced in the joint measurement space of range, azimuth, and the class related feature. The purpose of this region is to determine, before the measurement set is processed, which measurements are retained as plausible target originated measurements. Since the target class is uncertain, a validation region based on only one class conditioned prediction could reject measurements that are plausible under another class hypothesis. The validation region should therefore cover the dominant probability mass of the relevant present target class hypotheses.

This validation step cannot be based directly on the Bayesian decision regions introduced in chapter 3. These decision regions depend on the realized measurement set and, in particular, on the number of validated measurements received in the scan. As a result, they are not fixed pre-measurement gates and cannot be used to define a constant validation volume before the measurement set is observed.

One possible approach would be to construct separate class conditioned validation regions and use their combined region. However, the corresponding validation volume would then require accounting for the overlap between the class conditioned regions in the joint measurement space. This overlap volume is not analytically tractable in the present setting. Since the clutter density requires a single validation volume, a common effective validation region is used instead.

The common effective validation region is constructed by moment-matching [47] the class conditioned joint Gaussian measurement distributions in Equation 4.29. This yields one validation gate and one corresponding volume $V_k^{(0)}$, while still retaining measurements that are plausible under the dominant present target class hypotheses. The subsequent likelihood evaluation remains class conditioned through the individual single measurement densities $p_i(\mathbf{z}_k \mid Z_{1:k-1})$.

The normalized predicted class weights are defined as

$$w_{i,k} = \frac{\pi_{i,k|k-1}}{\sum_{j=1}^{N_c} \pi_{j,k|k-1}}, \qquad i = 1, \ldots, N_c. \tag{4.31}$$

These weights represent the relative predicted probability of each present target class hypothesis, conditioned on the target being present.

The effective predicted measurement is approximated as the class probability-weighted mean

$$\hat{\mathbf{z}}_{k|k-1}^{(0)} = \sum_{i=1}^{N_c} w_{i,k} \hat{\mathbf{z}}_{k|k-1}^{(i)}. \tag{4.32}$$

Similarly, the effective validation covariance is approximated as

$$\Sigma_k^{(0)} = \sum_{i=1}^{N_c} w_{i,k} \left[\Sigma_k^{(i)} + \left(\hat{\mathbf{z}}_{k|k-1}^{(i)} - \hat{\mathbf{z}}_{k|k-1}^{(0)} \right) \left(\hat{\mathbf{z}}_{k|k-1}^{(i)} - \hat{\mathbf{z}}_{k|k-1}^{(0)} \right)^\top \right]. \tag{4.33}$$

The common effective validation region can then be defined as the ellipsoidal gate

$$\mathcal{G}_k = \left\{ \mathbf{z}_k : \left(\mathbf{z}_k - \hat{\mathbf{z}}_{k|k-1}^{(0)} \right)^\top \left(\Sigma_k^{(0)} \right)^{-1} \left(\mathbf{z}_k - \hat{\mathbf{z}}_{k|k-1}^{(0)} \right) \le F_{\chi_3^2}^{-1}(P_G) \right\}. \tag{4.34}$$

The corresponding effective validation volume is

Gating Cit

$$V_k^{(0)} = \frac{4}{3}\pi \left(F_{\chi_3^2}^{-1}(P_G) \right)^{3/2} \sqrt{\left| \Sigma_k^{(0)} \right|}. \tag{4.35}$$

The gate probability P_G is chosen sufficiently large such that the dominant probability mass of the relevant class conditioned measurement distributions is contained in the effective region.

4.5.3. Absent Hypothesis Clutter Model

For the absent hypothesis $H_{0,k}$, validated measurements are assumed to be clutter-originated. The clutter density is modeled as uniform over the common effective validation region,

$$p_0(\mathbf{z}_k \mid Z_{1:k-1}) = \frac{1}{V_k^{(0)}}. \tag{4.36}$$

This density assigns equal likelihood to all validated measurements inside the common effective validation region.

The number of clutter measurements inside the common effective validation region is modeled as a Poisson random variable, which is a standard assumption in target tracking and probabilistic data association for representing randomly occurring false alarms in clutter [48]. Thus,

$$M_{\mathrm{FA},k} \sim \mathrm{Poisson}\left(\lambda_{\mathrm{FA},k}\right), \tag{4.37}$$

with

$$\lambda_{\mathrm{FA},k} = \beta_{\mathrm{FA}} V_k^{(0)}. \tag{4.38}$$

Here, β_{FA} denotes the false alarm density in the considered measurement space, and $V_k^{(0)}$ denotes the volume of the common effective validation region. The corresponding false alarm count probability is

$$p_F(m) = P(M_{\mathrm{FA},k} = m \mid Z_{1:k-1}) = \frac{\lambda_{\mathrm{FA},k}^m e^{-\lambda_{\mathrm{FA},k}}}{m!}. \tag{4.39}$$

Together, the uniform density and the Poisson cardinality model define the clutter model used in the recursive filtering and measurement-set likelihood evaluation.

4.6. Recursive Filtering and Probabilistic Data Association

The measurement-set densities introduced in Equation 3.13 define the probabilistic measurement model used throughout the recursive estimation and decision process. As described in subsection 4.5.2, the radar measurements are first validated using the common effective validation region $\mathcal{G}_k$. The recursive filter therefore operates on the validated measurement set

$$Z_k = \left\{\mathbf{z}_k^{(1)}, \ldots, \mathbf{z}_k^{(M_k)}\right\}, \tag{4.40}$$

where each validated measurement may be either target originated or clutter originated.

Because the radar measurement model in Equation 4.6 is nonlinear, an Extended Kalman Filter (EKF) is used. Since the origin of each validated measurement is unknown, probabilistic data association (PDA) is used to account for measurement-origin uncertainty [48]. In the joint tracking and classification framework considered here, a separate PDA-EKF is run for each present state class hypothesis $H_{i,k}$, with $i = 1, \ldots, N_c$. Each class conditioned filter maintains its own kinematic state estimate and covariance.

In the PDA framework, each validated measurement is assigned an association probability. In this work, the association likelihood includes both the kinematic measurement and the feature measurement through the target generated likelihood in Equation 4.28. For notational compactness, this likelihood is written as

$$L_{i,k}^{(m)} \triangleq p_i(\mathbf{z}_k^{(m)} \mid Z_{1:k-1}). \tag{4.41}$$

The class dependent association probability of validated measurement m is approximated by

$$\beta_{i,k}^{(m)} = \frac{P_{D,k}^{(i)} L_{i,k}^{(m)}}{\lambda_{\mathrm{FA},k} + P_{D,k}^{(i)} \sum_{\ell=1}^{M_k} L_{i,k}^{(\ell)}}, \qquad m = 1, \ldots, M_k, \tag{4.42}$$

where $P_{D,k}^{(i)}$ is the detection probability under hypothesis $H_{i,k}$, and $\lambda_{\mathrm{FA},k}$ represents the expected number of false alarms in the validation region. The no-association probability is given by

$$\beta_{i,k}^{(0)} = \frac{\lambda_{\mathrm{FA},k}}{\lambda_{\mathrm{FA},k} + P_{D,k}^{(i)} \sum_{\ell=1}^{M_k} L_{i,k}^{(\ell)}}. \tag{4.43}$$

For every validated measurement $\mathbf{z}_k^{(m)}$, the corresponding kinematic innovation under hypothesis $H_{i,k}$ is

$$\boldsymbol{\nu}_{i,k}^{(m)} = \mathbf{y}_k^{(m)} - \hat{\mathbf{y}}_{k|k-1}^{(i)}. \tag{4.44}$$

The corresponding EKF gain is

$$K_k^{(i)} = P_{k|k-1}^{(i)} \left(H_k^{(i)} \right)^\top \left(S_k^{(i)} \right)^{-1}. \tag{4.45}$$

For each class conditioned filter, the combined PDA innovation is computed as the weighted sum of the validated kinematic innovations,

$$\bar{\boldsymbol{\nu}}_{i,k} = \sum_{m=1}^{M_k} \beta_{i,k}^{(m)} \boldsymbol{\nu}_{i,k}^{(m)}. \tag{4.46}$$

Although the association probabilities are computed using both the kinematic and feature components of the measurement, only the kinematic innovation vector is used in the EKF correction, since the tracked state vector contains only the position and velocity components.

The class conditioned posterior state estimate is updated according to

$$\hat{\mathbf{s}}_{k|k}^{(i)} = \hat{\mathbf{s}}_{k|k-1}^{(i)} + K_k^{(i)} \bar{\boldsymbol{\nu}}_{i,k}. \tag{4.47}$$

The corresponding covariance update consists of the usual EKF covariance reduction step, but corrected for association uncertainty:

$$P_{k|k}^{(i)} = P_{k|k-1}^{(i)} - \left(1 - \beta_{i,k}^{(0)} \right) K_k^{(i)} S_k^{(i)} \left(K_k^{(i)} \right)^\top + K_k^{(i)} C_{\nu,i,k} \left(K_k^{(i)} \right)^\top, \tag{4.48}$$

where

$$C_{\nu,i,k} = \sum_{m=1}^{M_k} \beta_{i,k}^{(m)} \boldsymbol{\nu}_{i,k}^{(m)} \left(\boldsymbol{\nu}_{i,k}^{(m)} \right)^\top - \bar{\boldsymbol{\nu}}_{i,k} \bar{\boldsymbol{\nu}}_{i,k}^\top \tag{4.49}$$

is the spread of the validated kinematic innovations due to association uncertainty under hypothesis $H_{i,k}$.

5 Simulations

5.1. Simulation setup

This chapter evaluates the proposed Bayes risk-based RRM framework in a multi-target radar scenario. The goal of the simulations is to investigate how the adaptive allocation strategy responds to different operational cost structures and target presence conditions.

Throughout the simulations, N targets are observed within the surveillance region. Each target evolves independently and is assigned a fixed class during the full observation period. The radar system receives both target originated detections and clutter generated measurements, after which the joint target existence and classification beliefs are updated according to the filtering equations introduced in the previous chapters.

All scenarios use the same underlying radar, filtering, and optimization settings unless stated otherwise. Differences in the simulation results can therefore be attributed to the specific scenario variation, such as changes in the user cost matrix, range dependent costs, or the appearance and disappearance of targets. Although the proposed framework can have class dependent model parameters, such as different measurement noise levels, these parameters are kept the same for all classes. This choice isolates the effect of the operational cost structure on the radar budget allocation behavior.

5.2. Parameter settings

The simulations are performed using a common set of radar, filtering, and optimization parameters throughout all considered scenarios. Since the underlying models and parameter definitions were introduced extensively in the previous chapters, only the most relevant numerical settings used during the simulations are summarized here.

The simulation observation period consists of $K = 20$ discrete timesteps with a sampling interval of $T = 5$. Resource allocation optimization is also performed every $T_{\text{opt}} = 5$ seconds. The total normalized radar time budget is fixed to $B_{\max} = 1$.

Table 5.1: General simulation parameters.

Parameter	Value
T	5
K	20
T_{opt}	5
$B_{\max}$	1
SNR_0	1
r_0	50×10^3

Clutter is modeled using the uniform Poisson process inside the validation gate. A high gate probability P_G is selected to ensure that target originated measurements remain inside the validation region with high probability.

Table 5.2: Detection and clutter parameters.

Parameter	Value
β_{FA}	10^{-8}
P_G	0.99999997
P_F	10^{-4}

The optimization parameters are selected empirically to ensure stable convergence of the projected iterative updates.

Table 5.3: Optimization parameters.

Parameter	Value
ρ_{base}	0.5
ϵ	10^{-10}
$N_{\mathrm{iter,max}}$	5000

The Bayes risk functions are estimated using Monte Carlo simulations. A fixed set of Monte Carlo samples is reused during optimization and for the comparison between optimized and equal budget allocation. This use of common random numbers reduces stochastic variation between risk evaluations, making differences in realized risk more directly attributable to the allocation strategy rather than to Monte Carlo sampling noise.

For the surrogate risk construction, the dwell time grid is evaluated over the interval $[0.01, 7]\,\mathrm{s}$. The lower bound is chosen slightly above zero because a dwell time of exactly zero would lead to an ill-defined radar measurement model. In particular, the SNR becomes zero, which would make the detection probability and measurement noise expressions numerically problematic. The upper bound of $7\,\mathrm{s}$ is chosen a bit larger than the maximum feasible dwell time used during optimization. Since the revisit interval is fixed at $T = 5\,\mathrm{s}$ and the normalized budget satisfies $B_{\max} = 1$, a single target can receive at most $5\,\mathrm{s}$ of dwell time within the feasible allocation. The additional grid range above $5\,\mathrm{s}$ can thus not be used as a feasible allocation region, but is included to make the smoothing and interpolation near the upper feasible boundary less sensitive to Monte Carlo noise at the final grid point.

Table 5.4: Monte Carlo parameters.

Parameter	Value
N_{mc}	5000
τ grid	200 uniformly spaced values on $[0.01, 7]\,\mathrm{s}$
Feasible τ interval	$[0.01, 5]\,\mathrm{s}$

The considered target model consists of two possible target classes. The target class feature distributions are modeled using Gaussian class conditionals with class means given by μ_c.

Table 5.5: Target and class parameters.

Parameter	Value
C	2
m	2
μ_ξ	$[1\ \ 2]^T$
π_{birth}	$\left[0, \frac{1}{2}, \frac{1}{2}\right]^T$

Target existence and classification beliefs evolve according to the joint state first-order Markov transition model. The Markov probabilities for class switching $P_{1|2}$ and $P_{2|1}$ are set on 0, as this is assumed impossible. Also the target birth probabilities $P_{1|0}$ and $P_{2|0}$ are set to 0, as we assume that new targets will come from an external searching and detection algorithm and do not have to be picked up by the joint track maintenance and classification algorithm.

Table 5.6: Joint state Markov transition matrix.

$\mathbf{T}$
$\begin{bmatrix} 1.0 & 0.02 & 0.02 \\ 0.0 & 0.98 & 0.0 \\ 0.0 & 0.0 & 0.98 \end{bmatrix}$

Finally, the radar measurement noise and class feature noise settings are summarized in Table 5.7. The process noise σ_w is set to zero in order to isolate the effects of measurement uncertainty and adaptive radar resource allocation.

Table 5.7: Noise parameters.

Parameter	Value
$\sigma_{r,0}$	20
$\sigma_{\theta,0}$	$1°$
$\sigma_{\xi,0}$	15
σ_w	0

5.3. Performance metrics

The performance of the proposed RRM framework is evaluated using the realized Bayes risk. The realized risk directly quantifies the expected operational cost associated with the decisions made by the radar system at each timestep.

The realized risk is selected as the primary performance metric because the proposed framework is itself formulated as a Bayes risk minimization problem. Both the tracking uncertainty and classification uncertainty are already incorporated into the cost aware Bayesian objective function through the operational cost matrices introduced in Chapter 3. As a result, additional performance measures such as classification accuracy, tracking root-mean-square error, or target confusion rates are not required separately, since their operational relevance is already encoded within the Bayes risk formulation with their associated user chosen costs.

An important part of the evaluation is the interpretability of the resulting allocation behavior. The goal is not only to determine whether the proposed RRM framework reduces cumulative risk, but also to understand why certain targets receive more resources than others. This interpretability is important because it shows how the user defined cost matrix influences the allocation and how the costs can be adjusted according to mission specific priorities. Therefore, each scenario is discussed in terms of both the achieved risk reduction and the allocation behavior that emerges from the chosen mission dependent cost function.

5.4. Simulation Scenarios and Results

The simulation scenarios are based on a common three target engagement geometry. The targets are initialized at ranges of $15\,\mathrm{km}$, $16\,\mathrm{km}$, and $17\,\mathrm{km}$ from the radar, respectively, and move directly toward the radar with equal velocity. The initial separation of $1\,\mathrm{km}$ between consecutive targets provides similar, but not identical, sensing conditions.

The same underlying target trajectories are used throughout all scenarios. This allows differences in the adaptive resource allocation and realized Bayes risk to be attributed to the considered operational variation, instead of differences in target motion or engagement geometry. The symmetric cost scenario A-1 is used as the reference configuration. An overview of all scenarios and their main variations is provided in Table 5.8.

Table 5.8: Overview of the considered simulation scenarios.

Scenario	Short name	Main variation
A-1	Baseline costs	Symmetric tracking and classification costs
A-2	Tracking-focused	Classification errors between target classes are not penalized
A-3	Class-priority costs	Errors involving one target class are not penalized
B-1	Target appearance	A new target appears during the simulation
B-2	Target disappearance	One target disappears during the simulation
C	Range-dependent costs	Costs increase according to a linear range-dependent cost function
D	Active classification region	Classification costs are active only inside a predefined region

The use of three targets is particularly relevant for the target appearance and target disappearance scenarios. In these cases, the number of active targets changes between two and three. Starting with this number of targets ensures that the allocation problem remains nontrivial when one target is absent. With only two possible targets, an appearance or disappearance event would create a period with a single active target, during which no meaningful balancing of radar resources between targets would be required.

At the same time, the number of targets is deliberately kept small. This makes it possible to interpret the interaction between the individual target beliefs, allocated dwell times, and resulting risks. The three target geometry therefore provides a balance between maintaining a meaningful resource allocation problem and preserving interpretability in the observed framework behavior.

The targets move toward the radar so that their sensing conditions gradually improve as their distances to the radar decrease. This is particularly relevant for the classification component, since the quality of the class dependent measurements improves as a target approaches the radar. The selected initial ranges create a gradual transition from uncertain to more confident class beliefs. If the targets were initialized too close to the radar, the class belief probabilities could converge after only a small number of updates, causing the classification dependent allocation behavior to disappear very early in the simulation.

The results are visualized using four complementary plot types. Together, these plots show the target geometry, the evolution of the Bayesian beliefs, the resulting resource allocation, and the realized Bayes risk.

The trajectory plots show the motion of the targets within the surveillance region. The initial position of each target is indicated by a marker. Depending on the scenario, additional elements such as classification regions or spatially varying cost fields are included to illustrate the corresponding operational setting.

The prior belief plots show the evolution of the joint prior probabilities associated with target absence and the considered target classes. A probability close to one indicates strong confidence in the corresponding hypothesis, while a probability close to zero indicates that the hypothesis is considered unlikely.

The budget allocation plots show the fraction of the available radar time budget assigned to each target. Because the radar resources are shared, these plots show how the optimizer redistributes the available dwell time according to the predicted Bayes Risk reduction for the individual targets.

Finally, the risk plots show the realized Bayes risk for the optimized allocation and the fixed equal budget strategy. These plots visualize how the difference between the two allocation strategies develops over time in each scenario. To ensure a fair comparison, both simulations use the same common random numbers for the stochastic measurement process. The belief states will evolve separately according to the allocation strategy used in each run, as the budget allocation influences the quality of the measurements.

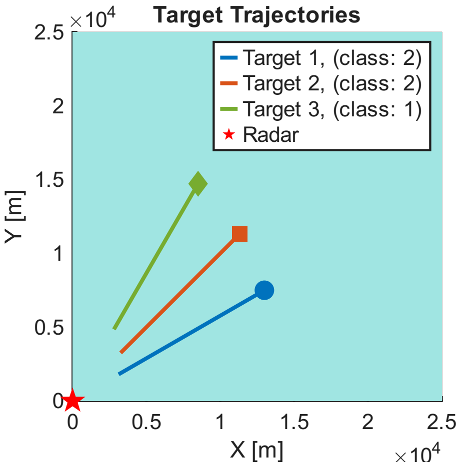

(a) Target trajectories.

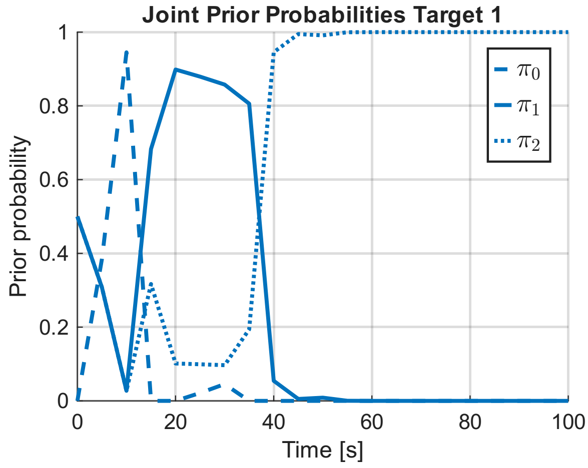

(b) Joint prior probabilities for target 1.

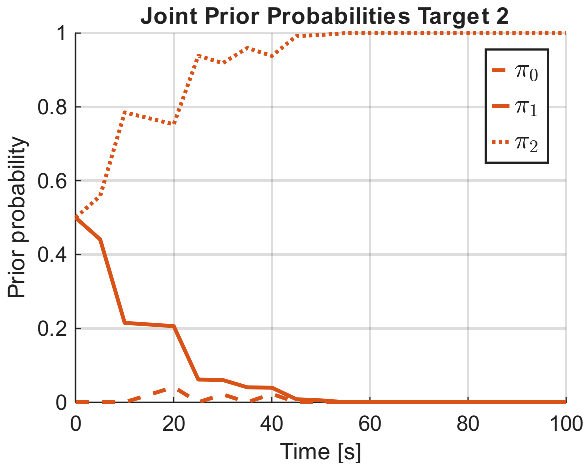

(c) Joint prior probabilities for target 2.

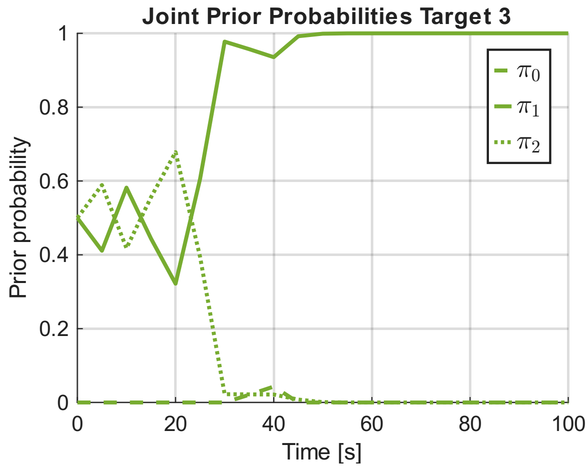

(d) Joint prior probabilities for target 3.

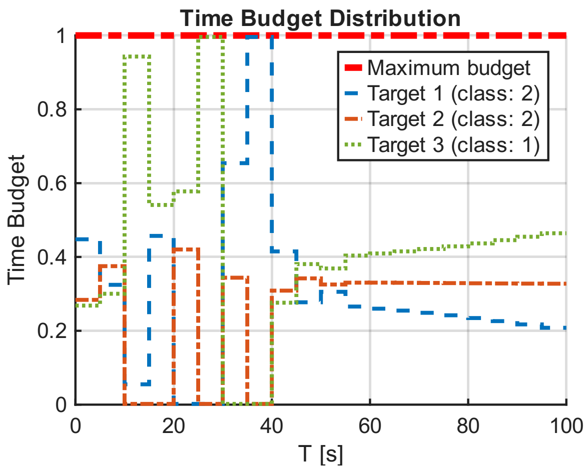

(e) Adaptive time-budget allocation over time.

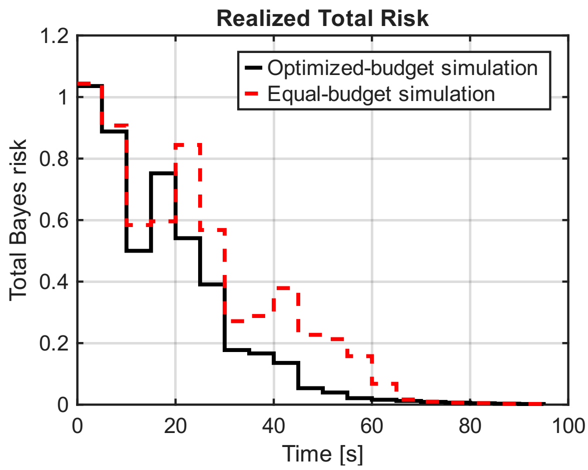

(f) Bayes risk compared to equal-budget simulation.

Figure 5.1: Simulation results for Scenario A-1. The top row shows the target trajectories and the joint prior probabilities for target 1. The middle row shows the joint prior probabilities for targets 2 and 3. The bottom row shows the adaptive radar resource allocation and the resulting Bayes risk compared to the equal-budget simulation.

5.5. Scenario A-1: Baseline symmetric costs

Scenario definition

Scenario A-1 is used as the baseline scenario. Symmetric operational costs are assigned to all incorrect target-state decisions, so target existence errors and classification errors are considered equally costly and no target class or location is preferred.

The joint cost matrix used in this scenario is given by

Table 5.9: Cost matrix for Scenario A-1.

$$\mathbf{C}_{\mathrm{A1}} = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}.$$

Because all off diagonal entries are equal, the optimizer is expected to be driven by the current belief uncertainty and measurement quality. This scenario is therefore used as the reference case for evaluating the different cost matrices and target presence changes introduced in the following scenarios.

Results and interpretation

The simulation results for Scenario A-1 are shown in Figure 5.1. As the baseline scenario, the adaptive allocation shows how radar resources are distributed when all incorrect decisions are considered equally costly.

During the first part of the simulation, the joint priors still contain uncertainty about both target existence and target class. The time budget allocation therefore changes strongly, as the optimizer assigns dwell time to targets for which additional measurements are expected to reduce the immediate Bayes risk most effectively. This behavior is consistent with the myopic nature of the current implementation, since the allocation is optimized one step ahead and does not explicitly account for the longer term evolution of the belief states.

The allocation is driven by both target existence and classification uncertainty. At $t = 35$, for example, the absence probability of the corresponding target is already close to zero, while the class probabilities are still uncertain. The optimizer nevertheless assigns the maximum dwell time to this target, indicating that the allocation is mainly driven by the remaining classification uncertainty. This shows that, under the symmetric baseline cost matrix, the framework actively allocates radar resources to improve classification when classification errors contribute to the Bayes risk.

The results also show the myopic nature of the one step ahead optimization. Since the optimizer only minimizes the immediate expected Bayes risk, it can assign a large fraction, or even all, of the available budget to a single target. Consequently, other targets may temporarily receive little dwell time, which can increase their uncertainty later. At $t = 15$, this is visible in the realized risk plot, where the optimized allocation gives a larger realized risk than the equal budget allocation.

The equal budget strategy can therefore outperform the optimized strategy at some timesteps. This does show that an allocation that is beneficial for the immediate risk does not necessarily remain beneficial over the full observation time as the belief states evolve.

After approximately $t = 50$, all the joint priors have converged to high confidence target beliefs. The system is then certain about both target existence and target class, so the adaptive allocation becomes more stable.

In this stabilized phase, the budget allocation is mainly determined by target range and thus the measurement quality. The more distant target receives more dwell time, because its lower SNR leads to a higher probability of detection errors and larger measurement uncertainty. As a result, its contribution to the Bayes risk is larger, causing the optimizer to allocate more radar resources to that target.

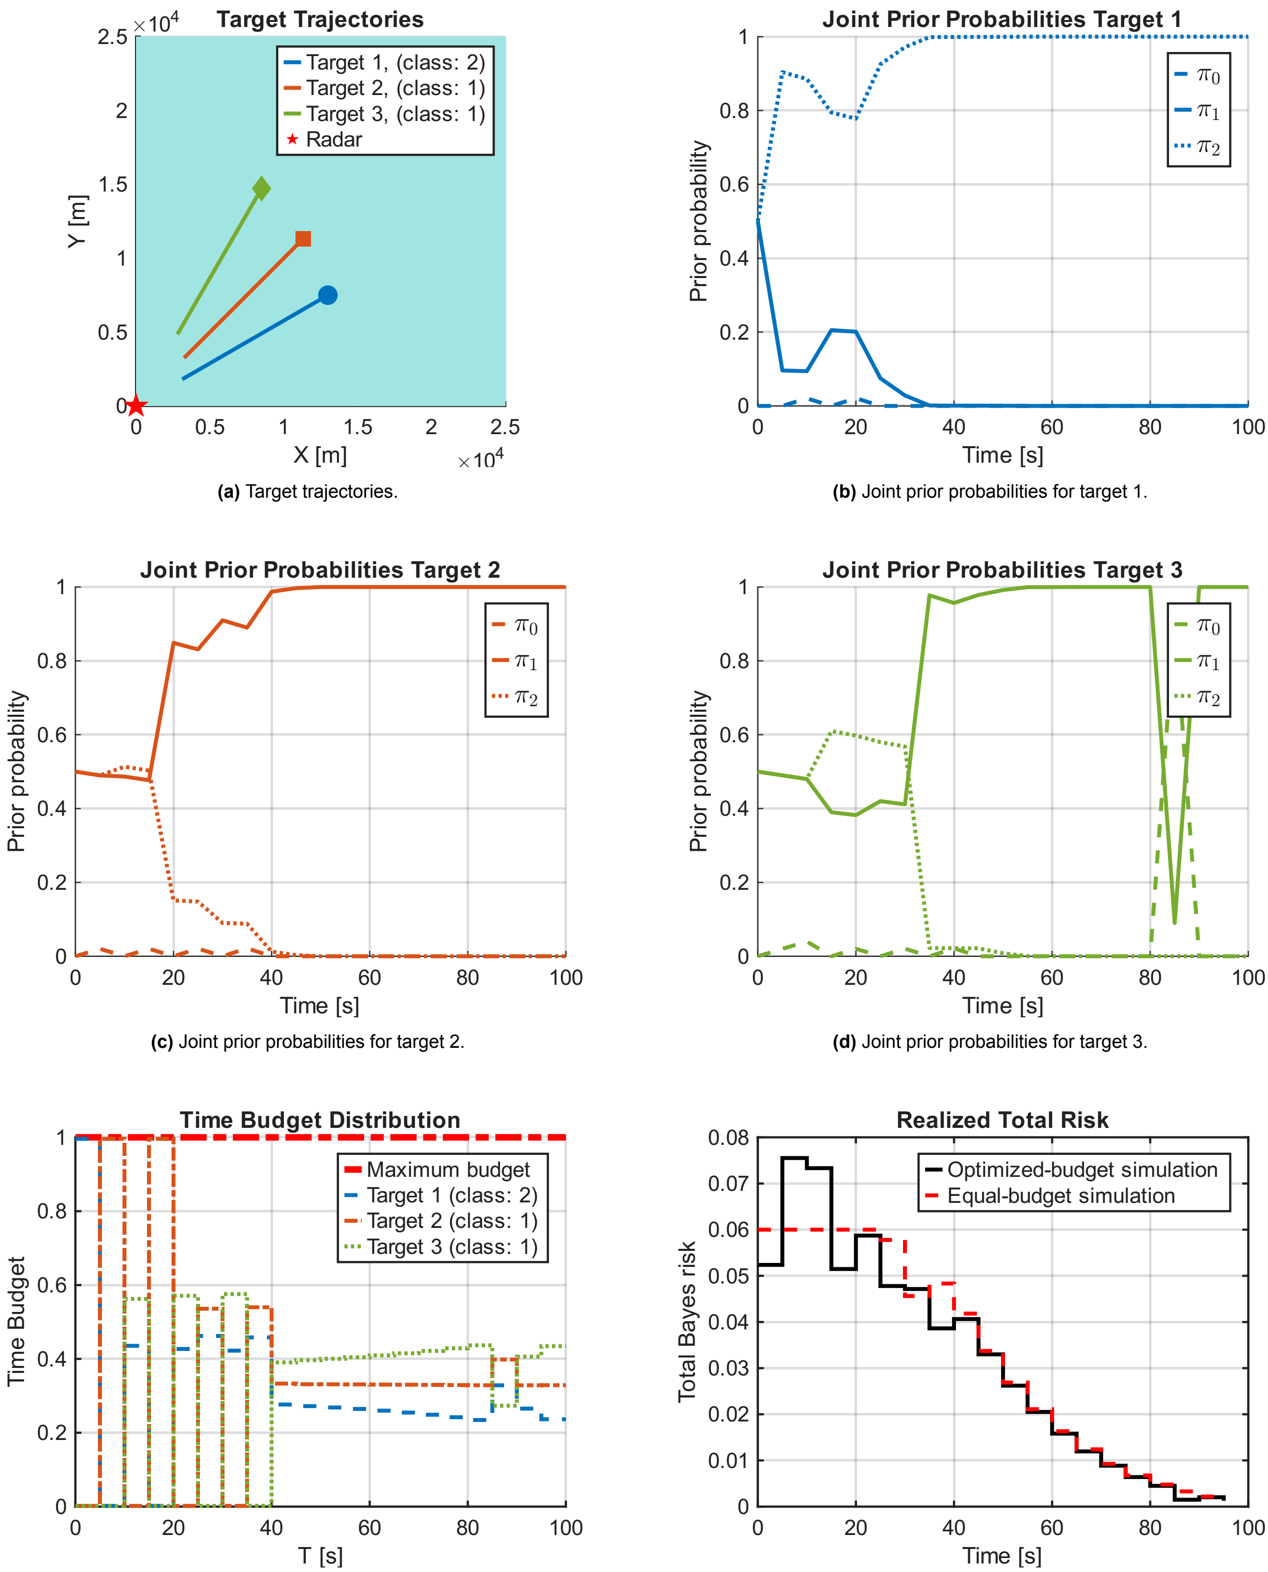

(a) Target trajectories.

(b) Joint prior probabilities for target 1.

(c) Joint prior probabilities for target 2.

(d) Joint prior probabilities for target 3.

(e) Adaptive time-budget allocation over time.

(f) Bayes risk compared to equal-budget simulation.

Figure 5.2: Simulation results for Scenario A-2. The top row shows the target trajectories and the joint prior probabilities for target 1. The middle row shows the joint prior probabilities for targets 2 and 3. The bottom row shows the adaptive radar resource allocation and the resulting Bayes risk compared to the equal-budget simulation.

5.6. Scenario A-2: Tracking-only costs

Scenario A-2 investigates a tracking only operational setting. In this scenario, classification errors between the target classes are not penalized. The operational costs should thus only be associated with distinguishing target existence from target absence, while confusion errors between the two target classes do not contribute to the Bayes risk.

This scenario evaluates whether the proposed framework shifts its sensing strategy toward track maintenance when the mission specific cost matrix assigns no cost to classification errors. In this case, classification uncertainty no longer contributes to the Bayes risk objective, illustrating that the user can decide through the mission objective whether classification performance should influence the resource allocation.

Table 5.10: Cost matrix for Scenario A-2.

$$\mathbf{C}_{\mathrm{A2}} = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}.$$

The simulation results for Scenario A-2 are shown in Figure 5.2. Compared to Scenario A-1, the classification related costs have been removed. As a result, the optimizer no longer has an incentive to allocate dwell time specifically to distinguish between target class 1 and target class 2.

This behavior can best be observed in the joint prior probability plots. Although the class probabilities still evolve over time due to the received feature measurements, the adaptive allocation is not driven by classification uncertainty. In contrast to Scenario A-1, the optimizer does not assign budget to targets because their class probabilities are still uncertain. This confirms that the framework responds to the user defined operational costs: when classification errors are not penalized, classification uncertainty no longer contributes to the resource allocation objective.

The dominant behavior in this scenario is instead related to target existence. Several missed detections occur during the simulation, which appear as spikes in the prior probability of target absence in the joint prior plots. These spikes can for example be seen at $t = 5, 15, 25$ and 35 for target 2. When such a spike occurs, the optimizer reacts strongly by assigning a large fraction, and in some cases all, of the available dwell time to the corresponding target. This behavior is expected, since the tracking only cost matrix penalizes incorrect decisions about whether a target is present or absent.

A clear example can be seen at $t = 85$, where the absence probability of one target increases sharply due to a missed detection. At this point, the prior belief temporarily suggests that the target may be absent. Since the no measurement and clutter only cases in the probability densities do not benefit from increased SNR in the same way as target originated detections, giving dwell time to this target is no longer an effective way to reduce the immediate Bayes risk. The optimizer can therefore allocate more budget to the other targets. In the next timestep, when a target originated measurement is received again, the belief in target existence is restored and the budget allocation returns into the same stabilized pattern as for Scenario A-1.

It is important to note that the framework can still classify the targets over time, even when classification errors are not included in the mission specific cost matrix. However, this classification is not caused by allocation behavior that explicitly aims to improve classification certainty. It just follows from the class related information contained in the measurements obtained under the tracking focused allocation.

Overall, Scenario A-2 demonstrates that the proposed Bayes risk-based framework adapts its sensing behavior according to the specified operational objective. Since classification errors are not penalized, the optimizer does not spend dwell time specifically to classify targets. Any convergence of the class probabilities is therefore a consequence of the measurements obtained while maintaining target tracks, rather than the result of classification driven resource allocation.

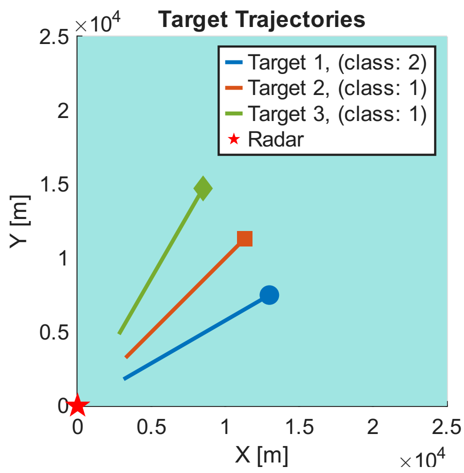

(a) Target trajectories.

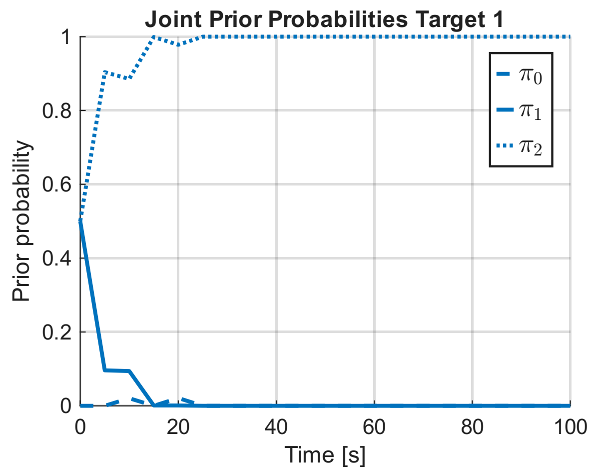

(b) Joint prior probabilities for target 1.

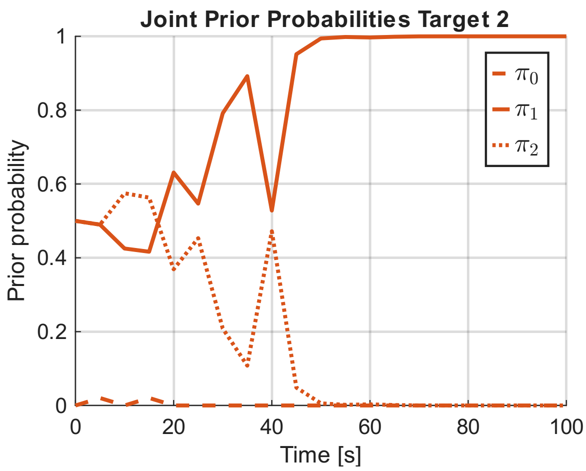

(c) Joint prior probabilities for target 2.

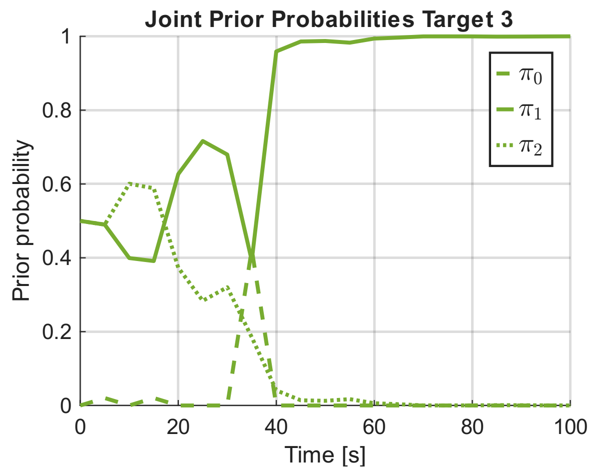

(d) Joint prior probabilities for target 3.

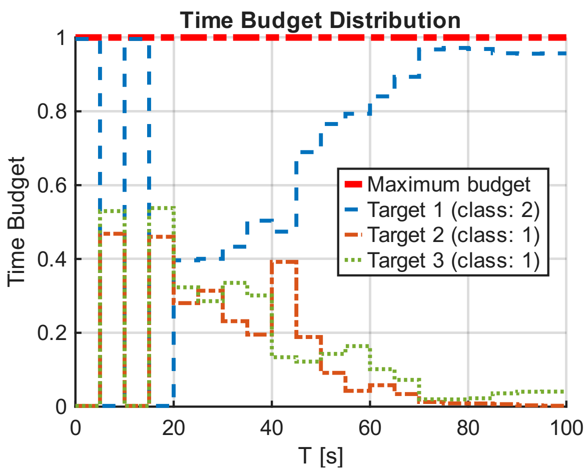

(e) Adaptive time-budget allocation over time.

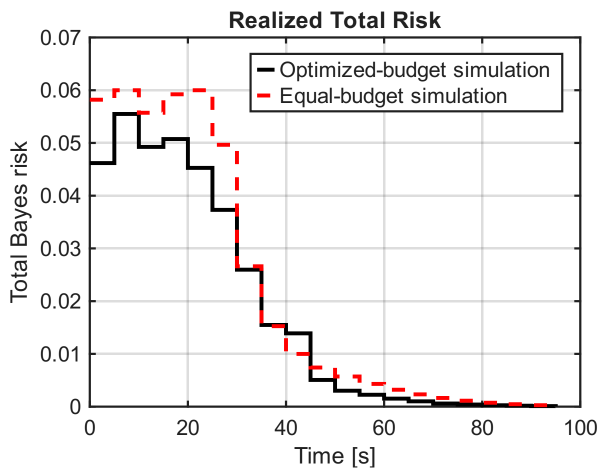

(f) Bayes risk compared to equal-budget simulation.

Figure 5.3: Simulation results for Scenario A-3. The top row shows the target trajectories and the joint prior probabilities for target 1. The middle row shows the joint prior probabilities for targets 2 and 3. The bottom row shows the adaptive radar resource allocation and the resulting Bayes risk compared to the equal-budget simulation.

5.7. Scenario A-3: Non-penalized class

Scenario A-3 considers an asymmetric operational cost setting in which errors involving class 1 are not penalized. In this scenario, the operational objective only assigns costs to decisions involving class 2 or target absence. As a result, once a target is believed to belong to class 1 with high confidence, incorrect decisions involving that target no longer contribute to the Bayes Risk.

This scenario investigates whether the proposed Bayes Risk-based framework reduces the allocated radar resources for targets that are believed to belong to a non penalized class, while continuing to allocate resources to targets that can belong to the penalized class.

Table 5.11: Cost matrix for Scenario A-3.

$$\mathbf{C}_{\mathrm{A3}} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}.$$

Results and interpretation

The simulation results for Scenario A-3 are shown in Figure 5.3. Compared to Scenario A-1, the cost matrix no longer penalizes errors involving class 1. Because of this, the adaptive allocation therefore becomes strongly dependent on whether a target is believed to belong to the penalized class or the non-penalized class.

This behavior can be observed clearly in the budget allocation plot. Target 1 belongs to class 2 and is classified very fast. This results in the fact that this target remains operationally relevant. As the simulation progresses, the optimizer increasingly allocates a large fraction of the available dwell time to this target, because errors involving class 2 are penalized by the cost function.

Targets 2 and 3 belong to class 1. Once the joint prior probabilities indicate with high confidence that these targets belong to class 1, the allocated budget for these targets decreases strongly. This is expected, because mistakes involving class 1 are not penalized in this scenario. The optimizer therefore has little incentive to spend dwell time on targets that are believed to belong to the non penalized class.

At $t = 40$, the class belief for one of the class 1 targets becomes more uncertain. In particular, the probability that the target may belong to class 2 temporarily increases. This causes a small increase in the allocated budget, because the target again contributes to the expected Bayes risk through the possibility of being class 2. Once the class belief returns to high confidence in class 1, the allocated budget decreases again.

Overall, Scenario A-3 again demonstrates that the proposed framework responds directly to the mission specific operational cost matrix. Targets that are confidently classified as class 1 receive little or no budget, because errors involving class 1 are not penalized. In contrast, targets that are believed to belong to class 2, or that still have a non negligible probability of belonging to class 2, receive more radar resources because they remain relevant to the Bayes risk objective.

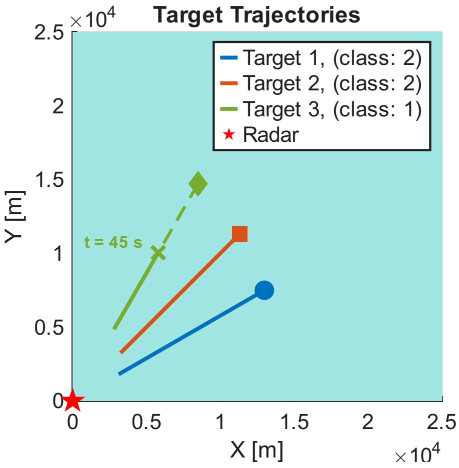

(a) Target trajectories for the target-appearance scenario. The cross indicates the position at which the newly initialized target enters the scenario at $t = 45\,\mathrm{s}$.

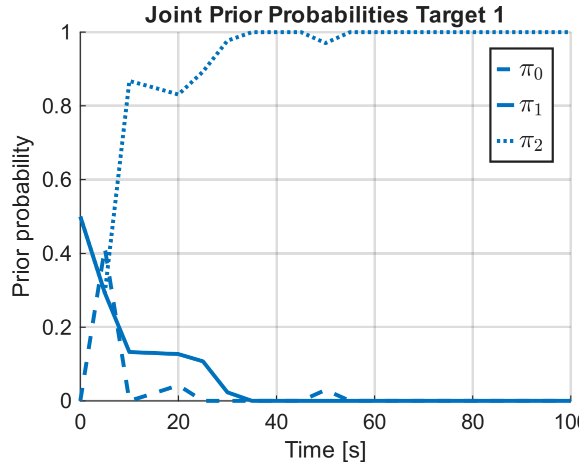

(b) Joint prior probabilities for target 1.

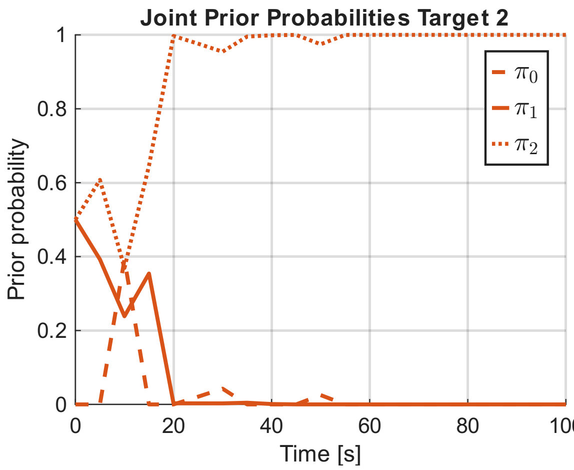

(c) Joint prior probabilities for target 2.

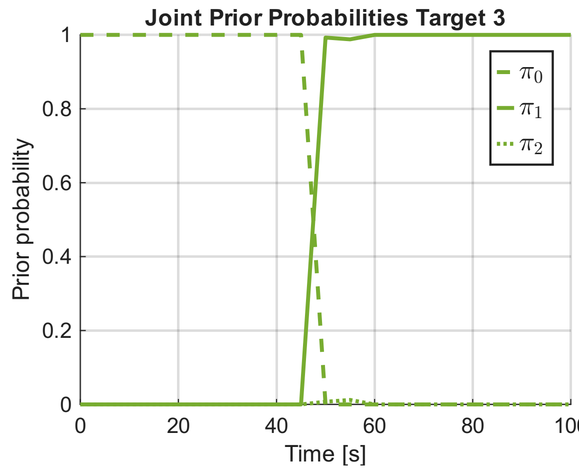

(d) Joint prior probabilities for target 3.

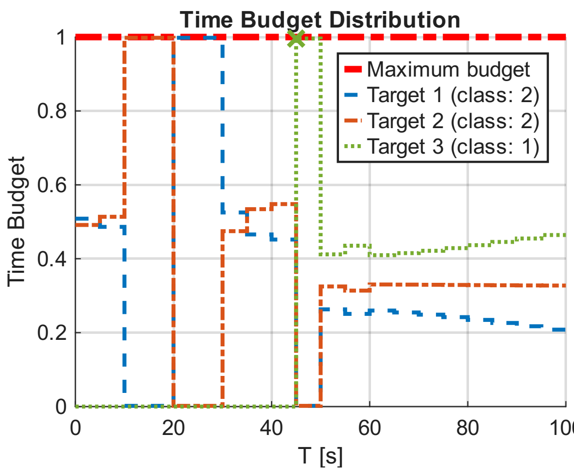

(e) Adaptive time-budget allocation over time. The cross on the target 3 budget curve marks the time at which target 3 is first detected and initialized as a known target in the resource-allocation process.

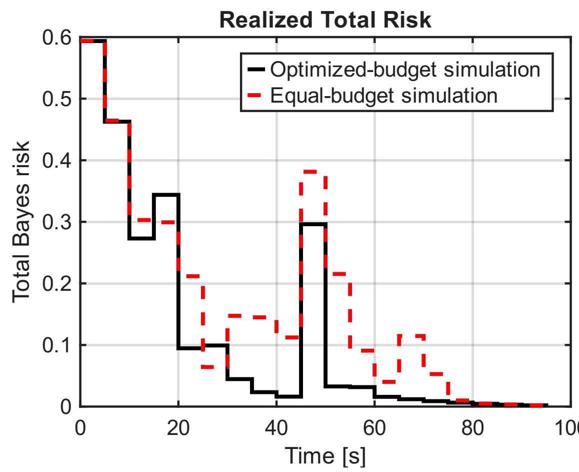

(f) Bayes risk compared to equal-budget simulation.

Figure 5.4: Simulation results for Scenario B-1. The top row shows the target trajectories and the joint prior probabilities for target 1. The middle row shows the joint prior probabilities for targets 2 and 3. The bottom row shows the adaptive radar resource allocation and the resulting Bayes risk compared to the equal-budget simulation.

5.8. Scenario B-1: Appearing target

Scenario B-1 investigates adaptive radar resource allocation when a previously unknown target is introduced into the radar decision process. In the simulated environment, Target 3 is already following its trajectory before $t = 45\,\mathrm{s}$, but it is not included in the radar system's prior target set, filtering process, or resource allocation before this time. At $t = 45\,\mathrm{s}$, the target is first detected and initialized as a known target. From that point onward, the radar system must update its joint target belief and redistribute the available sensing resources accordingly.

The operational cost matrix is the same as in Scenario A-1 and remains symmetric throughout the scenario. Target existence errors and classification errors are therefore penalized equally, which allows the effect of target appearance to be studied without changing the operational cost structure. This scenario evaluates how robust the proposed framework responds when a new target is introduced and how the additional target affects the resulting allocation behavior.

Results and interpretation

The simulation results for Scenario B-1 are shown in Figure 5.4. During the first part of the simulation, target 3 is absent. This is reflected in the joint prior probabilities, where the probability of target absence for target 3 remains equal to one until the target appears.

At $t = 45$, target 3 is introduced into the surveillance region. Its joint prior probabilities are initialized according to the birth probabilities π_{birth}, which creates a large uncertainty about the newly appearing target class. This uncertainty immediately increases the Bayes risk associated with target 3.

The effect of this appearance is clearly visible in the adaptive budget allocation. Directly after the target appears, the optimizer assigns all available dwell time to target 3. This allocation is used to rapidly classify the newly appearing target.

As a result, target 3 is incorporated into the framework almost immediately. The joint prior probabilities show that the target is classified within only a few timesteps after its appearance. This demonstrates that the proposed Bayes risk-based allocation strategy can react strongly to a newly introduced target when its uncertainty contributes significantly to the expected operational cost.

After this initial response, the allocation gradually returns toward the stabilized behavior observed in Scenario A-1. Once target 3 has been detected and classified with high confidence, the remaining budget allocation is again mainly determined by differences in target range, measurement quality, and the residual uncertainty of the different targets.

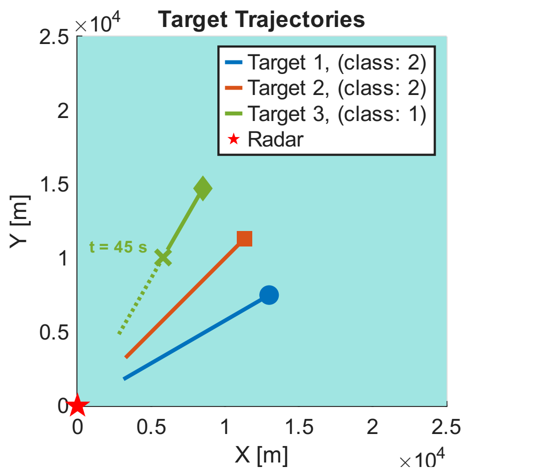

(a) Target trajectories. The cross marks the position at which target 3 disappears from the scenario at $t = 45\,\mathrm{s}$.

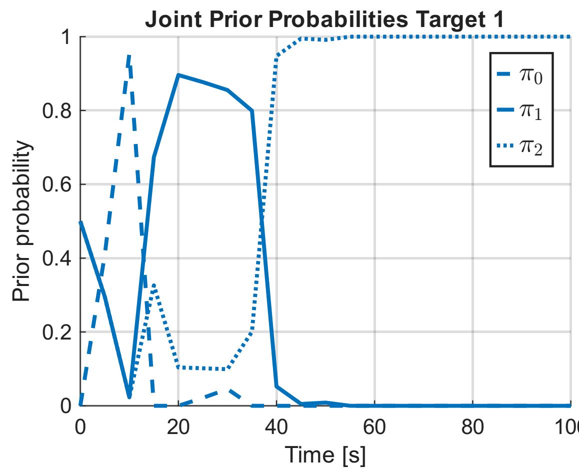

(b) Joint prior probabilities for target 1.

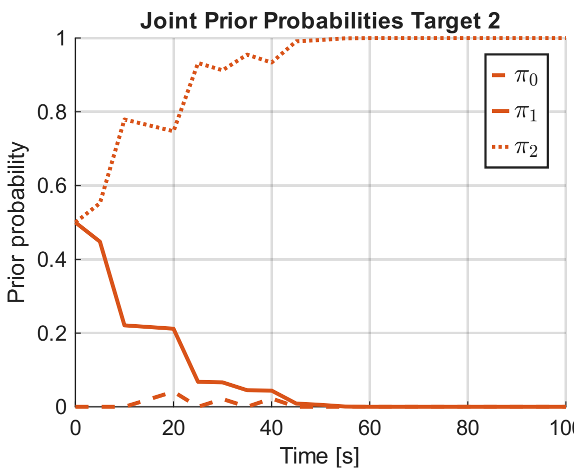

(c) Joint prior probabilities for target 2.

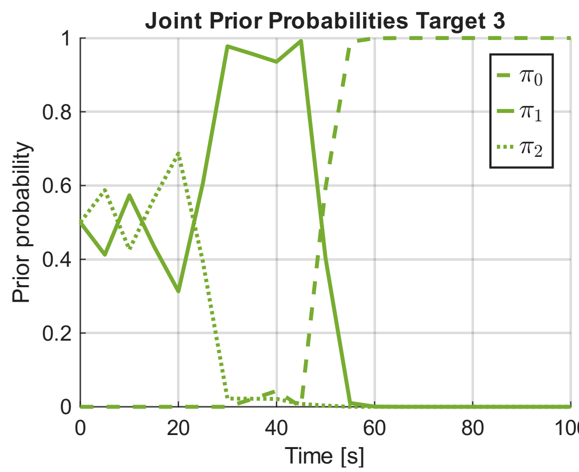

(d) Joint prior probabilities for target 3.

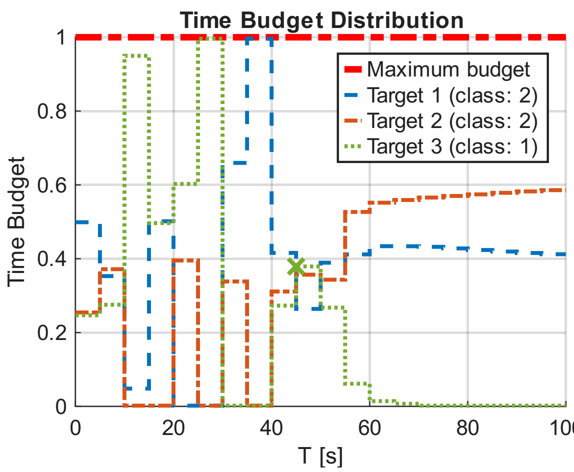

(e) Adaptive time-budget allocation over time. The cross marks the time at which target 3 disappears.

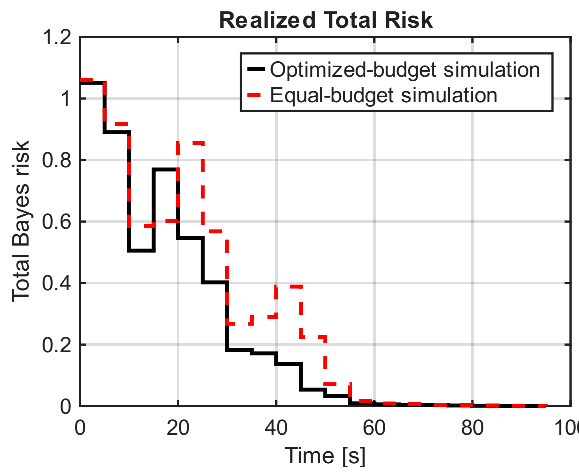

(f) Bayes risk compared to equal-budget simulation.

Figure 5.5: Simulation results for Scenario B-2. The top row shows the target trajectories and the joint prior probabilities for target 1. The middle row shows the joint prior probabilities for targets 2 and 3. The bottom row shows the adaptive radar resource allocation and the resulting Bayes risk compared to the equal-budget simulation.

5.9. Scenario B-2: Disappearing target

Scenario B-2 investigates adaptive radar resource allocation in the presence of a dynamically disappearing target. Initially, all three targets are present within the surveillance region. At $t = 45$, target 3 is assumed to have disappeared and it stops generating target originated measurements. The radar system should then update its joint target belief based on the absence of detections and redistribute the available sensing resources accordingly.

The operational cost matrix is again the same as in Scenario A-1 and remains symmetric throughout the scenario.

Results and interpretation

The simulation results for Scenario B-2 are shown in Figure 5.5. During the first part of the simulation, all three targets are present and the adaptive allocation behaves similarly to the baseline Scenario A-1. Target 3 is detected and classified with high confidence before the disappearance event occurs.

At $t = 45$, target 3 disappears and stops generating target originated measurements. Since the probability of detection is relatively high at this point, the absence of measurements provides strong evidence that the target is no longer present. This is visible in the joint prior probabilities of target 3, where the probability of target absence rapidly increases after the disappearance.

The adaptive allocation responds directly to this change in belief. As the belief in target absence increases, the expected Bayes risk associated with target 3 decreases. Consequently, the optimizer gradually reduces the dwell time allocated to target 3 rather than assigning additional resources to it. Once the framework is confident that the target is absent, allocating radar resources to this target no longer provides significant expected risk reduction and thus goes to the lower bound on τ.

After the disappearance is as good as certain, the system enters a new stable situation with only the two remaining targets. From this point onward, the available radar budget is redistributed between targets 1 and 2. The stabilized allocation then follows the same principle as in the baseline scenario: the remaining target with lower measurement quality, typically the more distant target, receives more dwell time because its contribution to the Bayes Risk is larger.

Overall, Scenario B-2 demonstrates that the proposed Bayes Risk-based framework can adapt not only to newly appearing targets, but is also robust to targets that stop generating measurements.

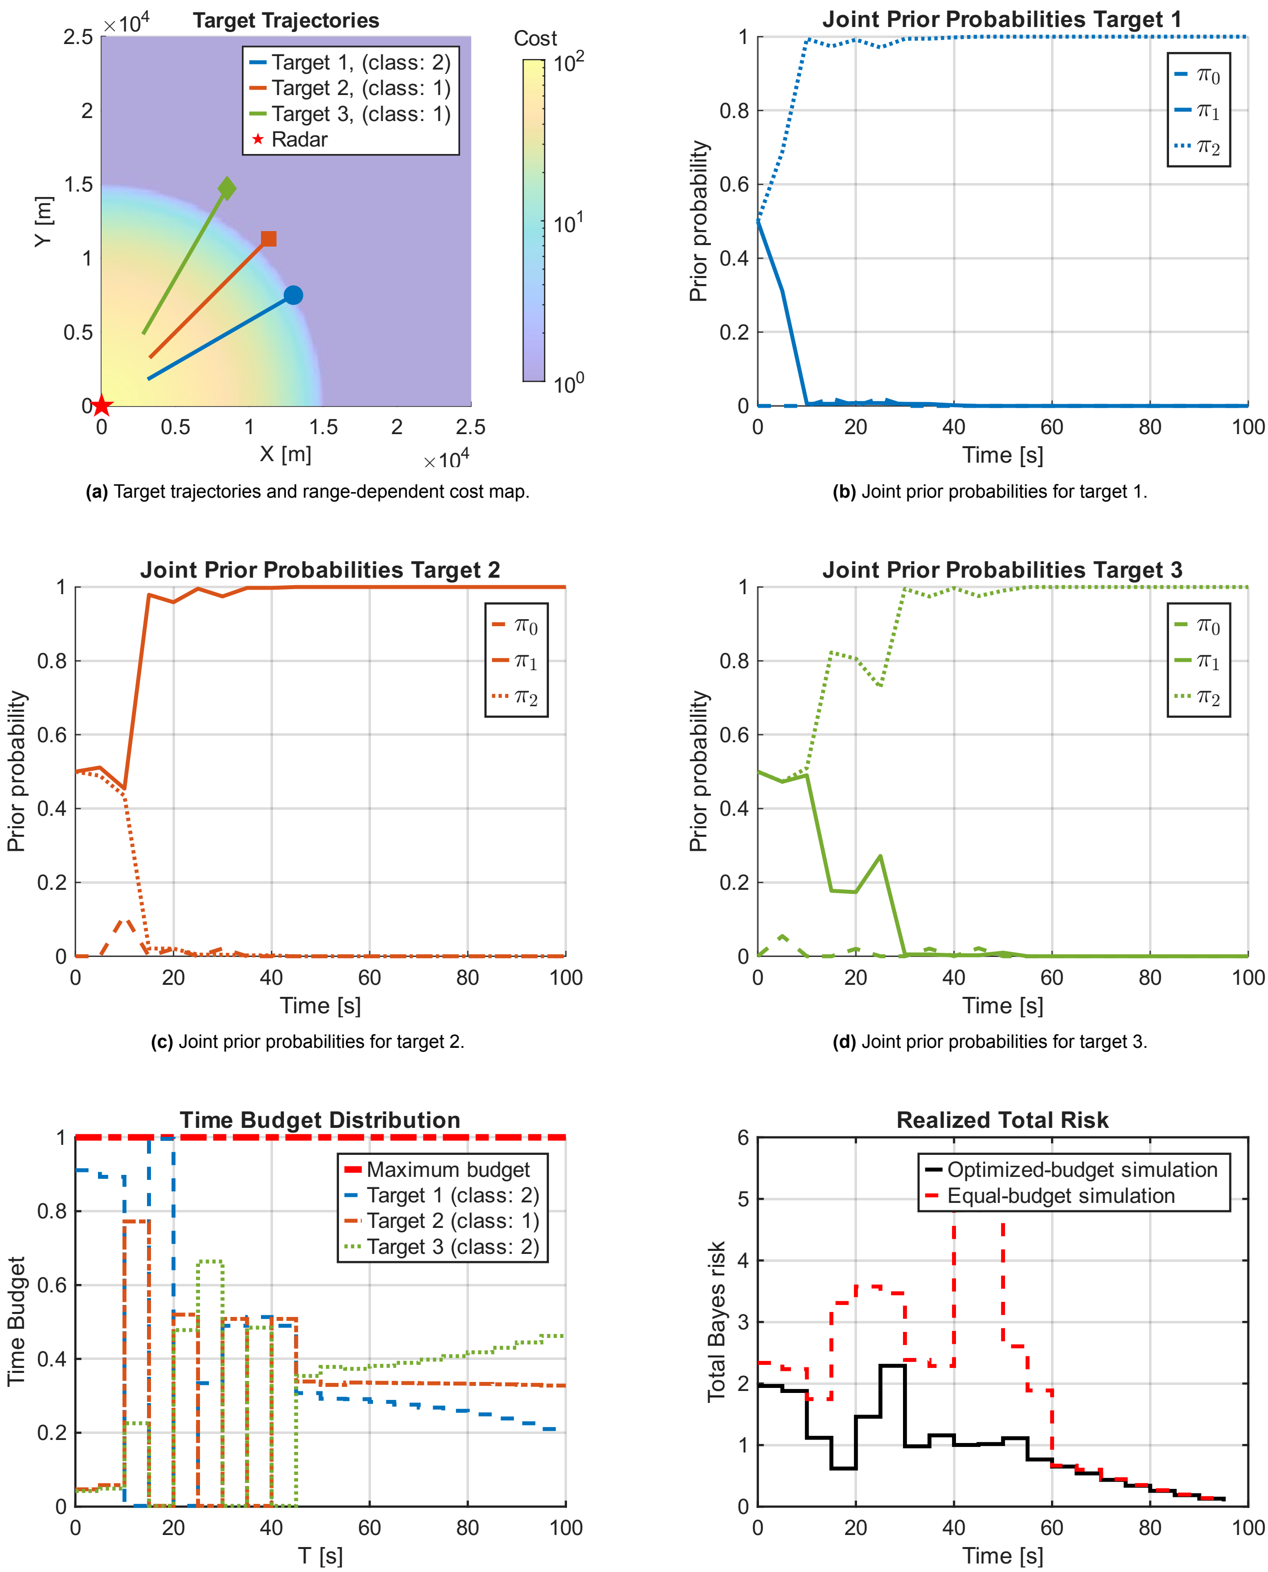

(a) Target trajectories and range-dependent cost map.

(b) Joint prior probabilities for target 1.

(c) Joint prior probabilities for target 2.

(d) Joint prior probabilities for target 3.

(e) Adaptive time-budget allocation over time.

(f) Bayes risk compared to equal-budget simulation.

Figure 5.6: Simulation results for Scenario C. The top row shows the target trajectories with the range-dependent cost structure and the joint prior probabilities for target 1. The middle row shows the joint prior probabilities for targets 2 and 3. The bottom row shows the adaptive radar resource allocation and the resulting Bayes risk compared to the equal-budget simulation.

5.10. Scenario C: Range-dependent costs

Scenario C investigates a range dependent operational cost function. As base joint cost matrix the symmetric cost matrix is used, but the penalty for incorrect decisions increases as a target approaches the radar. This represents an operational setting in which nearby targets are considered more critical than distant targets. As a result, now range affects the allocation behavior in two ways: through the measurement uncertainty, as in Scenario A-1, but now also through the operational range dependent cost assigned to incorrect decisions.

The range dependent joint cost matrix is constructed from the matrix $\mathbf{C}_{\mathrm{A1}}$ from Scenario A-1. For each target, an additional range-dependent penalty is multiplied with the cost matrix according to

$$w(r) = \max\left(\alpha\left(1 - \frac{r}{d_{\mathrm{ref}}}\right), 0\right), \tag{5.1}$$

where r denotes the target range, $\alpha = 100$ determines the strength of the range dependent cost increase, and $d_{\mathrm{ref}} = 15\mathrm{km}$ is the reference range beyond which no additional cost is added. The resulting cost matrix for each target is given by

$$\mathbf{C}_{\mathrm{C}}(r) = (1 + w(r))\,\mathbf{C}_{\mathrm{A1}}, \tag{5.2}$$

Results and interpretation

In Figure 5.6, the simulation results for Scenario C are shown. The adaptive allocation is influenced not only by the uncertainty associated with the targets, but also by their operational importance as determined by their distance to the radar.

From the beginning of the simulation, target 1 receives a relatively large fraction of the available radar budget. This target is closest to the radar, and therefore its operational cost is increased the most by the range dependent cost function. As a result, incorrect decisions involving target 1 contribute more strongly to the Bayes risk than they would under the symmetric baseline cost matrix. The optimizer therefore assigns more dwell time to this target in order to reduce the increased instantaneous risk.

This behavior shows that the adaptive allocation is no longer driven only by uncertainty reduction or measurement quality. In Scenario A-1, the more distant targets tend to receive more budget in the stabilized phase because their lower SNR leads to larger measurement uncertainty. In Scenario C, however, the range dependent cost function counteracts this effect by increasing the operational importance of nearby targets. The allocation is therefore determined by a trade-off between measurement uncertainty and range dependent operational cost.

The effect of this trade-off can also be observed during the transient part of the simulation. When the joint prior probabilities still contain uncertainty about target existence and target class, the optimizer allocates dwell time to targets that have both significant belief uncertainty and high operational cost. As a result, nearby targets can receive a large amount of budget even when other targets are more uncertain, because the cost of making an incorrect decision for the nearby target is larger.

After approximately $t = 50$, the joint prior probabilities have mostly converged and the system enters a more stable allocation regime. In this stabilized phase, the influence of the range-dependent cost function remains visible. Compared to Scenario A-1, the closer target receives substantially more budget, while the allocation to the more distant targets is reduced. This demonstrates that the range dependent cost function directly affects the adaptive budget allocation even after the main classification and existence uncertainties have been resolved.

Overall, Scenario C demonstrates that the proposed Bayes Risk-based framework can incorporate spatially dependent operational importance into the resource allocation process. Nearby targets are prioritized because errors involving these targets are assigned a larger operational cost, while distant targets are still considered through their larger measurement uncertainty. The resulting allocation therefore balances sensing quality against the range-dependent cost of incorrect decisions.

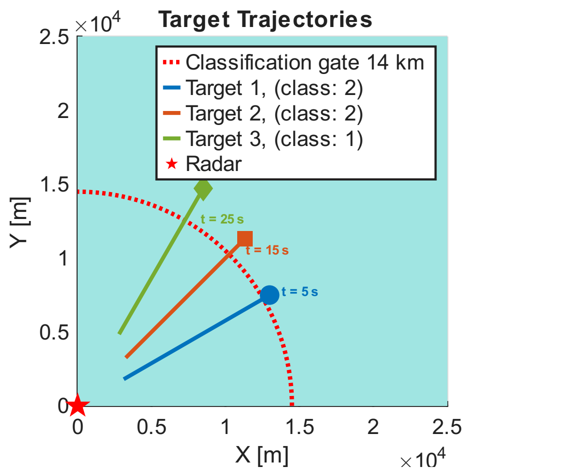

(a) Target trajectories and classification activation boundary. The timestamps indicate the time at which the targets enter the active classification region.

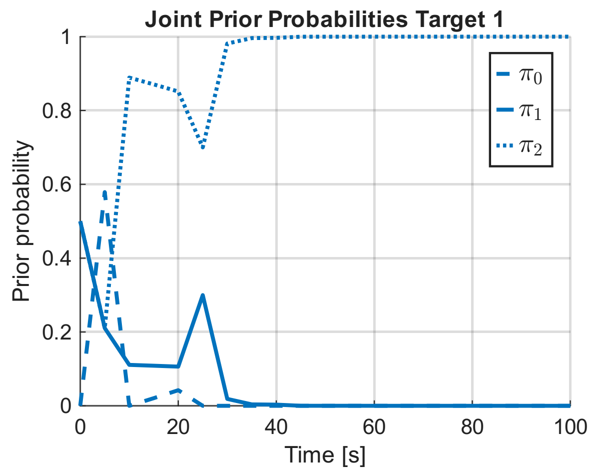

(b) Joint prior probabilities for target 1.

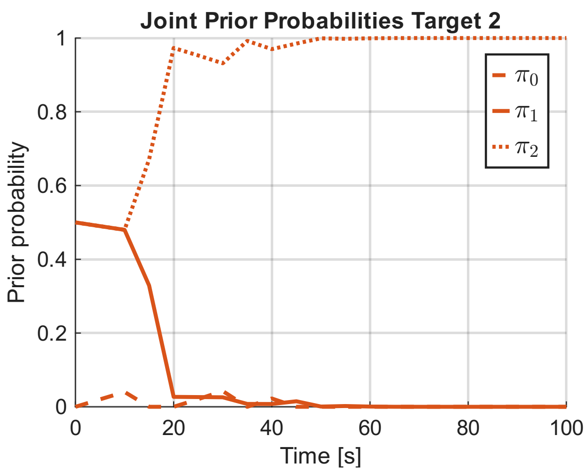

(c) Joint prior probabilities for target 2.

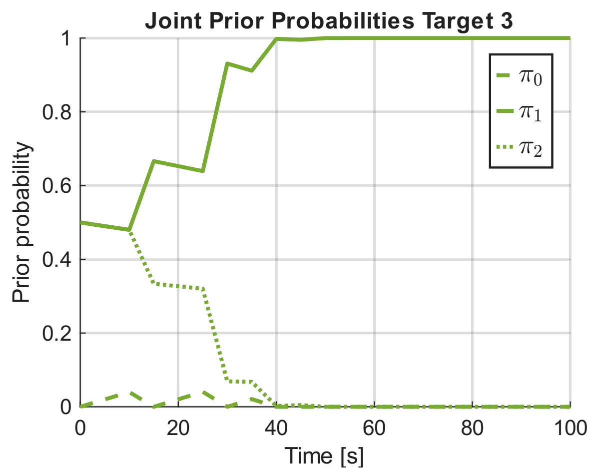

(d) Joint prior probabilities for target 3.

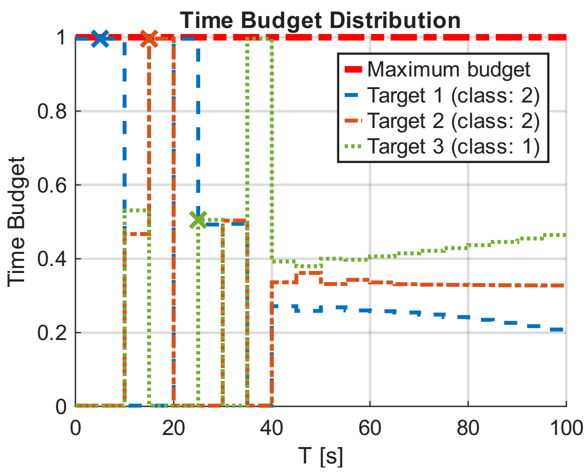

(e) Adaptive time-budget allocation over time. The crosses mark the times at which the corresponding targets enter the active classification region.

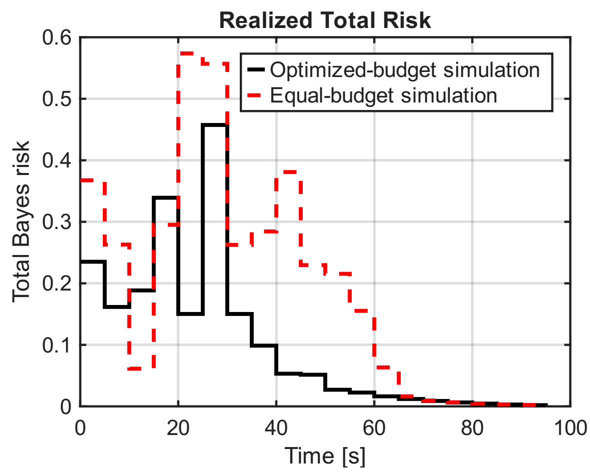

(f) Bayes risk compared to equal-budget simulation.

Figure 5.7: Simulation results for Scenario D. The top row shows the target trajectories with the classification activation boundary and the joint prior probabilities for target 1. The middle row shows the joint prior probabilities for targets 2 and 3. The bottom row shows the adaptive radar resource allocation and the resulting Bayes risk compared to the equal-budget simulation.

5.11. Scenario D: Active classification region

Scenario D investigates an operational setting in which classification is only relevant inside a predefined active region. Outside this region, the radar is only required to maintain target awareness, while distinguishing between the two target classes does not contribute to the operational cost. Once a target enters the active region, classification becomes operationally relevant and classification errors are included in the Bayes risk.

The classification related costs are only activated when the target range lies inside the active region. The boundary is chosen close to the starting points, to prevent the targets to be already classified before entering the region, such that the framework can change it's behavior due to the classification uncertainty. The indicator function can be expressed using the indicator function

$$\chi_{\text{border}}(r) = \begin{cases} 1, & r_{\min} \leq r \leq r_{\max}, \\ 0, & \text{otherwise}, \end{cases} \tag{5.3}$$

where $r_{\min} = 0$ and $r_{\max} = 14.5\text{km}$ define the inner and outer boundary of the region in which classification is operationally relevant. The resulting cost matrix becomes

$$\mathbf{C}_{\text{D}}(r) = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & \chi_{\text{border}}(r) \\ 1 & \chi_{\text{border}}(r) & 0 \end{bmatrix}. \tag{5.4}$$

Results and interpretation

In Figure 5.7, the simulation results for Scenario D are shown. The trajectory plot shows the active classification boundary. Outside this region, classification errors are not penalized, while inside the region the cost of confusing class 1 and class 2 becomes active.

The effect of the active classification region is visible in the adaptive budget allocation. When a target enters the active region, the optimizer immediately assigns a large fraction of the available dwell time to that target. This occurs because the target class suddenly becomes operationally relevant. As a result, reducing classification uncertainty now contributes directly to reducing the Bayes risk.

This behavior can be observed at the moments where the individual targets enter the active classification region. At $t = 5$, target 1 enters the active region and keeps receiving the biggest allocated dwell time. At this point, classification errors for target 1 start contributing to the Bayes risk, as well as multiple missed detections, causing the optimizer to prioritize this target. At $t = 15$, target 2 enters the active region and the budget allocation shifts toward this target in order to reduce its classification uncertainty, which we are sure of, because at this timestep, it is certain of the target existence because $\pi_0 \approx 0$. At $t = 25$, target 3 enters the active region and also receives a large fraction of the available dwell time.

At $t = 25$, the allocation must be shared between multiple competing targets. Target 3 has just entered the active classification region, while target 1 becomes more uncertain again. Since both targets now contribute to the Bayes risk, the optimizer divides the available dwell time between them instead of assigning all resources to a single target. This illustrates how the framework balances classification uncertainty across targets when several targets are operationally relevant at the same time.

After the classification uncertainty has been reduced, the allocation stabilizes. Once the framework is confident about the target classes, no dwell time is needed anymore for the classification. The remaining allocation is then determined by target-existence uncertainty, target range, and the associated measurement quality, just like in Scenario A-1.

Overall, Scenario D demonstrates that the proposed Bayes Risk-based framework can incorporate location dependent classification relevance. Targets outside the active region do not receive additional radar resources for classification, because classification errors are not penalized there. Once a target enters the active region, classification becomes part of the Bayes risk objective, and the optimizer reallocates radar resources accordingly, adding another mission specific situation to the toolkit.

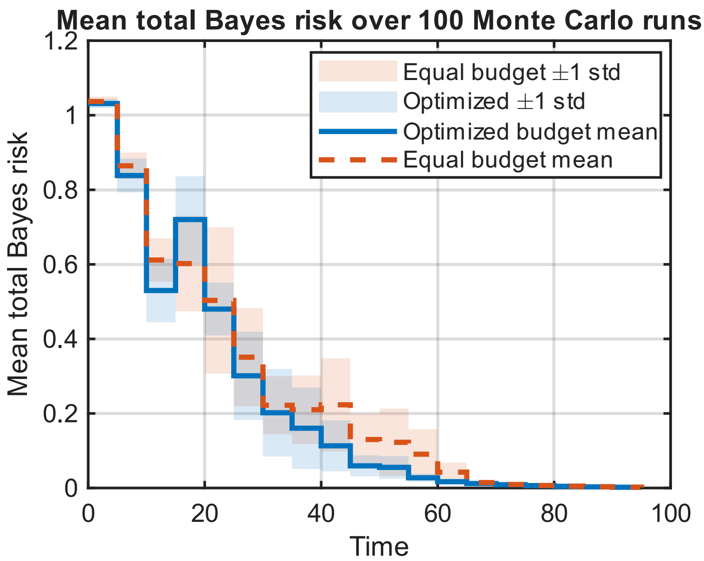

Figure 5.8: Mean total Bayes risk over 100 Monte Carlo runs for the optimized and equal-budget allocation strategies in Scenario A-1. The shaded regions show one standard deviation around the mean.

Table 5.12: Monte Carlo comparison between optimized and equal-budget allocation for all simulation scenarios.

Scenario	$\bar{J}_{\text{opt}}$	$\bar{J}_{\text{eq}}$	ΔJ	**Rel. impr. [%]**	**Better runs**	p**-value**
A-1	4.5732	5.0524	0.4792	6.50	71/100	2.95×10^{-6}
A-2	0.6172	0.6427	0.0255	3.94	99/100	3.30×10^{-34}
A-3	0.3364	0.4013	0.0649	15.68	86/100	2.82×10^{-17}
B-1	2.2295	2.7246	0.4950	16.88	86/100	9.35×10^{-19}
B-2	4.0420	4.7297	0.6877	12.94	78/100	1.42×10^{-13}
C	26.5972	39.1500	12.5528	29.98	96/100	2.74×10^{-26}
D	2.3033	2.8237	0.5204	12.22	87/100	7.14×10^{-9}

5.12. Analysis of Performance

To be able to prove the improved performance of the RRM framework, the performance of the optimized allocation strategy is evaluated using paired Monte Carlo simulations. For each Monte Carlo run r, the same underlying random realization is used for both the optimized-budget and fixed equal budget simulations. This makes it possible to compare the two allocation strategies directly, since differences in performance are then mainly caused by the allocation strategy rather than by different random measurement or clutter realizations.

For each Monte Carlo run, the paired risk difference is defined as

$$d_r = J_{\mathrm{eq},r} - J_{\mathrm{opt},r}, \tag{5.5}$$

where $J_{\mathrm{eq},r}$ and $J_{\mathrm{opt},r}$ denote the cumulative realized Bayes risk of the fixed equal budget and optimized budget simulations, respectively. A positive value of d_r therefore indicates that the optimized allocation achieves a lower cumulative realized risk than the equal budget allocation.

To assess whether the observed improvement is statistically significant, a paired two sided t-test is applied to the paired differences d_r. The null hypothesis is that the mean paired difference is zero,

$$H_0 : \mu_d = 0, \tag{5.6}$$

while the two-sided alternative hypothesis is

$$H_1 : \mu_d \neq 0. \tag{5.7}$$

The corresponding test statistic is given by

$$t = \frac{\bar{d}}{s_d/\sqrt{N_{\mathrm{MC}}}}, \tag{5.8}$$

where $\bar{d}$ and s_d denote the sample mean and sample standard deviation of the paired differences, respectively, and N_{MC} is the number of Monte Carlo runs. The corresponding p-value is obtained from a t-distribution with $N_{\mathrm{MC}} - 1$ degrees of freedom. A small p-value indicates that the observed difference between the optimized and equal-budget allocations is unlikely to be caused by Monte Carlo variation alone, and we can thus conclude that the realized risk decreases compared to the equal budget case.

The resulting performance metrics for all considered scenarios are summarized in Table 5.12. Since the individual scenario behavior has already been discussed in the preceding sections, the Monte Carlo analysis is primarily used here to evaluate the overall performance difference between the optimized and equal-budget allocation strategies. To illustrate the typical time evolution of the realized Bayes risk, a representative scenario is shown in Figure 5.8. The figure shows the mean realized Bayes risk over 100 Monte Carlo runs, together with a one standard deviation band around the mean. This provides insight into both the average performance and the variability of the two allocation strategies.

The optimized allocation generally achieves a lower cumulative realized Bayes risk than the equal budget allocation. However, the mean realized risk does not decrease monotonically over time, which we would expect in this case as the targets in Scenario A-1 are all flying towards the radar. In particular, during the early part of the simulation, the mean optimized allocation shows visible fluctuations in the realized Bayes risk. This behavior can be explained by the myopic nature of the optimization strategy. Since the optimizer minimizes the expected Bayes risk one step ahead, it may assign a large fraction, or even all, of the available dwell time to the target that provides the largest immediate reduction in expected risk. As a result, other targets may temporarily receive little sensing effort, which increases their uncertainty and could lead to a higher realized risk at later time steps.

Due to the fact that this behavior remains visible even after averaging over 100 Monte Carlo runs, it can be concluded that they appear to be a systematic consequence of the one step ahead allocation strategy.

The equal budget strategy distributes the available dwell time evenly across the targets. This can result in smoother risk behavior, since no target is temporarily deprived of sensing resources. However, this strategy does not adapt to differences in uncertainty, target range, or operational cost. Consequently, it does not generally achieve the lowest cumulative realized risk over the full simulation.

Overall, the Monte Carlo results show that the proposed Bayes-risk-based allocation framework improves overall performance compared with equal budget allocation. At the same time, the observed risk fluctuations indicate a limitation of the current myopic implementation. This motivates the investigation of longer horizon allocation strategies in future work, such as policy rollout. Such methods could explicitly account for the downstream effect of current dwell time decisions on future belief states, potentially reducing early risk fluctuations and further improving the overall allocation performance.

The baseline scenario A-1 shows that the optimized allocation already reduces the cumulative realized Bayes risk compared with equal budget allocation in a symmetric cost setting. However, the benefit of the proposed framework becomes more pronounced in the more mission specific scenarios, where the cost function introduces specific operational priorities. This is especially visible in scenario A-3 and the range-dependent cost scenario C. In scenario A-3, the asymmetric cost matrix makes certain class dependent decision errors more important than others, causing the optimizer to allocate resources according to the operational relevance of the different target classes. In scenario C, the range dependent cost function introduces a spatial mission priority, such that the consequence of decision errors depends on the target location. These scenarios are less straightforward to interpret with a simple allocation rule, because the best allocation is no longer determined only by target uncertainty or sensing quality, but also by the mission dependent consequences of different decisions.

The larger improvements in these mission related scenarios show that the proposed framework is particularly useful when the user defined cost function creates a structured and non uniform risk landscape. In such cases, the optimized allocation can exploit the mission structure by assigning more resources to the targets, classes, or regions where additional sensing effort has the largest expected impact on the Bayes Risk. At the same time, the optimized allocation also improves performance in the baseline scenario, where equal budget allocation is already a reasonable reference strategy. This indicates that the method provides a consistent benefit, with the largest gains obtained when the mission objectives introduce clear operational priorities.

The absolute risk differences should be interpreted together with the scale of the corresponding risk values. For example, scenario A-2 has a smaller absolute difference than the baseline case because fewer decision errors are penalized by the modified cost matrix. Therefore, both the relative improvement and the number of better runs should be used to assess the consistency of the optimized allocation strategy.

6
Conclusion

6.1. Summary of contributions

In this thesis, a Bayes Risk-based RRM framework for joint target tracking and classification was proposed. The proposed framework from this thesis extends existing the Bayes Risk-based track maintenance formulations from Chioccarello *et al.* [15] towards a combined multi hypothesis framework in which track maintenance and classification are considered simultaneously under a time budget constraint.

A novel joint Bayesian hypothesis formulation was introduced in which target absence and target classes are represented within one combined belief state. Using this formulation, the radar system recursively updates the target hypotheses using radar measurements, containing spatial measurements, as well as class dependent feature information. By coupling this belief state directly to the radar resource allocation process, the framework can dynamically adapt the allocated resources according to both track uncertainty and it's operational relevance.

To be able to include mission specific priorities, expanded Bayesian cost matrices were incorporated into the framework such that different tracking and classification decision errors can be directly weighted with their operational costs. This allows the radar resource allocation process to be controlled very directly through operationally interpretable penalties rather than purely through estimation accuracy metrics or an intermediate threat variable. Resulting in predictable and physically interpretable time allocation behavior across the given different operational conditions.

The resulting Bayes Risk functions were approximated using Monte Carlo sampling methods. However, directly optimizing these stochastic risk estimates proved to be challenging due to the Monte Carlo noise that was present in the evaluated risk curves. To circumvent this issue, a surrogate risk optimization framework was introduced. With these surrogate risk functions, the noisy risk estimates were transformed into smooth monotonic surrogate functions using smoothing and monotonic regression techniques. These estimated surrogate functions finally enabled stable and computationally feasible optimization of the adaptive dwell time allocation problem while satisfying the constrained global radar time budget.

6.2. Main findings

The simulation results show that the proposed framework produces predictable and interpretable adaptive sensing behavior across multiple multi-target scenarios. As shown in chapter 5, the optimizer distributes the radar resources according to target uncertainty, sensing conditions, and the operational costs encoded in the Bayesian cost matrices.

The simulation results show that changing the entries of the mission specific cost matrix directly changes the allocation behavior. By assigning different costs to target-existence errors, classification errors, target classes, and spatial regions, certain mistakes can be made more important than others in the Bayes risk objective. When both target existence and classification errors are penalized, both types of uncertainty drive the allocation together. When classification errors are assigned little or no cost, the framework no longer allocates resources specifically to reduce classification uncertainty.

At the same time, asymmetric and range dependent costs shift resources toward targets with higher operational relevance. This demonstrates that the proposed framework adapts to the user specified operational cost matrix.

The Monte Carlo performance analysis further showed that the optimized allocation reduced the cumulative realized Bayes risk compared with the fixed equal budget allocation in all considered scenarios. The largest improvements were obtained in the more mission specific scenarios, where the cost function was changed to reflect different operational priorities. These cases are difficult to handle with simple heuristic allocation rules, because the preferred allocation depends not only on target uncertainty, but also on the cost assigned to different decision errors. This shows the value of the proposed framework: it can use the structure introduced by the mission specific cost matrix to allocate resources where they are expected to reduce operational risk the most.

The target appearance and disappearance scenarios showed that the framework is robust under changing target presence conditions. Newly appearing targets were incorporated into the adaptive allocation process through the joint belief update, while targets that stopped generating measurements were gradually assigned less radar resources once the belief in target absence increased. This demonstrates that target existence estimation is naturally integrated into the same Bayes Risk-based framework as tracking and classification.

Taken together, the results show that the main research question of this thesis can be answered positively. A Bayes risk-based objective function for joint track maintenance and classification can be designed by representing target absence and target classes within a shared Bayesian hypothesis space, evaluating the expected decision cost over all hypotheses and decisions, and using this risk as the target-dependent objective in the radar resource allocation problem. Operationally different decision errors can then be incorporated directly through direct user defined costs defined in Bayesian cost matrices, which makes the resulting allocation behavior mission dependent, interpretable, and controllable.

It is also shown how tracking and classification can be coupled within a single resource management objective. Both track quality and classification quality influence the predicted belief state and depend on the allocated sensing resources through the measurement likelihood, detection probability, and measurement quality. As a result, the optimizer automatically balances track maintenance and classification requirements according to their expected contribution to the Bayes risk.

At the same time, the simulation results showed several important limitations of the current implementation. Most notably, the current optimization framework operates using a one step optimization horizon, which can lead to myopic allocation behavior. As discussed in the performance analysis in chapter 5, the optimizer primarily minimizes the instantaneous expected Bayes risk at each optimization interval without explicitly accounting for the longer term evolution of the belief states. The resulting allocation choices are interpretable, but they do not necessarily minimize the realized cumulative risk over all time intervals.

The results of this thesis also show an important limitation related to Bayes Risk-based classifiers. Although Bayes risk formulations provide strong operational interpretability and flexibility, there exist no practical optimization methods for obtaining analytical decision boundaries under complex Bayes risk formulations. As a result, the framework currently relies heavily on computationally expensive Monte Carlo evaluation methods to approximate the resulting risk behavior. This motivates the use of surrogate risk approximations and demonstrates the direction of developing computationally efficient optimization techniques for Bayes Risk-based adaptive sensing frameworks.

6.3. Recommendations for future work

Several directions for future work have been found during the research. A particularly important extension would be the application of multi step optimization strategies or policy rollout approaches within the underlying POMDP framework. It is assumed that such approaches could allow the optimizer to explicitly account for future uncertainty evolution and therefore reduce the myopic behavior observed in the current implementation. Such approaches have already been shown in previous work by Schöpe *et al.* [39], suggesting that these methods could likely also improve the long term allocation behavior.

Future work could also investigate analytical approximations or boundary based optimization methods for Bayes Risk-based classifiers. The current absence of practical analytical boundary formulations for non binary Bayes Risk problems increases the computational complexity of the optimization process. This currently motivates the continued use of Monte Carlo-based surrogate methods.

Another direction for future work is the joint optimization of multiple radar resource variables simultaneously. In the current implementation, the optimization was limited to adaptive dwell time allocation only, while the other controllable radar parameter the revisit interval was kept fixed during the simulations. Although this simplification allowed stable evaluation of the proposed Bayes risk framework, jointly optimizing both resource variables could improve the radar performance even more.

A related practical limitation is that no scheduling model was implemented in this thesis. The optimization produces desired dwell time allocations under a global time budget constraint, but it does not verify whether the resulting dwell and revisit time choices can always be scheduled exactly by a physical radar system. Especially when using different revisit times for all the targets, more difficulties are expected. Future work could thus include a scheduler that translates the optimized allocation into a feasible radar timeline. This would make it possible to investigate how closely a practical scheduler can approximate the optimized solution, and how scheduling constraints influence the achieved Bayes risk.

Finally, the results of this thesis demonstrated that the behavior for the resource allocation is strongly mission dependent, since the selected operational cost structure directly determines which decision errors are considered most important. This cost sensitivity should therefore be interpreted as a key feature of the proposed framework: by translating mission objectives in the Bayesian cost matrix, the radar resource manager can steer the sensing process toward the decisions with the highest operational relevance. Future work into mission dependent cost formulations, and operationally derived Bayesian objectives may therefore provide valuable further insights into practical radar resource management system design and it’s applications.

To conclude, this thesis demonstrated that Bayes Risk-based radar resource management provides a flexible, interpretable, and operationally controllable framework for the joint optimization of tracking and classification tasks under a limit amount of sensing resources. The proposed framework forms a promising foundation for future adaptive multi-function radar management systems in which operational decision making objectives are directly integrated into the sensing process itself.

References

[1] S. Haykin, “Cognitive radar: A way of the future,” *IEEE signal processing magazine*, vol. 23, no. 1, pp. 30–40, 2006.

[2] A. O. Hero and D. Cochran, “Sensor management: Past, present, and future,” *IEEE Sensors Journal*, vol. 11, no. 12, pp. 3064–3075, 2011.

[3] R. J. Mailloux, *Phased array antenna handbook*. Artech house, 2017.

[4] J. E. Stailey and K. D. Hondl, “Multifunction phased array radar for aircraft and weather surveillance,” *Proceedings of the IEEE*, vol. 104, no. 3, pp. 649–659, 2016. DOI: `10.1109/JPROC.2015.2491179`.

[5] A. Charlish, F. Hoffmann, C. Degen, and I. Schlangen, “The development from adaptive to cognitive radar resource management,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 35, no. 6, pp. 8–19, 2020. DOI: `10.1109/MAES.2019.2957847`.

[6] P. Moo and Z. Ding, *Adaptive radar resource management*. Academic Press, 2015.

[7] M. E. Weber, J. Y. N. Cho, and H. G. Thomas, “Command and control for multifunction phased array radar,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 10, pp. 5899–5912, 2017. DOI: `10.1109/TGRS.2017.2716935`.

[8] M. I. Schöpe, “Approximately optimal resource management for multi-function radar: Algorithmic solutions using a generic framework,” PhD thesis, Delft University of Technology, Delft, The Netherlands, 2021. DOI: `10.4233/uuid:2ba49af1-fa17-476c-88f4-97783ca4e39a`.

[9] W. Yi, Y. Yuan, R. Hoseinnezhad, and L. Kong, “Resource scheduling for distributed multi-target tracking in netted colocated mimo radar systems,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 1602–1617, 2020.

[10] C. Kreucher, K. Kastella, and A. O. Hero Iii, “Sensor management using an active sensing approach,” *Signal Processing*, vol. 85, no. 3, pp. 607–624, 2005.

[11] K. White, J. Williams, and P. Hoffensetz, “Radar sensor management for detection and tracking,” in *2008 11th International Conference on Information Fusion*, IEEE, 2008, pp. 1–8.

[12] S. M. Sowelam and A. H. Tewfik, “Optimal waveform selection for radar target classification,” in *Proceedings of International Conference on Image Processing*, IEEE, vol. 3, 1997, pp. 476–479.

[13] F. Katsilieris, H. Driessen, and A. Yarovoy, “Threat-based sensor management for target tracking,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 4, pp. 2772–2785, 2015.

[14] F. Katsilieris, H. Driessen, and A. Yarovoy, “Threat-based sensor management for joint target tracking and classification,” in *2015 18th International Conference on Information Fusion (Fusion)*, IEEE, 2015, pp. 435–442.

[15] S. Chioccarello, H. Driessen, and A. Yarovoy, “Exploration of bayes risk-based radar resource management for multi-target track maintenance,” in *2025 IEEE Radar Conference (RadarConf25)*, IEEE, 2025, pp. 1128–1133.

[16] M. I. Schöpe, H. Driessen, and A. Yarovoy, “Optimal threat-based radar resource management for multitarget joint tracking and classification,” *Journal of Advances in Information Fusion*, vol. 17, no. 1, p. 46, 2022.

[17] S. Chioccarello, H. Driessen, and A. Yarovoy, *Risk-based radar resource management for multi-target track maintenance*, Under review, 2026.

[18] N. J. Gordon, S. Maskell, and T. Kirubarajan, “Efficient particle filters for joint tracking and classification,” in *Signal and Data Processing of Small Targets 2002*, SPIE, vol. 4728, 2002, pp. 439–449.

[19] S. Challa and G. Pulford, “Joint target tracking and classification using radar and esm sensors,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, no. 3, pp. 1039–1055, 2001. DOI: `10.1109/7.953266`.

[20] X. Cheng, H. Ji, and Y. Zhang, “Multiple extended target joint tracking and classification based on gps and lmb filter,” *IEEE Signal Processing Letters*, vol. 31, pp. 1139–1143, 2024. DOI: `10.1109/LSP.2024.3383954`.

[21] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*. Springer, 2006, vol. 4.

[22] G. C. Anagnostopoulos, “Svm-based target recognition from synthetic aperture radar images using target region outline descriptors,” *Nonlinear Analysis: Theory, Methods & Applications*, vol. 71, no. 12, e2934–e2939, 2009.

[23] N. Ibrahim, R. Abdullah, and M. Saripan, “Artificial neural network approach in radar target classification,” *Journal of Computer Science*, vol. 5, no. 1, p. 23, 2009.

[24] A. Huizing, M. Heiligers, B. Dekker, J. de Wit, L. Cifola, and R. Harmanny, “Deep learning for classification of mini-uavs using micro-doppler spectrograms in cognitive radar,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 34, no. 11, pp. 46–56, 2019. DOI: `10.1109/MAES.2019.2933972`.

[25] T. Vercauteren, D. Guo, and X. Wang, “Joint multiple target tracking and classification in collaborative sensor networks,” *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 4, pp. 714–723, 2005. DOI: `10.1109/JSAC.2005.843540`.

[26] D. Angelova and L. Mihaylova, “Joint target tracking and classification with particle filtering and mixture kalman filtering using kinematic radar information,” *Digital Signal Processing*, vol. 16, no. 2, pp. 180–204, 2006.

[27] E. K. Chong, C. M. Kreucher, and A. O. Hero III, “Partially observable markov decision process approximations for adaptive sensing,” *Discrete Event Dynamic Systems*, vol. 19, no. 3, pp. 377–422, 2009.

[28] S. Ross, J. Pineau, S. Paquet, and B. Chaib-Draa, “Online planning algorithms for pomdps,” *Journal of Artificial Intelligence Research*, vol. 32, pp. 663–704, 2008.

[29] H. V. Poor, *An introduction to signal detection and estimation*. Springer Science & Business Media, 2013.

[30] F. Nielsen, “Generalized bhattacharyya and chernoff upper bounds on bayes error using quasi-arithmetic means,” *Pattern Recognition Letters*, vol. 42, pp. 25–34, 2014.

[31] D. Andrade and Y. Okajima, “Efficient bayes risk estimation for cost-sensitive classification,” in *The 22nd international conference on artificial intelligence and statistics*, PMLR, 2019, pp. 3372–3381.

[32] M. Lázaro and A. R. Figueiras-Vidal, “A bayes risk minimization machine for example-dependent cost classification,” *IEEE transactions on cybernetics*, vol. 51, no. 7, pp. 3524–3534, 2019.

[33] M. Hellman and J. Raviv, “Probability of error, equivocation, and the chernoff bound,” *IEEE Transactions on Information Theory*, vol. 16, no. 4, pp. 368–372, 1970.

[34] P. A. Devijver, “On a new class of bounds on bayes risk in multihypothesis pattern recognition,” *IEEE Transactions on Computers*, vol. 100, no. 1, pp. 70–80, 1974.

[35] S. Kirmani, “A lower bound on bayes risk in classification problems,” *Annals of the Institute of Statistical Mathematics*, vol. 28, no. 1, pp. 385–387, 1976.

[36] F. D. Garber and A. Djouadi, “Bounds on the bayes classification error based on pairwise risk functions,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 2, pp. 281–288, 1988.

[37] A. Wisler, V. Berisha, D. Wei, K. Ramamurthy, and A. Spanias, “Empirically-estimable multi-class classification bounds,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2016, pp. 2594–2598.

[38] L. A. Dalton and M. R. Yousefi, “On optimal bayesian classification and risk estimation under multiple classes,” *EURASIP Journal on Bioinformatics and Systems Biology*, vol. 2015, no. 1, p. 8, 2015.

[39] M. I. Schöpe, H. Driessen, and A. Yarovoy, “Multi-task sensor resource balancing using lagrangian relaxation and policy rollout,” in *2020 IEEE 23rd International Conference on Information Fusion (FUSION)*, IEEE, 2020, pp. 1–8.

[40] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[41] A. Savitzky and M. J. Golay, “Smoothing and differentiation of data by simplified least squares procedures.,” *Analytical chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.

[42] J. De Leeuw, K. Hornik, and P. Mair, “Isotone optimization in r: Pool-adjacent-violators algorithm (pava) and active set methods,” *Journal of statistical software*, vol. 32, pp. 1–24, 2010.

[43] F. N. Fritsch and J. Butland, “A method for constructing local monotone piecewise cubic interpolants,” *SIAM journal on scientific and statistical computing*, vol. 5, no. 2, pp. 300–304, 1984.

[44] F. N. Fritsch and R. E. Carlson, “Monotone piecewise cubic interpolation,” *SIAM Journal on Numerical Analysis*, vol. 17, no. 2, pp. 238–246, 1980.

[45] H. Meikle, *Modern radar systems*. Artech House, 2008.

[46] P. Swerling, “Probability of detection for fluctuating targets,” *IRE Transactions on Information Theory*, vol. 6, no. 2, pp. 269–308, 1960. DOI: `10.1109/TIT.1960.1057561`.

[47] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2001.

[48] T. Kirubarajan and Y. Bar-Shalom, “Probabilistic data association techniques for target tracking in clutter,” *Proceedings of the IEEE*, vol. 92, no. 3, pp. 536–557, 2004.