

**Collision Avoidance System Optimization for
Closely Spaced Parallel Operations through
Surrogate Modeling**

by

Kyle A. Smith

B.S., United States Air Force Academy (2011)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

This material is declared a work of the U.S. Government and is not
subject to copyright protection in the United States.

Author
Department of Aeronautics and Astronautics
April 10, 2013

Certified by.....
Mykel J. Kochenderfer
Technical Staff, Surveillance Systems Group, MIT Lincoln Laboratory
Thesis Supervisor

Certified by.....
Jonathan P. How
Richard C. Maclaurin Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by.....
Eytan H. Modiano
Professor of Aeronautics and Astronautics
Chair, Graduate Program Committee

Disclaimer:

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

Collision Avoidance System Optimization for Closely Spaced Parallel Operations through Surrogate Modeling

by

Kyle A. Smith

Submitted to the Department of Aeronautics and Astronautics
on April 10, 2013, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

The Traffic Alert and Collision Avoidance System (TCAS) is mandated worldwide to protect against aircraft mid-air collisions. One drawback of the current TCAS design is limited support for certain closely spaced parallel runway operations. TCAS alerts too frequently, leading pilots to often inhibit Resolution Advisories during approach. Research is underway on the Airborne Collision Avoidance System X (ACAS X), a next-generation collision avoidance system that will support new surveillance systems and air traffic control procedures. ACAS X has been shown to outperform TCAS for enroute encounter scenarios. However, the design parameters that are tuned for the enroute environment are not appropriate for closely spaced parallel operations (CSPO).

One concept to enhance the safety of CSPO is a procedure-specific mode of the logic that minimizes nuisance alerts while still providing collision protection. This thesis describes the application of surrogate modeling and automated search for the purpose of tuning ACAS X for parallel operations. The performance of the tuned system is assessed using a data-driven blunder model and an operational performance model. Although collision avoidance system development normally relies on human judgment and expertise to achieve ideal behavior, surrogate modeling is efficient and effective in tuning ACAS X for CSPO as the tuned logic outperforms TCAS in terms of both safety and operational suitability.

Thesis Supervisor: Mykel J. Kochenderfer

Title: Technical Staff, Surveillance Systems Group, MIT Lincoln Laboratory

Thesis Supervisor: Jonathan P. How

Title: Richard C. Maclaurin Professor of Aeronautics and Astronautics

Acknowledgments

I would like to thank MIT Lincoln Laboratory for supporting my education and research over the past two years, especially Col. (ret) John Kuconis, Division 4 Director Dr. Israel Soibelman, Group 42 Leader Gregory Hogan, and Associate Group 42 Leader Dr. Gregg Shoults.

I am thankful for Dr. Wesley Olson's leadership and support as head of the Group 42 ACAS X program, and for continually guiding and supporting my work in every aspect. This work was sponsored by the FAA TCAS Program Office AJM-233, and I gratefully acknowledge Neal Suchy for his leadership and support.

I am extremely grateful to Professor Jonathan How as my advisor and thesis supervisor for looking out for my best interests over the past two years and ensuring that I stayed on target in successfully completing my coursework and research on time.

I am indebted to Dr. Mykel Kochenderfer for his continual guidance during my entire time at Lincoln Laboratory, including helping to provide direction for my research and invaluable critique of this thesis throughout the entirety of its development. Likewise, Dr. Adan Vela stood by my side during critical moments of research and writing and dedicated extensive effort and time to helping me complete quality work. I would not have succeeded without the hard work and commitment exhibited by these two gentlemen.

I am thankful to numerous members of Group 42 and others within Lincoln Laboratory who offered their expertise when I needed it most. Specifically, I wish to thank Jessica Holland, Dylan Asmar, James Chryssanthocopoulos, Tomas Elder, and Thomas Billingsley for their insights and assistance in completing this work. Chung Lee of Georgia Tech was a great help due to his early input regarding potential optimization tools. I am also appreciative of my officemate Melvin Stone for his seemingly endless supply of stories, advice, peanuts, and tea.

Finally, the love and support shown by my family and friends never ceases to amaze me. Throughout the years, they have played a critical role in building the foundations for who I am today and every one of my accomplishments, and I am forever grateful.

Contents

1	Introduction	11
1.1	Traffic Alert and Collision Avoidance System	11
1.2	Closely Spaced Parallel Operations	14
1.3	Collision Avoidance System Design	17
1.4	ACAS X	19
1.5	Objective: Tuning ACAS Xo for CSPO	20
1.6	Literature Review: Optimization Methods	21
1.7	Contributions and Outline	23
2	Worst-Case Analysis of Airborne Collision Avoidance Systems	25
2.1	Problem Description	26
2.2	Worst-Case Analysis	30
2.3	Discussion	35
3	Optimization via Surrogate Modeling	37
3.1	Performance Metrics	38
3.2	Historical Encounter Set	39
3.3	Simulation and Evaluation	42
3.4	Defining the Objective Value	43
3.5	Screening Study	44
3.6	Sampling Plan	46
3.7	Surrogate Modeling	47
3.8	Searching the Model	49

3.9 Discussion	53
4 Performance Analysis	55
4.1 Historical Encounter Set Performance	57
4.2 Individual Parameter Analysis	61
4.3 Generated Encounter Set Performance	63
4.4 Operational Performance Analysis	68
4.5 Policy Examples	73
4.6 Application of Worst-Case Analysis to ACAS Xo Development	76
4.7 Discussion	82
5 Conclusions and Further Work	85
A Objective Value Weighting Sweep	89
Bibliography	102

List of Figures

1-1	Example TA cockpit display and annunciation	12
1-2	Example RA cockpit display and annunciation	12
1-3	Aircraft trajectories and TCAS commands in the Überlingen mid-air collision	13
1-4	Example parallel approach scenario	15
1-5	Example dependent approach geometry	16
1-6	Example overtake geometry	16
1-7	Example oscillation geometry	17
1-8	Example TCAS pseudocode	18
1-9	Example slow-closure geometry	19
1-10	ACAS X optimization and development process	20
2-1	Encounter timeline	27
2-2	Climb maneuver	28
2-3	Constant-drift blunder	28
2-4	Constant-turn blunder	29
2-5	Worst-case available time to alert (t_c) vs. drift angle	31
2-6	Drift angle vs. perpendicular deviation	32
2-7	Worst-case available time to alert (t_c) vs. turn rate	33
2-8	Worst-case available time to alert (t_c) vs. turn rate	34
2-9	Turn rate vs. perpendicular deviation	34
3-1	Overview of the optimization via surrogate modeling process	38
3-2	Example encounter generation	40

3-3	Example aircraft behavior	41
3-4	Objective values of eighty infill points (solid) and the cumulative minimum value achieved (dashed)	51
4-1	Sample encounter	60
4-2	Sample encounter	61
4-3	Effect of cycles parameter	63
4-4	Generated encounter set configuration	65
4-5	Probability (in percent) of alert	69
4-6	No-blunder CSPO configuration (vertical sweep)	73
4-7	Example policy plots for Aircraft 1	74
4-8	No-blunder CSPO configuration (horizontal sweep)	75
4-9	Example horizontal alert probabilities for Aircraft 1	75
4-10	Sample encounter	77
4-11	Sample encounter	78
4-12	Sample encounter	79
4-13	Sample encounter	80
4-14	Worst-case available time to alert (t_c) vs. drift angle; Actual ACAS Xo alert time vs. drift angle	81
4-15	Worst-case available time to alert (t_c) vs. turn rate; Actual ACAS Xo alert time vs. turn rate	81
A-1	Risk ratio (blue) and nuisance alert rate (red) at the optimal point found for each weighting combination	101

List of Tables

3.1	Optimization Performance Metrics	39
3.2	Design Parameters Analyzed in the Screening Study	44
3.3	Design Parameter Screening Study Results (Mean)	45
3.4	Design Parameter Screening Study Results (Standard Deviation)	45
3.5	Design Parameter Ranges used for Optimization	47
3.6	Evaluation Results for the First Forty Infill Points	52
3.7	Evaluation Results for the Second Forty Infill Points	53
4.1	ACAS Xo Parameter Value Changes Relative to ACAS Xa	56
4.2	Metric Results for the Historical Blunder Encounter Set	58
4.3	Metric Results using ACAS Xa Settings for One Parameter at a Time	62
4.4	Metric Results for the Generated Blunder Encounter Set	66
4.5	Metric Results for the Generated Nominal Encounter Set	67
A.1	Evaluation Results for the First Forty Infill Points (97.5%/2.5%)	91
A.2	Evaluation Results for the Second Forty Infill Points (97.5%/2.5%)	92
A.3	Evaluation Results for the First Forty Infill Points (92.5%/7.5%)	93
A.4	Evaluation Results for the Second Forty Infill Points (92.5%/7.5%)	94
A.5	Evaluation Results for the First Forty Infill Points (90%/10%)	95
A.6	Evaluation Results for the Second Forty Infill Points (90%/10%)	96
A.7	Evaluation Results for the First Forty Infill Points (85%/15%)	97
A.8	Evaluation Results for the Second Forty Infill Points (85%/15%)	98
A.9	Evaluation Results for the First Forty Infill Points (80%/20%)	99
A.10	Evaluation Results for the Second Forty Infill Points (80%/20%)	100

List of Acronyms and Abbreviations

ACAS	Airborne Collision Avoidance System
CAT	Category
CSPO	Closely Spaced Parallel Operations
deg	degrees
ft	feet
ILS	Instrument Landing System
IMC	Instrument Meteorological Conditions
kt	knots
lbs	pounds
Max EI	Maximum Expected Improvement
min	minute
MLE	Maximum Likelihood Estimate
NAS	National Airspace System
NextGen	Next Generation Air Transportation System
NM	Nautical Mile
NMAC	Near Mid-Air Collision
NTZ	No Transgression Zone
PRM	Precision Runway Monitor
RA	Resolution Advisory
RNP	Required Navigation Performance
TA	Traffic Alert
TCAS	Traffic Alert and Collision Avoidance System

Chapter 1

Introduction

This thesis asserts that a surrogate modeling and automated search process can tune an airborne collision avoidance system for ideal alerting behavior during closely spaced parallel operations (CSPO). The goal is to produce a system that maintains or enhances the safety level of the current operational system while reducing unnecessary alerts, thereby decreasing pilot workload during parallel approaches. Improved alerting behavior will allow for greater operational efficiency during CSPO, leading to greater airport throughput and cost savings especially during Instrument Meteorological Conditions in which pilots and controllers cannot rely on visual separation between aircraft.

1.1 Traffic Alert and Collision Avoidance System

The Traffic Alert and Collision Avoidance System (TCAS) provides pilots with vertical avoidance maneuver guidance when a near mid-air collision is imminent. TCAS is currently mandated on all commercial aircraft with more than thirty seats or a maximum takeoff weight greater than 33,000 lbs, in addition to being used extensively by many business jet type aircraft [9]. TCAS is able to provide pilots with both Traffic Alerts (TAs) and Resolution Advisories (RAs) for any aircraft with an operating transponder. TAs warn pilots of nearby traffic that poses a potential threat and facilitate pilots in visually acquiring the traffic. TAs consist of an aural annunciation

(“Traffic, Traffic”) and traffic information such as relative altitude on a visual display in the cockpit, as depicted in Figure 1-1.



Figure 1-1: Example TA cockpit display and annunciation




Figure 1-2: Example RA cockpit display and annunciation

If a threat becomes imminent, an RA is issued to prevent a collision by commanding the pilot to execute a vertical avoidance maneuver (e.g., Climb, Descend, Level

Off). As shown in Figure 1-2, RAs include visual instructions such as target vertical speeds or pitch angles, as well as aural annunciations (e.g., “Climb, Climb”). TCAS functions independently of ground-based systems and air traffic controllers and relies on surveillance equipment onboard the aircraft [9].

If an aircraft has begun a maneuver in concordance with an RA, the ground controller is no longer responsible for separation between that aircraft and any other aircraft, airspace, or obstacle. Once the controller is aware of the RA, he or she is not authorized to instruct the pilot to maneuver contrary to the RA. Controller responsibility for separation is reinstated when either the aircraft returns to its previously assigned altitude or the pilot indicates that the RA maneuver is completed. The pilot in command is ultimately responsible for the safe operation of flight and should immediately respond to any RA unless a response would jeopardize safety or is deemed unnecessary by visual contact with the target aircraft [9]. To illustrate the importance of these responsibilities, Figure 1-3 shows the aircraft trajectories and TCAS commands in the 2002 mid-air collision over Überlingen, Germany. The conflicting instructions of TCAS (Climb) and air traffic control (Descend) resulted in the Russian Tu-154 passenger jet following the controller’s instructions and descending into the path of the Boeing B-757 [23].




Figure 1-3: Aircraft trajectories and TCAS commands in the Überlingen mid-air collision [23]

1.2 Closely Spaced Parallel Operations

The National Airspace System (NAS) is operated by the Federal Aviation Administration and consists of various units that regulate, coordinate, and supervise aircraft travel over the United States. The NAS is designed to enable safe and efficient aircraft operations through the interaction of airspace structure, technologies, policies, and standard procedures [40]. One such set of procedures is CSPO, in which multiple aircraft simultaneously approach parallel runways for landing at a given airport to increase airport throughput. Depending on the airport, parallel runways can be laterally spaced as close as 700 ft [28]. Uninterrupted and efficient CSPO is normally feasible during Visual Meteorological Conditions when pilots are able to maintain visual separation with other aircraft. However, Instrument Meteorological Conditions (IMC) are defined by varying degrees of reduced visibility and cloud clearance, depending on the type of airspace. As visibility deteriorates in IMC, safety cannot be ensured at smaller runway spacings and CSPO is restricted to a smaller set of runway configurations. This limited CSPO capability decreases airport throughput, causes delays in the air and on the ground, and decreases the overall efficiency in the NAS [7].

Executing CSPO often allows for higher airport throughput than would otherwise be achievable [7]. However, such operations are sometimes prohibited due to safety concerns, such as during IMC. Between October 2010 and March 2011, CSPO accounted for 13% of all unique RAs detected in U.S. high-tempo regions. Among the most affected regions were San Francisco, where 92.6% of RAs were during parallel approaches, Atlanta (56.3%), and Denver (54.1%) [28]. Under current regulations, independent parallel approaches, in which no minimum diagonal spacing is required, may only be conducted in IMC for runways separated laterally by at least 4300 ft [25]. This requirement is relaxed down to a minimum of 3000 ft if certain conditions are met for approaches equipped with the Precision Runway Monitor (PRM) system. PRM employs a ground-based high-update-rate radar that allows an air traffic controller to closely monitor aircraft tracks on parallel approaches. If either aircraft deviates, or

blunders, from its correct trajectory into a predefined No Transgression Zone (NTZ) during PRM, as illustrated in Figure 1-4, the controller will issue an alert to prevent a collision if TCAS has not already issued an RA [3]. These separation standards were developed by simulating worst-case blunder scenarios and determining the minimum acceptable runway separations [25].

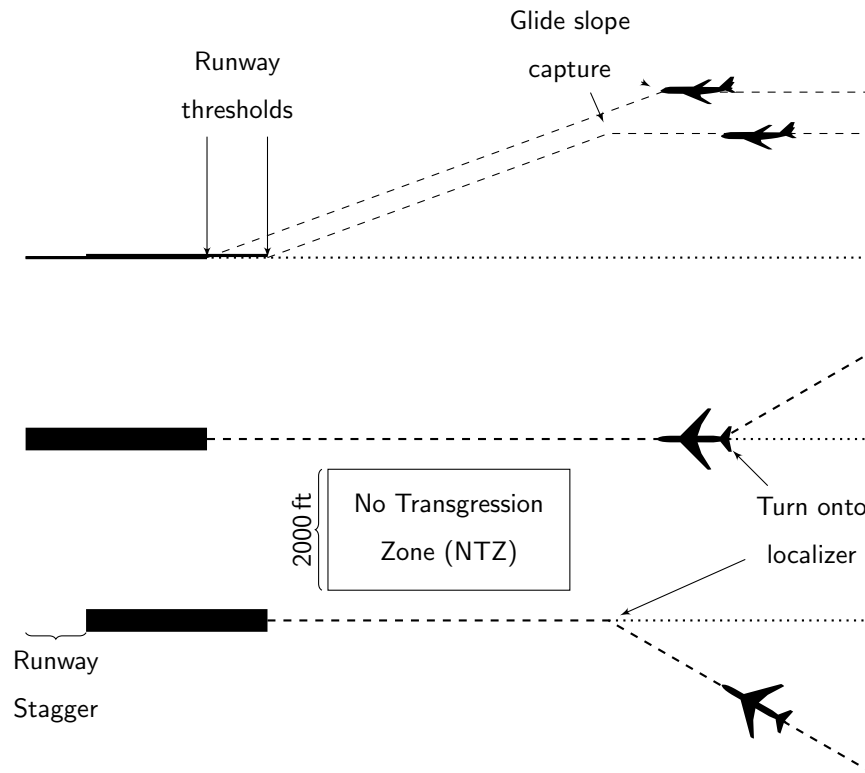


Figure 1-4: Example parallel approach scenario

PRM is only available at select U.S. airports and only allows dependent approaches, in which aircraft must maintain certain diagonal spacing in addition to in-trail spacing, for runway separations of at least 3000 ft [25]. Figure 1-5 gives an example of this dependent approach geometry.




Figure 1-5: Example dependent approach geometry [2]

PRM also requires an additional specially trained controller to monitor and provide separation assurance for the aircraft on parallel approach. This staffing requirement can be difficult to schedule, especially in the case of unanticipated weather events. Providing an alternate means of conducting CSPO in IMC for smaller runway separations, specifically an enhanced airborne collision avoidance system, would enable high throughput at a larger set of airports without requiring additional staffing or technology on the ground.

CSPO often involves unique aircraft geometries that require collision avoidance logic to respond with fine-tuned sensitivity. One example is an overtake scenario, shown in Figure 1-6, in which a trailing aircraft overtakes a lead aircraft on a parallel approach due to speed differential.




Figure 1-6: Example overtake geometry [9]

A second example involves navigation error during CSPO in IMC. For instance, an aircraft may appear to be initiating a blunder toward another aircraft on approach, as shown in Figure 1-7, but it may simply be oscillating along the approach path while

attempting to stabilize its track along the localizer during an Instrument Landing System (ILS) approach.




Figure 1-7: Example oscillation geometry

As a third example, an aircraft that seems to be aimed toward another aircraft on final approach may simply be angling to initially establish the localizer, as in Figure 1-4. These situations are typically safe and intentional, but the challenge lies in tuning the collision avoidance system such that it can distinguish between these safe scenarios and imminent threats due to deviations from the correct nominal flight paths.

1.3 Collision Avoidance System Design

There are a number of trade-offs that must be taken into account in the design of a collision avoidance system. For example, the system should alert frequently enough to maintain a desired level of safety. However, the system should also minimize nuisance alerts such that safe operations can be performed without being interrupted by an RA. There are several hazards associated with nuisance alerts. For example, when faced with too many unnecessary alerts during intentional and safe operations, pilots may become desensitized and distrust the collision avoidance system, potentially ignoring legitimate alerts [23, 31].

Collision avoidance systems are highly complex and a comprehensive understanding of their decision processes is not a trivial task for most common users. The fact

that pilots regularly utilize collision avoidance systems “implies that they must have made a leap of faith” [27]. This leap of faith is heavily based on experience with the system and its perceived predictability during risky events, leading the pilot to generalize the system’s dependability [24]. Perceived system unreliability can also lead to increased response times due to indecision about the best course of action [33]. Finally, excessive nuisance alerts may increase the potential for induced collisions where a pilot follows a nuisance RA, maneuvers unnecessarily, and creates an unsafe situation.

```

PROCESS Tau_calculation;

IF (ITF.R GT DMOD)
  THEN T1 = -(ITF.R-((DMOD**2)/ITF.R))/RDTEMP;
  ELSE T1 = P.MINTAU;
T3 = -ITF.R/RDTEMP;
IF (ITF.TAUCAP EQ $TRUE)
  THEN IF (T3 LT ITF.TRTUN)
    THEN ITF.TAURISE = 0;
    ELSE IF (ITF.R GT P.NAFRANGE OR (G.CLSTRONG EQ 8 AND
      G.CLSTROLD EQ 8) OR (G.DESTRONG EQ 8 AND
      G.DESTROLD EQ 8))
      THEN ITF.TAURISE = ITF.TAURISE + 1;
    ITF.TAUR = MIN(T1, ITF.TAUR);
    ITF.TRTRU = MIN(T3, ITF.TRTRU);
  ELSE ITF.TAUR = T1;
    ITF.TRTRU = T3;
    ITF.TAURISE = 0;
ITF.TAUR = MAX(P.MINTAU, ITF.TAUR);
ITF.TRTRU = MAX(P.MINTAU, ITF.TRTRU);
ITF.TRTUN = T3;
IF (ITF.HFIRM GE P.MININITHFIRM)
  THEN SET ITF.TAUCAP;

END Tau_calculation;

```

Figure 1-8: Example TCAS pseudocode [34]

TCAS follows a complex set of heuristic rules, as illustrated in the pseudocode of Figure 1-8, that are difficult to adjust in concordance with airspace and procedure changes over time. At its foundation, the TCAS logic uses the projected time until closest point of approach to determine when to issue an RA. However, during slow-closure situations, such as in Figure 1-9, the projected time until closest point of approach may be too large to provide adequate protection against an intruder aircraft. In these circumstances, TCAS uses an enlarged critical region based on the aircraft altitude. These ranges are defined as follows: 0.2NM below 2350 ft Above

Ground Level, 0.35 NM below 5000 ft Mean Sea Level, and 0.55 NM between 5000 ft and 10,000 ft Mean Sea Level [9]. Although the thresholds defining the critical region tend to provide an adequate level of safety, rigid rules like these can result in undesirable alert rates and behavior during CSPO [23, 30, 37].




Figure 1-9: Example slow-closure geometry [9]

1.4 ACAS X

As the airspace evolves with the introduction of the Next Generation Air Transportation System (NextGen), research is underway to field a system known as the Airborne Collision Avoidance System X (ACAS X) to replace TCAS. In contrast to the rule-based logic of TCAS, ACAS X uses probabilistic models and cost functions to determine an optimal collision avoidance action [21]. Four variants of ACAS X are in development for different users and objectives. The first variant is ACAS Xa, which uses active secondary radar supplemented with passive surveillance. ACAS Xa is the analogue to TCAS and is intended to be used by current TCAS-mandated users during the majority of flight operations. ACAS Xo, the subject of this thesis, is tuned for specific operations such as CSPO. ACAS Xo issues procedure-specific RAs for designated aircraft while ACAS Xa maintains global protection against all other aircraft in the vicinity. The initial version of ACAS Xo uses the same surveillance as ACAS Xa, while future versions may incorporate data obtained directly from other aircraft such as bank angle or autopilot status. The two other variants, ACAS Xp

and Xu, are designed for general aviation aircraft and unmanned aircraft, respectively.

The optimization and development process of ACAS X is depicted in Figure 1-10 [21]. ACAS X differs from TCAS in that it uses a probabilistic dynamic model of aircraft behavior and optimizes the alerting behavior with respect to an offline cost function defined by numerous design parameters (approximately forty). The aim of the cost function is to balance safety and operational goals such as minimizing collision risk and nuisance alerts. With an appropriate cost function and probabilistic model, a dynamic programming approach is used to precompute an offline look-up table and optimal system parameters (millions). Using the offline look-up table and additional online parameters in real-time, the optimal action is determined at each timestep.




Figure 1-10: ACAS X optimization and development process

1.5 Objective: Tuning ACAS Xo for CSPO

CSPO is the initial target application for ACAS Xo. The goal of ACAS Xo is to provide additional protection during CSPO, especially in IMC, while minimizing unnecessary alerts. Providing protection with minimal alerting will increase the efficiency of parallel approaches during IMC and help achieve NextGen goals such as fewer flight delays, greater cost savings, and increased throughput at high-volume airports [17]. TCAS has been able to maintain adequate safety levels during CSPO [23], but it is also plagued with excessive alert rates which cause pilots to inhibit RAs by switching to a TA-only mode [28]. To be operationally suitable for CSPO, ACAS Xo must maintain the safety level of TCAS while improving the overall alerting behavior, specifically by reducing nuisance alerts.

The performance of ACAS Xo depends on design parameters that can be tuned to modify various behaviors of the collision avoidance logic. Hand-tuning these design parameters is challenging since the impact of changing parameters can be difficult to predict, usually resulting in sub-optimal system behavior. Instead, optimization via surrogate modeling is applied to tune parameters for desirable ACAS Xo performance during CSPO. The safety and alerting performance of ACAS Xo on a set of CSPO encounters is essentially a “black-box” function that is dependent on the individual design parameter values. Using surrogate modeling to tune ACAS Xo allows for a more efficient and effective process that results in ideal alerting behavior during CSPO encounters.

The global optimization process used here efficiently selects data points to test in order to model the “black-box” function that is unknown to the user [11]. The function value at a given point is unknown until a simulation is performed for a certain combination of parameter settings. Exploring the design space improves the function prediction, thereby allowing the search method to identify promising areas of the design space (i.e., parameter settings) to search [18]. Eventually the design space is explored thoroughly enough to determine the location of the global optimum.

The objective of this thesis is to apply the aforementioned global optimization to ACAS Xo such that it achieves acceptable performance on the encounter set used for tuning as well as additional independent CSPO encounter sets that represent various aircraft behaviors and procedures. Furthermore, the tuning process should result in ACAS Xo outperforming TCAS in terms of both safety and operational suitability in all encounter sets used for analysis.

1.6 Literature Review: Optimization Methods

Variations of this global optimization process have been used in many different application areas. One example is a model for flame velocity (output) based on five different chemical reaction rates (input). The objective was to tune a computer fluid-dynamics model to match physical data for flame velocity. High computational cost

for each run dictated the need for careful selection of data points to test [35]. Another application has been in minimizing the drag coefficient of an airfoil (output) defined by various inputs such as thickness-to-chord ratio. Expensive computer simulations must be run to determine the drag coefficient at each distinct input setting [10].

There are many different types of surrogate models that can be used for tuning. Kriging models are often used due to the flexibility and high dimensionality that they allow. Kriging was originally developed for geostatistics, where core samples were taken at carefully selected positions with the goal of predicting the concentration of a mineral at a given location [18]. The model has been applied to a wide range of topics and represents a special case of radial basis functions. The Kriging model is used extensively in work by Forrester et al., Jones et al., Sacks et al., and Schonlau, who sometimes refer to the model as a “stochastic process model” [10, 11, 18, 35, 36]. Sacks et al. explored various methods of determining the Kriging model parameters that best represent a given data set, including Integrated Mean Squared Error and Maximum Mean Squared Error [35].

Jones et al. developed a process called “Efficient Global Optimization” which uses a Kriging model and search criteria that maximizes expected improvement for each evaluation point [18]. Forrester et al. expanded the scope of global optimization work by applying Kriging and max expected improvement criteria to noisy computer experiments and adjusting the estimate of prediction error [10, 11]. Regis and Shoemaker developed a global optimization method subject to nonlinear constraints. Additional constraints were imposed which force the search method to cycle between global and local search by specifying the next evaluation point’s minimum distance from previous data points [32]. This thesis mainly utilizes methods developed by Forrester et al. to show that global optimization can be used to tune ACAS Xo for optimal performance during CSPO [11].

On a broader level, there are numerous stochastic optimization methods that do not rely on an already existing model of the “black-box” function. A genetic algorithm is one such method which uses ideas derived from evolution to find an optimal solution, as outlined by Goldberg and Whitley [12, 42]. A process of selection is applied to each

population to determine the best solutions. Recombination and mutation processes are applied to the population to generate a new population that is ideally comprised of more promising solutions.

Alternatively, Kirkpatrick et al. and Press et al. describe Simulated Annealing as an optimization method that bases its implementation on thermodynamic principles [19, 29]. The goal is to navigate around many local optima within a large design space to find the approximate location of the global optimum. An annealing schedule is generated which dictates the direction and magnitude of change in parameter settings. An important feature of Simulated Annealing is that, in a minimization problem for example, the algorithm may accept sample points with higher objective values, but with decreasing probability over time.

Gaussian Adaptation is an additional optimization method that is similar to Simulated Annealing with respect to the decreasing probability of accepting less desirable points over time [20]. Drawing from its name, Gaussian Adaptation adapts a Gaussian distribution to regions of feasible points in the design space. Certain regions in the design space are designated as feasible based on whether associated samples satisfy a criterion that becomes more demanding over time. This increasingly demanding criterion drives the aforementioned decreasing probability of accepting points with higher objective values (when minimization is desired).

1.7 Contributions and Outline

This thesis offers several contributions to the problem of designing airborne collision avoidance systems:

- Worst-case scenarios are identified for various blunder types and severities, and the available alert time is calculated for each scenario. The worst-case available alert times help guide the development of collision avoidance systems by defining alerting requirements at various parallel approach configurations.
- Optimization via surrogate modeling is proven as a viable global optimization

process for tuning collision avoidance systems. Given a collision avoidance logic dependent on a small number of design parameters, the logic's performance can be tuned to behave optimally for specific operations and procedures.

- The tuning process is shown to not only result in acceptable ACAS Xo performance, but also in exceptional performance on independent encounter sets compared to that of TCAS. As ACAS Xo development continues in the future, the tuning process can be reapplied to update the design parameter settings based on revised versions of the logic.

The remainder of this thesis is organized as follows:

- Chapter 2 completes a worst-case analysis of CSPO. The alerting limits of collision avoidance systems during these worst-case scenarios are determined for varying aircraft behaviors, runway configurations, and pilot responses. These collision avoidance limits help guide the development of ACAS Xo and establish realistic expectations for performance in worst-case scenarios.
- Chapter 3 provides an overview of the surrogate modeling and tuning process, as well as the encounter set used for simulations. The results of the tuning process are presented in preparation for more detailed performance analysis in Chapter 4.
- Chapter 4 analyzes and compares the safety and alerting performance of TCAS and ACAS Xo. Additional encounter sets are utilized to compare TCAS and ACAS Xo alerting ranges and probabilities in certain approach configurations. Performance contributions of individual ACAS Xo parameters are analyzed, and the ACAS Xo logic is tested in various worst-case scenarios from Chapter 2. The alerting behavior of ACAS Xo is also compared to the limits determined in Chapter 2.
- Chapter 5 concludes and suggests areas for further work related to the limits of collision avoidance and tuning ACAS Xo for CSPO.

Chapter 2

Worst-Case Analysis of Airborne Collision Avoidance Systems

During CSPO, there may exist certain blunders and aircraft configurations that comprise “worst-case” scenarios in which a collision avoidance system may be unable to alert in time to prevent a near mid-air collision (NMAC). This inability to alert in time may be caused by several factors, most notably: a small period of time between blunder initiation and NMAC, and the desire to minimize nuisance alerts. Minimizing nuisance alerts is only desirable until safety is undermined. The worst-case scenarios analyzed in this chapter aid in understanding where this threshold lies. Understanding the alerting requirements of collision avoidance systems during CSPO is helpful in guiding the development of ACAS Xo and setting realistic expectations for its performance.

To extract certain worst-case scenarios, two types of basic aircraft blunders relevant to CSPO are simulated for the intruder aircraft, as well as a corresponding vertical avoidance maneuver for the own aircraft. By simulating these aircraft dynamics within feasible parallel approach configurations, scenarios are extracted for which an NMAC is only avoidable if the collision avoidance system issues an RA within a small timeframe.

2.1 Problem Description

A general encounter framework is applied from which worst-case scenarios can be extracted. Each encounter consists of two aircraft on final approach to parallel runways separated by distance D . The own aircraft is defined to be on final approach to the left runway and only strays from its nominal flight path to execute a vertical avoidance maneuver. The intruder aircraft is defined to be on approach to the right runway and is simulated to blunder into the own aircraft's flight path. Both aircraft maintain a constant speed (v_1 and v_2 , respectively) and a 3° glide slope, with the exception of the own aircraft which may deviate from its glide slope to execute an avoidance maneuver. Only the own aircraft is simulated as equipped with a collision avoidance system, as the intruder does not terminate its blunder maneuver once initiated.

Two types of blunders are simulated: constant-drift and constant-turn. The constant-drift blunder is characterized by the intruder flying a normal final approach and suddenly altering its course to the left by a constant drift angle β . The intruder maintains its new course across the own aircraft's approach path. The constant-turn blunder is characterized by the intruder flying a normal final approach and then initiating and maintaining a constant-rate turn to the left with turn rate ω .

The possible avoidance maneuvers for the own aircraft are limited to a constant-rate climb or descent. The own aircraft dynamics [43] are modeled such that the pilot applies a constant-acceleration control input until the target rate of climb or descent is achieved. Unless otherwise stated, the target vertical rate for climbs/descents is fixed at ± 2500 ft/min since that is the greatest vertical rate commanded by TCAS and ACAS Xo. Furthermore, a 5 s pilot response delay (PD) is assumed unless otherwise stated. This delay is the elapsed time between the initial RA and the initial pilot response to the RA. For simplicity, no horizontal avoidance maneuvers are simulated. This simplification aligns with the implementation of TCAS and ACAS Xo, both of which restrict RAs to the vertical dimension [9, 21].

The basis of measurement for the limits of collision avoidance is the available time

to alert in order to avoid an NMAC, which is defined as when aircraft come within 500 ft horizontally (h_{nmac}) and 100 ft vertically (v_{nmac}) of each other. This calculation requires knowledge of the theoretical time until NMAC, assuming the own aircraft does not execute an avoidance maneuver, and the time required for the own aircraft to respond to an RA and climb or descend to a clear-of-conflict altitude. These values are represented by T and t_m , respectively. The difference of these two values results in the time available for the collision avoidance system to issue an RA before an NMAC is unavoidable, represented by t_c :

$$t_c = T - t_m. \tag{2.1}$$

The point at which $t_c = 0$ indicates where an NMAC is unavoidable unless an RA is issued before the blunder is initiated. Figure 2-1 illustrates the relationship between T , t_m , and t_c .

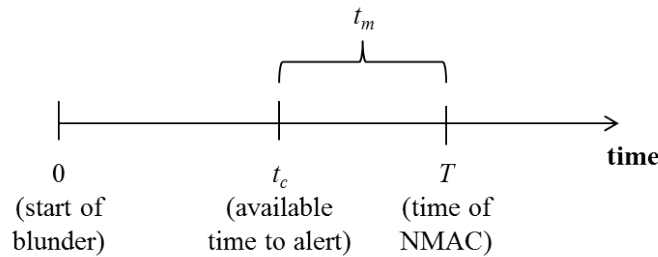


Figure 2-1: Encounter timeline

Figure 2-2 illustrates an example vertical profile over time of the own aircraft (blue) and intruder aircraft (red) after the intruder has begun to blunder. After the intruder begins to blunder at time zero, the own aircraft receives an RA at time t_c and initiates a climb avoidance maneuver after the pilot response delay PD . In this case, since the RA was issued at time t_c , the own aircraft achieves clear-of-conflict vertical separation v_{nmac} at time T . If the RA had been issued any later, the own aircraft would not have had enough time to maneuver to a safe altitude. If the RA had been issued earlier, separation would have been achieved with altitude to spare.




Figure 2-2: Climb maneuver




Figure 2-3: Constant-drift blunder

Figure 2-3 illustrates the horizontal profile of an example constant-drift blunder scenario. For the constant-drift blunder, T is defined as a quadratic such that

$$T = \frac{-b - \sqrt{b^2 - 4ac}}{2a}, \quad (2.2)$$

where

$$a = v_1^2 + v_2^2 - 2v_1v_2 \cos(\beta)$$

$$b = 2v_2(S \cos(\beta) - D \sin(\beta)) - 2Sv_1$$

$$c = D^2 + S^2 - h_{nmac}^2.$$




Figure 2-4: Constant-turn blunder

Figure 2-4 illustrates the horizontal profile of an example constant-turn blunder scenario. T cannot be calculated directly for the constant-turn blunder and requires a series of intermediate steps. The intruder's turn radius and velocity help calculate when the intruder enters and exits the horizontal region defined by a perpendicular distance of h_{nmac} from either side of the own aircraft's flight path. The boundaries

of this region represent the outer bounds for the possible values of T . Discretizing and searching over this time range yields the time when the intruder is first within an absolute horizontal distance of h_{nmac} from the own aircraft.

2.2 Worst-Case Analysis

With the previous aircraft dynamics and blunders developed, an analysis of the physical collision avoidance limits can be completed. For a constant-drift blunder, the minimum (worst-case) available time to alert (t_c) is calculated for a given runway separation and intruder drift angle β , for 10–90° in 5° intervals, over all possible combinations of velocity, longitudinal runway stagger, and longitudinal aircraft stagger. The velocity for each aircraft is searched between 110 and 170 kt in 5 kt increments. Longitudinal runway and aircraft stagger are both searched between –4000 and 4000 ft in 200 ft increments. The vertical aircraft stagger is not directly varied, but rather derived from the longitudinal runway and aircraft stagger distances.

Figure 2-5 displays the worst-case available time to alert (t_c) versus drift angle β for a 5 s pilot response delay and runway separations of 2500, 3000, 3500, and 4000 ft. As drift angle β increases, the plot lines become slightly wavy due to the velocity and stagger discretization eliminating potentially worse scenarios from consideration. The smallest drift angle β for which the worst-case configuration precludes NMAC resolution ($t_c < 0$) is 50° at 2500 ft runway separation. The worst-case scenario at 3000 ft runway separation that guarantees an NMAC occurs at a 70° drift angle. However, even at a 90° drift angle, worst-case scenarios at runway separations of 3500 and 4000 ft do not quite reach $t_c = 0$.

Although t_c does not reach zero until $\beta = 50^\circ$ for 2500 ft runway separation, the available time to alert is still limited for smaller, more realistic drift angles. For example, even at $\beta = 25^\circ$, the collision avoidance system must recognize the blunder and decide to alert less than 10 s after the blunder begins. This 10 s threshold is crossed at $\beta = 30^\circ$ for 3000 ft runway separation. Alerting within 10 s is manageable, though not trivial, for a system like ACAS Xo.

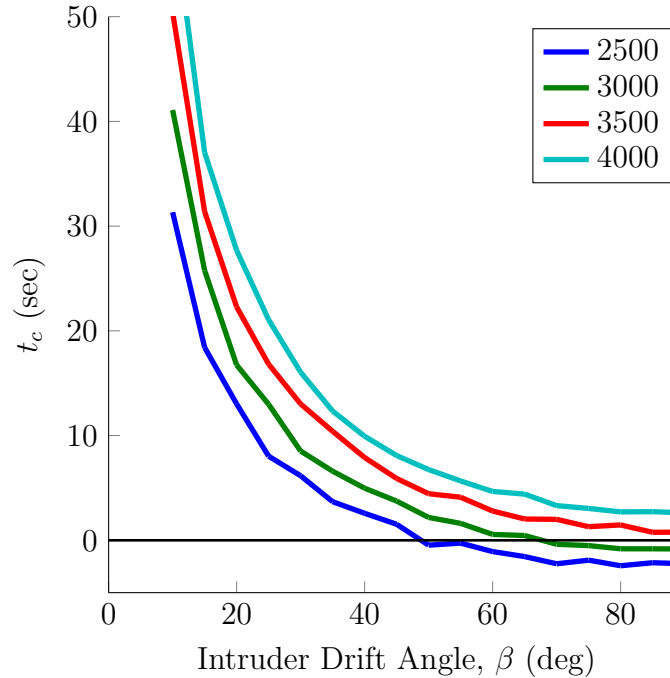


Figure 2-5: Worst-case available time to alert (t_c) vs. drift angle

Figure 2-6 displays the intruder’s perpendicular deviation from the runway at the corresponding t_c from Figure 2-5, with varying pilot response delays, for 3000 ft runway separation. Again, the lack of smoothness in the plot lines is due to the same discretization as in Figure 2-5. Notice that pilot response delay has a larger effect on t_c , and thus the perpendicular distance from the nominal flight path at t_c , as the intruder’s drift angle increases.

Figure 2-6 allows for comparison to a PRM approach, in which the controller would issue breakout instructions after the intruder violates a 2000 ft wide NTZ. This 2000 ft wide NTZ would begin 500 ft from the intruder’s nominal flight path when considering the 3000 ft runway separation of Figure 2-6. Considering the 5 s pilot response delay curve in Figure 2-6, the collision avoidance system would be able to effectively alert later than during PRM for intruder drift angles below 50° (where the curve crosses the 500 ft mark).




Figure 2-6: Drift angle vs. perpendicular deviation

For a constant-turn blunder, the minimum (worst-case) available time to alert (t_c) is calculated for a given runway separation and intruder turn rate ω , between 1 and 5 deg/s in 1 deg/s intervals, over all possible combinations of velocity, longitudinal runway stagger, and longitudinal aircraft stagger. The velocity for each aircraft is searched between 110 and 170 kt in 10 kt increments. Longitudinal runway and aircraft stagger are both searched between -4000 and 4000 ft in 400 ft increments. These larger increments, compared to those of the constant-drift analysis, are due to computational limitations.

Figure 2-7 displays the worst-case available time to alert (t_c) versus turn rate ω for a 5 s pilot response delay and runway separations of 2500, 3000, 3500, and 4000 ft. The filled plot areas represent the t_c range for avoidance maneuver vertical rates between 1500 and 2500 ft/min. For a given runway separation and intruder turn rate, an increased avoidance vertical rate from 1500 ft/min to 2500 ft/min only adds, at most, about 1 s to the available alert time t_c . Since no curve passes below the $t_c = 0$ line, there are no worst-case scenarios which absolutely preclude NMAC resolution. For 3000 ft runway separation, the worst-case available time to alert drops below 10 s

at intruder turn rates greater than or equal to 3 deg/s. Furthermore, each additional 500 ft of runway separation allows for approximately 2 s of additional time to alert.

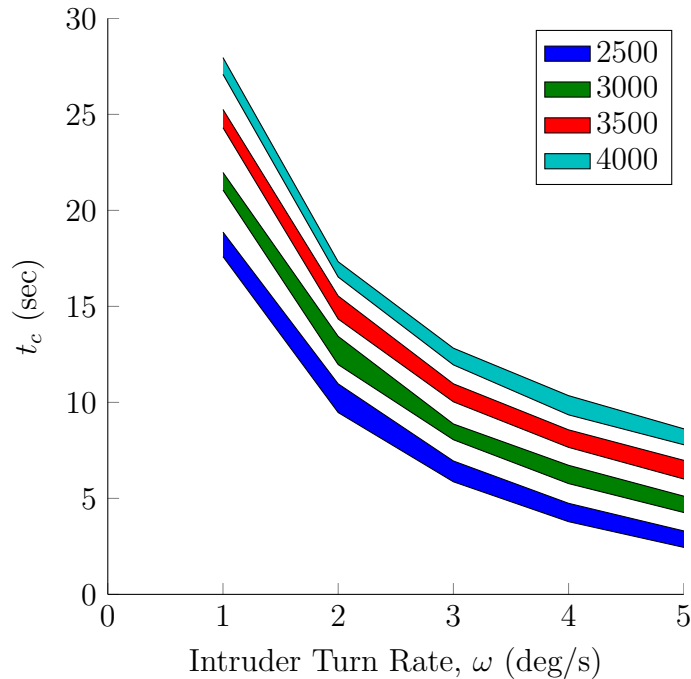


Figure 2-7: Worst-case available time to alert (t_c) vs. turn rate

Figure 2-8 displays the worst-case available time to alert (t_c) versus turn rate ω for 3000 ft runway separation and pilot response delays of 0, 2.5, and 5 s. The trends are similar to those observed in Figure 2-7. As expected, each additional 2.5 s of pilot response delay decreases the available time to alert by 2–3 s.

Figure 2-9 displays the intruder’s perpendicular deviation from the runway at the corresponding t_c from Figure 2-8, with varying pilot response delays, for 3000 ft runway separation and a 2500 ft/min avoidance maneuver. Comparing the 5 s pilot response delay curve in Figure 2-9 with a PRM approach, the collision avoidance system would be able to effectively alert later than when using PRM for intruder turn rates less than 4 deg/s. With a standard 2000 ft NTZ and 3000 ft runway separation, PRM alerts when the intruder deviates 500 ft from centerline. Smaller pilot response delays allow for even greater advantages over PRM alert thresholds, as any intruder turn rate up to at least 5 deg/s does not require an RA until after the 500 ft mark.

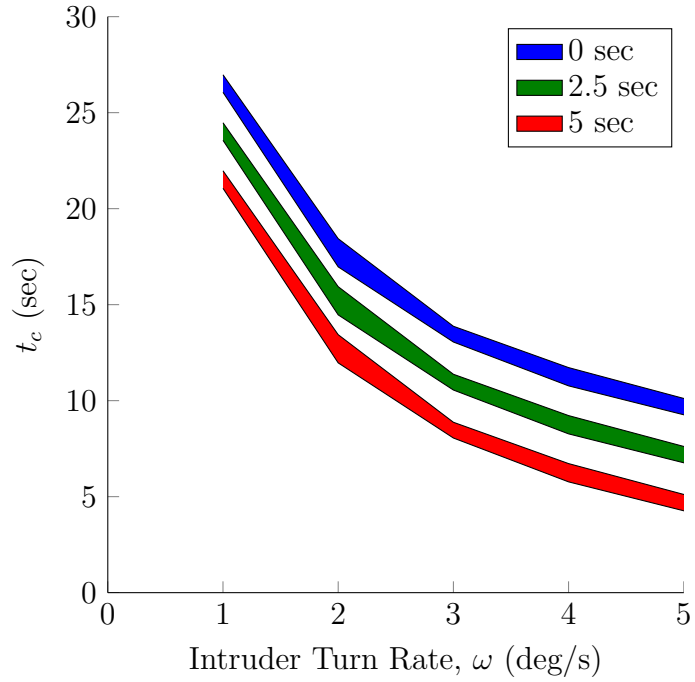


Figure 2-8: Worst-case available time to alert (t_c) vs. turn rate




Figure 2-9: Turn rate vs. perpendicular deviation

2.3 Discussion

According to Figures 2-5 and 2-6, for reasonable blunders (around 30° drift angle or 3 deg/s turn rate) and assuming a 5 s pilot response delay and 3000 ft runway separation, the worst-case scenarios allow 5–10 s to alert after the blunder is initiated. These worst-case alerting requirements correspond to path deviations equal to or greater than those that would trigger a controller-issued alert during PRM operations, as displayed in Figures 2-6 and 2-9. There are no cases in which it is too late to avoid an NMAC ($t_c < 0$) for any turn rate or any drift angle less than 70° at 3000 ft runway separations. However, Figure 2-6 reveals that there are some high-angle drift scenarios at 3000 ft runway separation that require an alert within only a few seconds of the start of a blunder.

Taking these conclusions into consideration, it is feasible that ACAS Xo can be tuned to meet the alerting criteria for all reasonable blunders and potentially more severe constant-turn blunders. Even in the worst-case and with the most conservative pilot response assumption (5 s), ACAS Xo could likely detect the blunder and alert in time to resolve any conflicts. Furthermore, in the case of the more severe drift blunders, even with a few seconds of available alert time, the blunder may be obvious enough from the outset that ACAS Xo is able to alert in time. The true abilities of ACAS Xo, however, cannot be realized until it is tuned for CSPO through surrogate modeling and automated search. This optimization process is described in the next chapter.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Optimization via Surrogate Modeling

To achieve optimal ACAS Xo performance with respect to a set of CSPO encounters, optimization via surrogate modeling is used to tune applicable design parameters. Optimization via surrogate modeling is an approach that aims to reduce computation time and increase optimization effectiveness in a large design space by testing data points that are most likely to lead to the global optimum. The goal of this study is to tune ACAS Xo for optimal performance during CSPO with respect to both safety and alerting behavior. The logic can be tuned by altering the settings of multiple design parameters, the effects and interactions of which are difficult to infer by human judgment.

The automated tuning procedure applied here begins with a screening study of a large initial set of design parameters to reduce the dimensionality of the search to a smaller subset of parameters. Latin hypercubes are optimized to generate a sampling plan that effectively samples the design space [15]. Each design parameter setting in the sampling plan is evaluated on 100,000 CSPO encounters using Monte Carlo simulation. Using the resulting sample data, a surrogate model is generated that estimates the model parameters, and thus the performance function (the “black-box” function) at each point in the design space. The sample data and surrogate model are updated after testing each new infill point (additional sample point) that meets

the Maximum Expected Improvement criterion. With sufficient data, this iterative process converges on the optimum [11]. Figure 3-1 outlines the process.

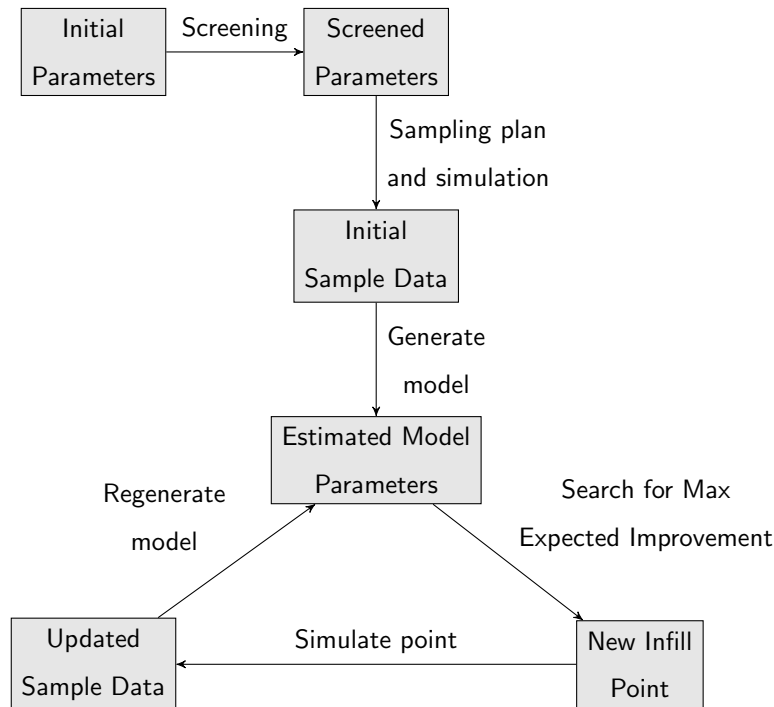


Figure 3-1: Overview of the optimization via surrogate modeling process

3.1 Performance Metrics

As previously stated, the ultimate goal is to tune ACAS Xo for optimal safety and alerting behavior during CSPO. Safety and alerting behavior can generally be measured by two performance metrics:

$$\text{Risk Ratio} = \frac{\text{NMACs with collision avoidance}}{\text{NMACs without collision avoidance}}, \quad (3.1)$$

$$\text{Nuisance Alert Rate} = \frac{\text{Alerts in non-NMAC encounters}}{\text{Non-NMAC encounters}}. \quad (3.2)$$

The output of the Monte Carlo simulations includes information about each encounter, such as whether an NMAC occurs and whether an RA is issued. The CSPO encounters are also simulated without any collision avoidance system to determine

which encounters nominally result in an NMAC. An NMAC is defined to occur when two aircraft come within 500 ft horizontally and 100 ft vertically of each other.

Additional performance metrics not involved in the optimization process may still be of interest when analyzing collision avoidance system performance. These additional metrics, referred to as secondary metrics, relate to specific behaviors during an RA, such as strengthenings and reversals. For example, a Climb (1500 ft/min) command could be strengthened to a 2500 ft/min climb rate or reversed to a Descend command (1500 ft/min). Though strengthenings and reversals are sometimes necessary, they should be minimized when possible due to their disruptive nature [21]. Table 3.1 lists the performance metrics considered in this study. The metrics in boldface correspond to the primary metrics of Equations 3.1 and 3.2 that are used in the tuning process. The performance metrics of Table 3.1 are also commonly used for ACAS Xa performance assessments.

Table 3.1: Optimization Performance Metrics

Metric	Description
Unresolved NMACs	# NMACs that occur when an alert is issued and when an NMAC would have occurred otherwise
Induced NMACs	# NMACs that occur when an alert is issued but when an NMAC would not have occurred otherwise
Missed Alerts	# NMACs that occur when no alert is issued
Risk Ratio	Ratio of total NMACs with collision avoidance to NMACs without collision avoidance
Nuisance Alert Rate	Rate of alerting in encounters where no NMAC occurs both with and without collision avoidance
Strengthenings	# encounters that include a strengthening command
Reversals	# encounters that include a reversal command

3.2 Historical Encounter Set

To support the tuning efforts described in this thesis, an historical CSPO encounter set of 100,000 encounters was generated to enable Monte Carlo simulation. A technical report describes the historical CSPO encounter model development in further detail [41]. Individual encounters are used to simulate certain aircraft dynamics, con-

figurations, and procedures by specifying the states of multiple aircraft at regular time intervals. This encounter set includes significant aircraft blunders and is used to assess system safety as well as alerting behavior.

For every encounter, runways are separated by 3000 ft without stagger. A runway separation of 3000 ft is selected because it represents the minimum distance allowed for parallel operations without addressing the risk of wake turbulence. For each encounter, two aircraft are initialized so that they come within 2000 ft of each other's longitudinal position at some point along the approach path.

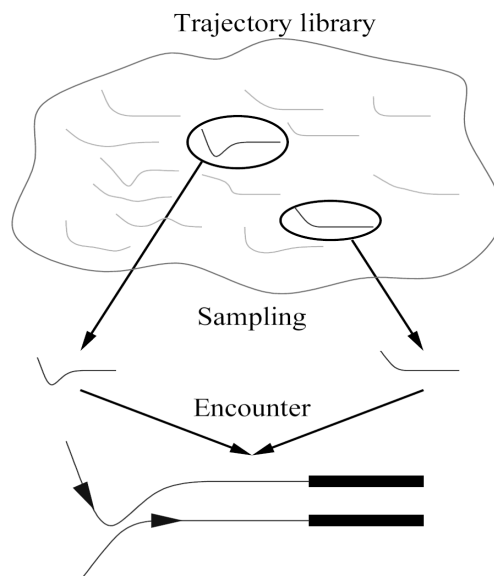


Figure 3-2: Example encounter generation

The 100,000 encounters are generated by randomly selecting two aircraft tracks from an historical trajectory library. The library contains approximately 140,000 tracks of aircraft on approach to landing during IMC at various airports in the NAS. Three main assumptions were required to create the historical trajectory library [41]:

- Only including trajectories from IMC periods results in only instrument approach trajectories.
- Lateral deviations about the localizer are symmetrically distributed.
- An approach path's glide slope is best determined by the mode of the discretized probability distribution.

Each aircraft track is transformed (rotated and/or flipped) such that the aircraft approaches a generic runway 09L or 09R (east-bound). To create an encounter, two trajectories are transformed to approach the desired parallel runways. An illustration of this process is provided in Figure 3-2.

Each sampled trajectory begins at least 10 NM from the runway threshold. The sampling process biases trajectories with large deviations over small deviations by up to a factor of 1000. Nominally, each pair of trajectories should be vertically separated by at least 1000 ft when entering the localizer beam, which guides each aircraft’s lateral position. Blunder types include incorrect altitude join-ups, incorrect ILS captures (Figure 3-3(a)), and break-out maneuvers in the direction of the paired aircraft (Figure 3-3(b)), each occurring at a rate of one per hundred approaches. This rate is chosen so that ACAS Xo is tested against a sufficiently high number of severe blunders.




Figure 3-3: Example aircraft behavior

Because the aircraft trajectories are sampled directly from the historical library, they exhibit a wide range of behavior, including intercepting the localizer at large angles (Figure 3-3(c)) and capturing the ILS at a variety of altitudes and distances from the runway. Some of these behaviors are outside those observed during actual CSPO operations in IMC. The goal of such an encounter set is to allow for simulations over a wide range of possible scenarios. This is the encounter set that is used to tune ACAS Xo for CSPO.

3.3 Simulation and Evaluation

Collision avoidance system performance is evaluated on a given encounter using Monte Carlo simulation. Both TCAS and ACAS Xo aircraft are simulated to be equipped with the same surveillance capabilities to determine the range, bearing, and altitude of other aircraft within radar range. The range and bearing measurements are corrupted with zero-mean Gaussian noise with 50 ft and 10° standard deviations, respectively, to account for surveillance uncertainty, while the altitude is quantized to 25 ft.

The state space within which the sensor measurements are taken must be discretized so that the ACAS X solution is computationally feasible to obtain. Variables within the horizontal and vertical dimensions are discretized such that the state space is adequately defined. A pilot response model must also be defined to accurately reflect pilot response delays observed in real life. For both TCAS and ACAS X simulations, the pilot response model includes a 5 s initial pilot response delay and a 3 s subsequent pilot response delay. In other words, the pilot will begin to respond 5 s after an initial RA is issued and 3 s after any additional RAs are issued.

After an encounter is simulated, a wide variety of performance metrics can be extracted for that encounter. For the purposes of tuning ACAS Xo, these metrics include:

- Whether an RA is issued for either aircraft
- Whether a strengthening RA is issued for either aircraft
- Whether a reversal RA is issued for either aircraft
- Whether the encounter results in an NMAC

A set of encounters can be simulated with any combination of unequipped aircraft (i.e., lacking a collision avoidance system) and aircraft equipped with TCAS or ACAS Xo. By simulating an encounter set with two unequipped aircraft, in addition to simulating with equipped aircraft, the performance metrics for each case can be combined to output additional meaningful metrics for each encounter, such

as specific NMAC types. For example, an NMAC resulting from a simulation of equipped aircraft is categorized as an induced NMAC if the unequipped simulation of the same encounter does not result in an NMAC. Unresolved NMACs and missed alerts can be similarly extracted from the raw output metrics.

3.4 Defining the Objective Value

The collision avoidance system performance for a given set of encounters is essentially a “black-box” function since the performance cannot be known until simulations are performed for each combination of design parameter settings. An underlying premise of surrogate modeling is that the performance of the system can be defined by a single function value, or objective value, at each point in the design space. Based on the objective values at different parameter settings, a surrogate model for the function can be estimated with increasing accuracy as more data points are intelligently selected and tested over time to find the “black-box” function’s global optimum.

Utility elicitation is one concept that helps generate an objective function which combines performance metrics into a single objective value [38]. Examples of utility elicitation methods include additive modeling, swing weight elicitation, and rank-weighting [8]. Although several methods are available to elicit preferences from users to determine the appropriate attribute rankings and relative weightings, utility elicitation is not within the scope of this study.

The objective function is formulated as a linear combination of two performance metrics: risk ratio and nuisance alert rate. A weighting sweep is performed to test various linear combinations of performance metrics and select the weighting that most effectively tunes ACAS Xo. The weighting sweep procedure and results are provided in more detail in Appendix A. Since safety is the primary concern, the risk ratio is given more weight than the nuisance alert rate, resulting in the following objective function:

$$\text{Objective Value} = 0.95 \times (\text{Risk Ratio}) + 0.05 \times (\text{Nuisance Alert Rate}). \quad (3.3)$$

Risk ratio also requires more weight due to its lower-magnitude values compared to those of nuisance alert rate. The selected linear combination in Equation 3.3 is cross-referenced and confirmed with a subjective hand-ranking of data points with respect to overall performance. Such a simple linear combination of performance metrics has been shown to perform well in decision-making processes [5].

3.5 Screening Study

To maximize computational efficiency in the optimization process, the scope of the tuning process is reduced. For the ACAS Xo problem, the scope correlates to the number of design variables that will be tuned. The ACAS X logic incorporates 46 distinct design parameters. Based on prior knowledge of ACAS X, the eight parameters listed in Table 3.2 were deemed appropriate to tune and potentially relevant to ACAS Xo behavior in CSPO encounters. These parameters are discussed in further detail in technical reports [1, 21]. The goal of the screening study is to extract the most significant parameters affecting the performance metrics of interest, specifically the primary metrics of risk ratio and nuisance alert rate.

Table 3.2: Design Parameters Analyzed in the Screening Study

Parameter	Description
r alert	Cost of issuing an RA to the pilot
r maintain < 1500 ft/min	Cost of issuing a “Maintain Vertical Rate” command when the current vertical rate is < 1500 ft/min
ddx ddy sigma	Intruder horizontal acceleration deviation
cycles	Number of seconds for which alerting is allowed when in horizontal conflict
horiz. conflict def.	Horizontal range that defines an NMAC for logic calculations
vert. conflict def.	Vertical range that defines an NMAC for logic calculations
slow clos. rho thresh.	Range threshold where time until conflict is calculated using horizontal and vertical information
slow clos. drho thresh.	Range rate threshold where time until conflict is calculated using horizontal and vertical information

To determine the most significant parameters relative to each performance metric, the elementary effect on the primary and secondary performance metrics is calculated;

that is, a sensitivity study is performed. The elementary effect for each parameter-metric pair is expressed by a mean and standard deviation. The mean value measures the change in the performance metric for a given change in the parameter. A large-magnitude mean indicates that the parameter has a strong effect on the given metric. A large standard deviation indicates that the parameter experiences strong interactions with other parameters [11].

Table 3.3: Design Parameter Screening Study Results (Mean)

Design Parameter	Min	Max	Risk	Alert	Strengthening	Reversal
			Mean	Mean	Mean	Mean
r alert	$-1.5 \cdot 10^{-2}$	$-3 \cdot 10^{-3}$	-0.0643	0.3047	-0.0198	-0.0002
r maintain < 1500 ft/min	-1	$-1 \cdot 10^{-2}$	0.0000	0.0000	-0.0002	0.0000
ddx ddy sigma	0	32	-0.0012	0.0411	-0.0132	-0.0006
cycles	0	5	-0.0684	0.4368	0.0337	0.0020
horiz. conflict def.	300	1,000	-0.0358	0.2039	-0.0087	-0.0002
vert. conflict def.	0	500	-0.0239	0.5672	0.0461	0.0008
slow clos. rho thresh.	1,500	3,000	0.0372	0.0742	0.0140	0.0000
slow clos. drho thresh.	0	500	0.0191	0.0131	0.0008	0.0000

Table 3.4: Design Parameter Screening Study Results (Standard Deviation)

Design Parameter	Min	Max	Risk	Alert	Strengthening	Reversal
			Std Dev	Std Dev	Std Dev	Std Dev
r alert	$-1.5 \cdot 10^{-2}$	$-3 \cdot 10^{-3}$	0.0556	0.1963	0.0174	0.0004
r maintain < 1500 ft/min	-1	$-1 \cdot 10^{-2}$	0.0000	0.0000	0.0002	0.0000
ddx ddy sigma	0	32	0.0009	0.0409	0.0127	0.0009
cycles	0	5	0.0820	0.2817	0.0349	0.0019
horiz. conflict def.	300	1,000	0.0378	0.1434	0.0262	0.0010
vert. conflict def.	0	500	0.0349	0.3387	0.0356	0.0009
slow clos. rho thresh.	1,500	3,000	0.0540	0.1082	0.0270	0.0000
slow clos. drho thresh.	0	500	0.0314	0.0251	0.0017	0.0002

Table 3.3 shows the eight parameter ranges and the mean elementary effect for each parameter-metric pair. As an example, the cycles parameter, which was tested between values of 0 and 5, exhibits a relatively large-magnitude elementary effect for each of the four metrics. Therefore, the cycles parameter has a significant effect on each of the performance metrics and is included in the optimization process moving

forward. The standard deviations are also included in Table 3.4. The results indicate a strong correlation between design parameters with a large standard deviation and those with a large mean.

The five parameters that exhibit significant elementary effects for the primary metrics are: (1) cost of alerting, (2) number of cycles, (3) horizontal conflict definition, (4) vertical conflict definition, and (5) slow closure rho threshold. The cost of alerting relates to the cost of issuing an initial RA to the pilot. The cycles parameter indicates the number of seconds the logic may continue issuing an RA while the two aircraft remain in horizontal conflict for an extended period of time. The cycles parameter is particularly relevant when considering CSPO encounters in which low closure rates are prevalent. The horizontal and vertical conflict definitions define the NMAC region on which the ACAS X logic bases its calculations, though the actual horizontal and vertical NMAC definitions remain unchanged (500 ft and 100 ft, respectively) for the performance metrics. The slow closure rho threshold parameter defines the range threshold between aircraft where ACAS X begins using both horizontal and vertical state information to calculate the estimated time until conflict. These five design parameters are used for the remainder of the tuning of ACAS Xo for CSPO.

3.6 Sampling Plan

The next step to prepare for the optimization is to generate a sampling plan that will provide a uniform, space-filling initial data set over the design space. This initial sampling plan will lead to a well informed surrogate model from which additional infill points can be intelligently selected for testing. The sampling plan is generated using Latin hypercubes that are optimized with respect to the Morris-Mitchell criteria. The result is a sampling plan that does not repeat any one parameter setting and comprehensively samples the design space [11]. Forty sample points, each representing a combination of parameter settings, are tested on the historical encounter set of 100,000 CSPO encounters. Table 3.5 lists the ranges used for the five design parameters.

Table 3.5: Design Parameter Ranges used for Optimization

Design Parameter	Min	Max
r alert	$-5 \cdot 10^{-3}$	$-2.5 \cdot 10^{-3}$
cycles	4	7
horiz. conflict def.	200	400
vert. conflict def.	0	100
slow clos. rho thresh.	1,500	2,500

3.7 Surrogate Modeling

After simulating the 100,000 CSPO encounters from the historical encounter set, and calculating objective values using Equation 3.1 through Equation 3.3 for each parameter setting from the sampling plan, a surrogate model can be generated that approximates the objective value of Equation 3.3 at every point in the design space.

The Kriging model is a Gaussian process based model that allows for the calculation of a prediction mean and variance at every point in the design space. The Kriging model is chosen for its ability to model complex functions and measure uncertainty at each point [22]. Specifically, the Kriging model is defined by the radial basis function of the form

$$\psi^{(i)} = \exp \left(- \sum_{j=1}^k \theta_j |x_j^{(i)} - x_j|^{p_j} \right). \quad (3.4)$$

The Kriging basis function only differs from a Gaussian basis function in that the Kriging basis allows the width and smoothness of the function to vary in each dimension. Holding the width and smoothness constant for every parameter yields the Gaussian basis function [11].

In Equation 3.4, ψ denotes the basis function (one centered at each data point), k is the number of design parameters (dimensionality), θ defines the basis function width (also known as the “activity” parameter) for each dimension, and p determines the function smoothness for each dimension. The sample data and their observed

responses (objective values) are defined, respectively, as follows:

$$\mathbf{X} = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)}\}^T, \quad (3.5)$$

$$\mathbf{y} = \{y^{(1)}, y^{(2)}, \dots, y^{(n)}\}^T. \quad (3.6)$$

The observed responses at each data point are considered to be derived from a stochastic process with mean $\mathbf{1}\mu$, where $\mathbf{1}$ is an n by 1 vector of ones:

$$\mathbf{Y} = \{Y(\mathbf{x}^{(1)}), \dots, Y(\mathbf{x}^{(n)})\}^T. \quad (3.7)$$

The basis function of Equation 3.4 can then be used to calculate the correlation between each pair of responses, resulting in an n by n correlation matrix Ψ , where an element $\Psi(i, l)$ of the matrix is defined by

$$\text{cor}[Y(\mathbf{x}^{(i)}), \dots, Y(\mathbf{x}^{(l)})] = \exp\left(-\sum_{j=1}^k \theta_j |x_j^{(i)} - x_j^{(l)}|^{p_j}\right). \quad (3.8)$$

The correlation between a pair of responses, expressed by Equation 3.8, depends on the distance between each point as well as the θ and p parameters [11].

The next step in generating the surrogate model is to calculate maximum likelihood estimates (MLEs) for the parameters μ and σ , which are base terms used to define the model. The MLEs of μ and σ are calculated by maximizing the likelihood of the existing sample data given the parameters:

$$L(\mathbf{Y}^{(i)}, \dots, \mathbf{Y}^{(n)} | \mu, \sigma) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left[-\frac{\sum_{i=1}^n (\mathbf{Y}^{(i)} - \mu)^2}{2\sigma^2}\right]. \quad (3.9)$$

Equation 3.9 can be expressed in terms of its natural logarithm and the existing data set:

$$\ln(L) = \frac{n}{2} \ln(2\pi) - \frac{n}{2} \ln(\sigma^2) - \frac{1}{2} \ln |\Psi| - \frac{(\mathbf{y} - \mathbf{1}\mu)^T \Psi^{-1} (\mathbf{y} - \mathbf{1}\mu)}{2\sigma^2}. \quad (3.10)$$

By setting the derivatives of Equation 3.10 to zero, the MLEs for μ and σ^2 can be extracted [11]:

$$\hat{\mu} = \frac{\mathbf{1}^T \Psi^{-1} \mathbf{y}}{\mathbf{1}^T \Psi^{-1} \mathbf{1}}, \quad (3.11)$$

$$\hat{\sigma}^2 = \frac{(\mathbf{y} - \mathbf{1}\hat{\mu})^T \Psi^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu})}{n}. \quad (3.12)$$

These MLEs can be substituted back into Equation 3.10 to derive the concentrated log-likelihood function (constant terms removed) which depends on θ and p via Ψ [11]:

$$\ln(L) \approx -\frac{n}{2} \ln(\hat{\sigma}^2) - \frac{1}{2} \ln |\Psi|. \quad (3.13)$$

The log-likelihood function in Equation 3.13 is not differentiable and requires numerical optimization to maximize the function and find MLEs for θ and p .

A genetic algorithm was applied throughout this process [12, 42]. The parameter θ was searched on a logarithmic scale from 10^{-3} to 10^3 , and p was searched from 1 to 2. After finding the MLEs for each model parameter, a prediction can be made at any point in the design space by maximizing the conditional likelihood of the existing data set and the point's predicted objective value. After defining a vector $\boldsymbol{\psi}$ of correlations between the prediction point and each existing observed value, the objective value prediction at a point \mathbf{x} in the design space can be calculated [11] by

$$\hat{y}(\mathbf{x}) = \hat{\mu} + \boldsymbol{\psi}^T \Psi^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu}). \quad (3.14)$$

3.8 Searching the Model

After developing the Kriging model to estimate the objective value at each point in the design space, a search algorithm must determine which point in the design space, known as an infill point, to test next and efficiently drive toward the optimum design setting with the best performance. This determination relies on a specific heuristic to be used in computing the expected reward of sampling a certain point. In searching for the global optimum, an heuristic can be designed to focus on global exploration,

local exploitation, or a combination of both. A possible exploration heuristic could aim to maximize the variance (uncertainty) of the function, so that information gain is likewise maximized in each iteration [39]. This search strategy ignores potential points of improvement (greedy search) in favor of increased learning of the function. Entropy Search is one such method that places an emphasis on exploration as opposed to exploitation [13].

In contrast to exploration, an exploitation-based heuristic places emphasis on finding local optima earlier in the process. One such heuristic could combine predicted mean and variance to calculate an optimistic predicted objective value (reward) at each point [39]. A final option is combining global exploration and local exploitation strategies into a balanced heuristic. One way of striking this balance is by searching for the point of Maximum Expected Improvement (Max EI) given each prediction's mean $\hat{y}(\mathbf{x})$ and variance $\hat{s}^2(\mathbf{x})$ [35]. After comparing results through preliminary optimization tests, the Max EI search method yielded the best results in tuning ACAS Xo. The expected improvement at a point \mathbf{x} is calculated [11] by

$$E[I(\mathbf{x})] = (y_{min} - \hat{y}(\mathbf{x})) \left[\frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{y_{min} - \hat{y}(\mathbf{x})}{\hat{s}(\mathbf{x})\sqrt{2}} \right) \right] + \hat{s}(\mathbf{x}) \frac{1}{\sqrt{2\pi}} \exp \left[\frac{-(y_{min} - \hat{y}(\mathbf{x}))^2}{2\hat{s}^2(\mathbf{x})} \right]. \quad (3.15)$$

After sampling each subsequent infill point, the Kriging model in Equation 3.4 is recomputed. The process is repeated until eighty infill points are computed and evaluated for a total of 120 evaluations. Figure 3-4 shows the objective value achieved after including each new infill point (solid). The dashed curve indicates the best value achieved up to and including that infill point. The 103rd data point (i.e., the 63rd infill point) results in the best performance with a risk ratio of 0.0014 and a nuisance alert rate of 12.9% for the historical encounter set. As mentioned in Section 3.4, this infill procedure is repeated using different weightings in the objective value calculation. The weighting sweep is performed in Appendix A, indicating that the 95%/5% weighting of the primary metrics (risk ratio and nuisance alert rate, respectively) used here results in the most effective tuning.




Figure 3-4: Objective values of eighty infill points (solid) and the cumulative minimum value achieved (dashed)

Tables 3.6 and 3.7 show the eighty infill points and their simulation results. The progression in Figure 3-4 indicates significant local search but also includes global exploration which is partially indicated by spikes in the objective values. The local search is often indicated by consecutive points with similar, ideally decreasing, values.

However, local search does not always yield improving results, especially during early infill points as the model is not yet well predicted in various areas of the design space. As an example, data points 55–58 represent a local search progression. Each pair of consecutive data points in this cluster contains similar values for two of the parameters, as shown in Table 3.6. The remaining three parameters are altered, based on the current Kriging model, in hopes of improving the objective value. The objective value improves with the 56th and 57th data points, but the 58th point allows too many alerts and increases the objective value. Notice that, following the 58th data point, the point of Max EI is no longer local, and global search is implemented for the 59th data point. In this case, the global search at the 59th data point yields an objective value similar to that of the 57th data point, albeit in a different area of the design space.

Table 3.6: Evaluation Results for the First Forty Infill Points

Data Pt	R Alert	Cyc.	HCD	VCD	Slow Clos. Rho	Risk Ratio	Nuis. Rate	Obj Value
41	-0.0050	7	293	0	1,562	0.2946	0.092	0.2845
42	-0.0025	7	224	5	1,653	0.0021	0.147	0.0094
43	-0.0026	5	312	91	1,500	0.0021	0.155	0.0098
44	-0.0050	7	215	0	2,012	0.4584	0.076	0.4392
45	-0.0037	4	400	5	2,281	0.0014	0.163	0.0095
46	-0.0025	4	400	50	1,500	0.0014	0.160	0.0093
47	-0.0025	4	200	39	2,500	0.0007	0.214	0.0113
48	-0.0025	4	200	91	2,500	0.0007	0.214	0.0113
49	-0.0025	4	308	44	2,500	0.0007	0.237	0.0125
50	-0.0050	4	200	57	2,500	0.0042	0.179	0.0129
51	-0.0050	7	400	71	1,500	0.0021	0.130	0.0085
52	-0.0025	7	395	20	2,500	0.0007	0.280	0.0146
53	-0.0050	5	262	11	1,769	0.0049	0.104	0.0098
54	-0.0050	4	300	0	2,314	0.5388	0.063	0.5150
55	-0.0027	4	400	6	2,188	0.0014	0.178	0.0102
56	-0.0031	4	400	3	1,500	0.0028	0.141	0.0097
57	-0.0025	4	200	83	1,500	0.0042	0.108	0.0094
58	-0.0025	7	400	12	1,500	0.0014	0.196	0.0111
59	-0.0038	4	300	7	2,000	0.0035	0.121	0.0094
60	-0.0050	4	200	85	1,500	0.0147	0.074	0.0177
61	-0.0025	4	400	4	2,063	0.0014	0.175	0.0101
62	-0.0025	4	201	66	2,500	0.0007	0.214	0.0113
63	-0.0037	7	200	5	1,816	0.0028	0.114	0.0084
64	-0.0025	4	200	28	2,500	0.0007	0.214	0.0113
65	-0.0050	7	200	52	1,500	0.0126	0.089	0.0164
66	-0.0038	4	230	99	1,750	0.0049	0.101	0.0097
67	-0.0050	7	322	4	2,000	0.0014	0.133	0.0080
68	-0.0050	7	324	5	1,926	0.0014	0.130	0.0078
69	-0.0025	7	400	29	1,500	0.0014	0.196	0.0111
70	-0.0025	7	400	97	1,500	0.0014	0.196	0.0111
71	-0.0050	4	398	100	1,577	0.0056	0.107	0.0107
72	-0.0050	7	299	0	2,095	0.4486	0.086	0.4304
73	-0.0050	4	400	8	2,500	0.0007	0.196	0.0105
74	-0.0050	7	400	59	2,500	0.0007	0.224	0.0119
75	-0.0050	4	400	27	2,078	0.0028	0.125	0.0089
76	-0.0025	4	400	72	2,500	0.0007	0.246	0.0130
77	-0.0025	7	200	65	1,500	0.0028	0.129	0.0091
78	-0.0050	7	200	61	1,500	0.0126	0.089	0.0164
79	-0.0042	4	309	57	2,250	0.0014	0.143	0.0085
80	-0.0034	4	312	100	2,225	0.0007	0.153	0.0083

The improving accuracy of the Kriging model over time is apparent, indicated in Figure 3-4 by consistently low objective values following data point 97. Careful analysis of Table 3.7 shows that the majority of consecutive data points after data point 97 do not share similar parameter values, indicating a search progression different from the local search experienced for data points 55–58. However, global search is also not prevalent, with the exception of data point 117, considering the relative consistency in the objective values achieved. A comparison of Tables 3.6 and 3.7 reveals that many of the parameter settings after data point 97 are similar to earlier infill points. This indicates that the search progression after data point 97 mostly contains isolated continuations of local search progressions from the first forty infill points.

Table 3.7: Evaluation Results for the Second Forty Infill Points

Data Pt	R Alert	Cyc.	HCD	VCD	Slow Clos. Rho	Risk Ratio	Nuis. Rate	Obj Value
81	-0.0050	4	217	100	1,633	0.0119	0.085	0.0155
82	-0.0050	4	300	0	1,875	0.4612	0.067	0.4415
83	-0.0025	5	400	4	1,718	0.0014	0.179	0.0103
84	-0.0025	4	400	41	2,000	0.0014	0.172	0.0099
85	-0.0041	5	311	25	1,984	0.0014	0.133	0.0080
86	-0.0036	5	309	1	1,716	0.0028	0.132	0.0092
87	-0.0049	4	288	100	2,062	0.0063	0.111	0.0115
88	-0.0047	7	319	27	1,959	0.0014	0.135	0.0081
89	-0.0035	5	375	44	2,375	0.0007	0.196	0.0104
90	-0.0050	4	300	0	2,188	0.5227	0.064	0.4998
91	-0.0046	5	338	10	1,754	0.0042	0.117	0.0098
92	-0.0038	4	334	28	1,990	0.0021	0.128	0.0084
93	-0.0041	6	304	0	2,091	0.4010	0.105	0.3862
94	-0.0044	7	335	74	1,813	0.0014	0.135	0.0081
95	-0.0038	6	400	48	1,612	0.0021	0.149	0.0095
96	-0.0034	4	262	2	1,500	0.0049	0.110	0.0101
97	-0.0025	4	267	1	1,563	0.2449	0.116	0.2385
98	-0.0038	5	400	3	1,500	0.0028	0.136	0.0095
99	-0.0048	4	284	100	2,094	0.0056	0.114	0.0110
100	-0.0050	4	394	100	2,156	0.0042	0.131	0.0105
101	-0.0025	4	335	100	1,783	0.0021	0.152	0.0096
102	-0.0036	6	231	66	2,187	0.0014	0.145	0.0086
103	-0.0038	5	278	90	2,000	0.0014	0.129	0.0078
104	-0.0042	6	237	43	2,164	0.0014	0.134	0.0080
105	-0.0025	4	393	100	1,502	0.0014	0.160	0.0093
106	-0.0031	4	282	48	2,334	0.0007	0.166	0.0090
107	-0.0050	6	397	45	1,500	0.0028	0.124	0.0088
108	-0.0050	6	252	99	1,562	0.0056	0.105	0.0105
109	-0.0035	4	318	100	2,250	0.0014	0.154	0.0090
110	-0.0040	5	334	24	1,828	0.0021	0.128	0.0084
111	-0.0044	7	345	55	2,102	0.0014	0.148	0.0087
112	-0.0029	5	300	69	1,625	0.0035	0.135	0.0101
113	-0.0037	7	275	100	1,514	0.0021	0.127	0.0084
114	-0.0041	4	400	72	1,676	0.0035	0.123	0.0095
115	-0.0041	4	288	100	2,078	0.0035	0.121	0.0094
116	-0.0034	6	371	39	2,250	0.0007	0.187	0.0100
117	-0.0050	5	300	1	2,043	0.4675	0.074	0.4478
118	-0.0037	4	200	42	1,500	0.0126	0.086	0.0163
119	-0.0025	7	233	7	1,521	0.0021	0.145	0.0092
120	-0.0037	5	200	81	1,500	0.0084	0.094	0.0127

3.9 Discussion

The procedure executed in this chapter attempted to locate the global optimum of a performance function related to the behavior of ACAS Xo during CSPO. An historical blunder set was utilized to compute the performance of ACAS Xo by testing the system on a wide variety of blunders and aircraft behaviors. A small set of design parameters was selected to tune such that the behavior of ACAS Xo during CSPO would be meaningfully affected. A Kriging model was fit to the sample data and subsequently searched to determine the point of Max EI for each infill iteration. After eighty iterations of the search and infill procedure, the 63rd infill point was

discovered to achieve the lowest objective value and, theoretically, the most desired performance of the existing sample data. However, a thorough investigation of the performance and behavior of ACAS Xo at this data point must be carried out. This investigation includes performance analyses on several distinct encounter sets and comparisons to the performance of TCAS, the current benchmark. This performance analysis and comparison is the objective of the next chapter.

Chapter 4

Performance Analysis

The tuning process of Chapter 3 results in an optimal combination of the five design parameters that emerged from the screening study: (1) cost of alerting, (2) number of cycles, (3) horizontal conflict definition, (4) vertical conflict definition, and (5) slow closure rho threshold. The optimal settings for ACAS Xo are found at the 103rd data point in Table 3.7. The cost of alerting defines the cost of issuing an RA to the pilot. The number of cycles indicates the number of seconds the logic may continue issuing an RA while the two aircraft are in horizontal conflict for an extended period of time. In normal circumstances, ACAS X ceases alerting once the aircraft are in horizontal conflict, but the cycles parameter allows for continued alerting during this time, which becomes necessary during many slow-closure scenarios such as CSPO. The horizontal and vertical conflict definitions define the NMAC region for the ACAS X logic, though the NMAC definition for subsequent performance analysis remains fixed at 500 ft horizontally and 100 ft vertically. The slow closure rho threshold parameter defines the aircraft range at which ACAS X calculates the estimated time until conflict using both horizontal and vertical state information, as opposed to solely horizontal state information.

Table 4.1 shows the relative parameter value changes from ACAS Xa to ACAS Xo. The cost of alerting is increased by 50%, which is expected since one of the main objectives is to limit the rate of nuisance alerts. However, the alerting cost is not increased to the point of adversely affecting safety and allowing additional NMACs to remain

unresolved, as is evident in this chapter’s performance analyses. To effectively reduce nuisance alerts, alerting must also be restricted to a select set of circumstances defined by the horizontal and vertical conflict definitions. These two parameters define the NMAC thresholds that the ACAS X logic attempts to avoid. If the conflict definition values are too large, the logic will alert more than necessary so that it maintains the required separation defined by the horizontal and vertical conflict definitions. Both of the conflict definition values decrease, especially the horizontal parameter, to reflect the smaller ranges experienced during CSPO. These decreases help accommodate CSPO with smaller runway separations and turns near parallel approach paths, while still enforcing adequate separation to maintain acceptable safety levels. The number of cycles increases to reflect the need to account for more slow-closure scenarios in which alerts may still be effective when two aircraft are converging at close range. Finally, the slow closure rho threshold parameter decreases to activate the specialized slow-closure logic earlier in the encounter.

Table 4.1: ACAS Xo Parameter Value Changes Relative to ACAS Xa

Parameter	Relative Change
r alert	50%
cycles	67%
horiz. conflict def.	−72%
vert. conflict def.	−10%
slow clos. rho thresh.	−33%

In this chapter, the optimized parameter settings for ACAS Xo are shown to balance safety and operational suitability. Although system experts could likely deduce the direction of change required for each of the five tuned parameters, the appropriate magnitude of change for each parameter is much more difficult to determine. The ability to solve such a complex problem is the strength of the automated tuning process carried out in the previous chapter. The surrogate model is able to infer the performance effects of multidimensional parameter changes and identify promising areas of the design space.

The remainder of this chapter compares the performance of ACAS Xo and TCAS. Section 4.1 compares performance on the historical encounter set described in Chapter 3 and provides two sample encounters. ACAS Xo is tuned for this encounter set and is expected to exhibit its best behavior for this encounter set. Section 4.2 investigates each parameter’s contribution to the improvement of ACAS Xo performance over that of ACAS Xa, according to the changes in Table 4.1. Section 4.3 describes two additional generated encounter sets that more accurately reflect typical aircraft behavior during CSPO, and uses them to compare the performance of ACAS Xo and TCAS. The performance of ACAS Xo for these two encounter sets helps demonstrate that the tuned logic performs well on a wide variety of encounter types and distributions.

Section 4.4 compares ACAS Xo and TCAS performance for an operational performance model. The operational performance model tests each system’s ability to minimize nuisance alerts in non-NMAC encounters with varying degrees of aircraft navigational error. Section 4.5 compares the specific alerting behaviors of ACAS Xo and TCAS at various horizontal and vertical configurations. The comparisons allow for a deeper understanding of each system’s sensitivity to changes in both the horizontal and vertical dimensions. Section 4.6 applies the ACAS Xo logic to four sample encounters derived from the worst-case analysis of Chapter 2. The actual alert times of ACAS Xo are compared to the theoretical required alert times calculated in Chapter 2. Finally, Section 4.7 discusses the chapter’s results and their implications.

4.1 Historical Encounter Set Performance

The optimized parameter settings of ACAS Xo exhibit improved performance over TCAS. Table 4.2 compares the performance of both systems for the historical blunder encounter set described in Chapter 3.

Table 4.2: Metric Results for the Historical Blunder Encounter Set

System	NMACs			Risk Ratio	Nuisance Rate	Strengthenings	Reversals
	Unresolved	Induced	Missed Alerts				
ACAS Xo	0	0	2	0.0014	0.129	596	236
TCAS	4	1	2	0.0049	0.358	1,562	2,283

As indicated in Table 4.2, TCAS results in mediocre performance compared to ACAS Xo. TCAS is unable to resolve four of the 1429 NMACs, misses two alerts, and induces one NMAC which would not have occurred had the aircraft continued their trajectories without interruption. The two missed alerts with both TCAS and ACAS Xo are for encounters that experience extremely low closure rates at low altitudes where alerting is inhibited due to operational concerns. These encounters are considered extraneous and are arguably encounters for which an alert above the inhibition altitude of approximately 1000 ft would be undesirable, due to sufficient aircraft separation down to 1000 ft of altitude.

The most noticeable advantage of ACAS Xo over TCAS is its lower nuisance alert rate. ACAS Xo achieves a rate of 12.9% compared to 35.8% with TCAS. Furthermore, ACAS Xo limits the number of encounters that include strengthenings or reversals. ACAS Xo only issues strengthenings and reversals in 596 and 236 encounters, respectively, whereas TCAS issues them in 1562 and 2283 encounters, respectively. Although they may be necessary in some situations, strengthenings and reversals should be minimized when possible [21]. ACAS Xo unquestionably outperforms TCAS in all areas under consideration.


To help illustrate the differences between ACAS Xo and TCAS, two sample encounters are provided in Figure 4-1 and Figure 4-2. Aircraft 1 is defined by the darker track and Aircraft 2 is defined by the lighter track. Figure 4-1 illustrates an encounter where TCAS induces an NMAC that would not have occurred had the aircraft continued their trajectories without interruption. The horizontal profile of Figure 4-1(a) illustrates a parallel approach in which Aircraft 1 overshoots the approach path, proceeds to establish final approach on the same runway as Aircraft 2, and overtakes Aircraft 2. Though the paths overlap, the aircraft avoid an NMAC as

their trajectories are spatially separated. During the overtake portion, the aircraft converge horizontally but remain vertically separated by no less than 140 ft during horizontal conflict, thus remaining clear of conflict.


The NMAC induced by TCAS occurs at the perpendicular crossing (approximately the 75-second mark), though the intended flight paths (indicated by the dashed lines) avoid an NMAC by a small margin at the point of crossing (118 ft of vertical separation), as shown in Figure 4-1(b). The NMAC is induced when TCAS issues a Climb command to both aircraft simultaneously, and then issues one reversal to Aircraft 1 and two reversals to Aircraft 2, all within 10 s. The result of these reversals within a small timeframe, compounded with a delayed pilot response, leads to an NMAC at the point of crossing. In a similar real-life scenario, one or both aircraft could potentially be commanded to terminate the approach. However, this analysis is solely concerned with how each collision avoidance system performs with respect to the aircraft geometry, independent of air traffic controller intervention or pilot judgment.

Figure 4-1(c) shows the same encounter when using ACAS Xo as the alerting system. Notice that ACAS Xo alerts far less frequently than TCAS and does not issue a reversal. ACAS Xo only issues Descend and Do-Not-Descend commands to the aircraft to resolve the potential perpendicular crossing conflict. The sample encounter demonstrates the efficiency and effectiveness of ACAS Xo in alerting to avoid a potential NMAC.


Figure 4-2 depicts an encounter in which TCAS issues a nuisance alert to both aircraft, whereas ACAS Xo does not issue any alerts. The ACAS Xo vertical trajectories are not shown since they are the same as the nominal trajectories, indicated by the dashed lines in Figure 4-2(b). The horizontal trajectories shown in Figure 4-2(a) maintain about 3000 ft lateral separation on parallel approach. TCAS, which relies on straight-line trajectory projections, issues an alert to each aircraft just as Aircraft 1 begins its turn onto final approach. ACAS Xo, on the other hand, is tuned such that it delays alerting until it is certain that an NMAC is imminent without intervention.



(a) Horizontal Profile

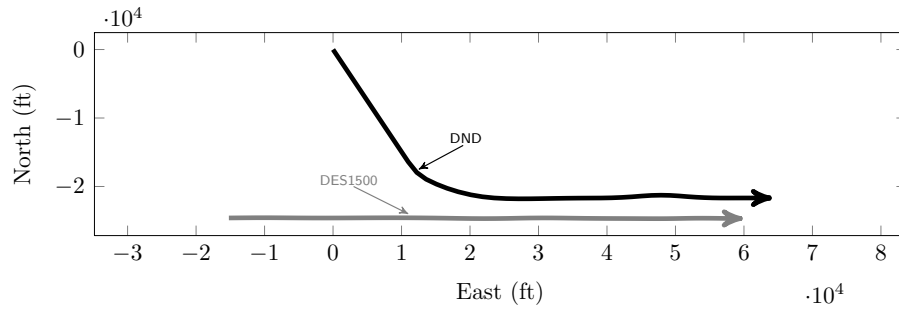


(b) TCAS Vertical Profile




(c) ACAS Xo Vertical Profile

Figure 4-1: Sample encounter



(a) Horizontal Profile



(b) TCAS Vertical Profile

Figure 4-2: Sample encounter

4.2 Individual Parameter Analysis

While observing the overall performance of the optimized ACAS Xo settings is critical, each parameter’s individual contribution to the improvement in performance should be understood so that parameter settings can be intelligently adjusted throughout future ACAS Xo development. Taking the final tuned ACAS Xo settings as a baseline, the 100,000 encounters of the historical blunder encounter set are re-simulated using the original ACAS Xa setting for a given parameter while holding all other parameters at the tuned ACAS Xo settings. Table 4.3 shows the performance metrics for each setting combination. For example, the first entry shows the performance metrics achieved using the ACAS Xa value for the alerting cost (r alert), while holding all

other values at the tuned ACAS Xo settings.

The results in Table 4.3 indicate that all of the ACAS Xa parameter settings, except for cycles, result in a higher nuisance alert rate than with the ACAS Xo settings. Compared to the ACAS Xa settings, the alerting cost of ACAS Xo increases, both conflict definitions decrease, and the slow closure range threshold decreases. Each of these changes results in more restricted alerting with ACAS Xo. In other words, the ACAS Xo settings for these parameters limit the set of scenarios in which issuing an RA is the optimal action. In contrast, the number of cycles increases with ACAS Xo compared to ACAS Xa, allowing alerting during a longer time period while the aircraft are in horizontal conflict.

Table 4.3: Metric Results using ACAS Xa Settings for One Parameter at a Time

Parameter	NMACs			Risk Ratio	Nuisance Rate	Strengthenings	Reversals
	Unresolved	Induced	Missed Alerts				
r alert	0	0	2	0.0014	0.205	279	296
cycles	7	1	2	0.0070	0.107	171	161
horiz conflict def	0	0	0	0.0000	0.516	196	797
vert conflict def	0	0	2	0.0014	0.241	524	487
slow clos rho thresh	0	0	0	0.0000	0.709	6,665	844

When using the ACAS Xa setting for the cycles parameter, the number of NMACs increases and safety decreases. More NMACs occur with the ACAS Xa cycles setting because the smaller cycles value inhibits alerting sooner after horizontal conflict has occurred. Once the aircraft are in horizontal conflict, the number of cycles determines the number of seconds for which RAs may continue to be issued, as illustrated in Figure 4-3. Earlier alerting inhibition increases the likelihood that a potentially resolvable NMAC remains unresolved. The remainder of the ACAS Xa parameter settings achieve an equal or better risk ratio compared to ACAS Xo, albeit with higher nuisance alert rates. With the exception of the effect that the number of cycles has on safety, the slow closure range threshold has the largest overall effect. When applying the ACAS Xa setting for this parameter, not only does the nuisance alert rate increase dramatically to 70.9%, but 6665 encounters include a strengthening command as opposed to 596 when using the ACAS Xo setting.

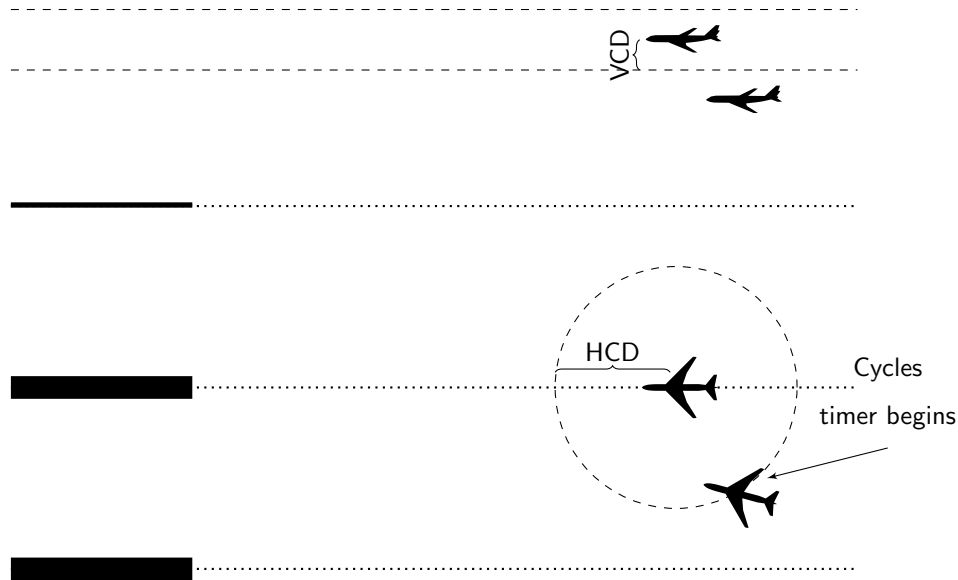


Figure 4-3: Effect of cycles parameter

4.3 Generated Encounter Set Performance

Two additional encounter sets are developed and utilized to further assess unnecessary alert rates with ACAS Xo and TCAS. In a similar fashion as the historical encounter set, the generated encounters sets are created by sampling trajectories from a trajectory library. The generated trajectory library differs from the historical trajectory library in that it contains trajectories intended to accurately reflect realistic IMC final approaches. In contrast, the historical trajectory library is intended to contain a wide range of possible scenarios, though many may not represent realistic IMC final approach trajectories. The generated encounter model, described in further detail in a technical report, uses a Markov representation through a dynamic Bayesian Network [41]. This approach results in the aircraft state (e.g., position, heading, speed, descent rate) being updated at each time step in a probabilistic manner. The Markov representation dictates that the probability distribution of the future state depends only on the current state.

Several assumptions were made during the development of the generated encounter model [41]:

- Final approach trajectories can be modeled by a Markov process.
- Aircraft vertical, lateral, and speed controls are independent.
- Final approach trajectories and course deviations are independent of airport, runway configuration, approach direction, and aircraft type.
- Glide slope deviations are independent of the glide slope.
- A pilot is in direct control of the aircraft with flight director guidance.
- Measurement noise and aircraft disturbances are modeled through randomness in the control process.
- Trajectories can be represented by heading, speed, and descent rate changes using first-order Euler update equations.
- Trajectories in the historical trajectory library are not interrupted by a collision avoidance system.

The two generated encounter sets are defined by the following airport configuration and requirements for aircraft entering the localizer beam, as illustrated in Figure 4-4:

- Aircraft begin their approach within 30 s of each other.
- Runways are spaced 3000 ft apart without stagger.
- The glide slope for both runways is 3°.
- The left aircraft intercepts the localizer 14–25 NM from the runway threshold at an altitude of 4000 ft.
- The right aircraft intercepts the localizer 14–25 NM from the runway threshold at an altitude of 5000 ft.
- Aircraft intercept their respective localizers at an angle of 20–45°.

Allowing for variance in the intercept angles allows for situations in which aircraft do not properly capture the ILS and large deviations that may occur later in the final approach.

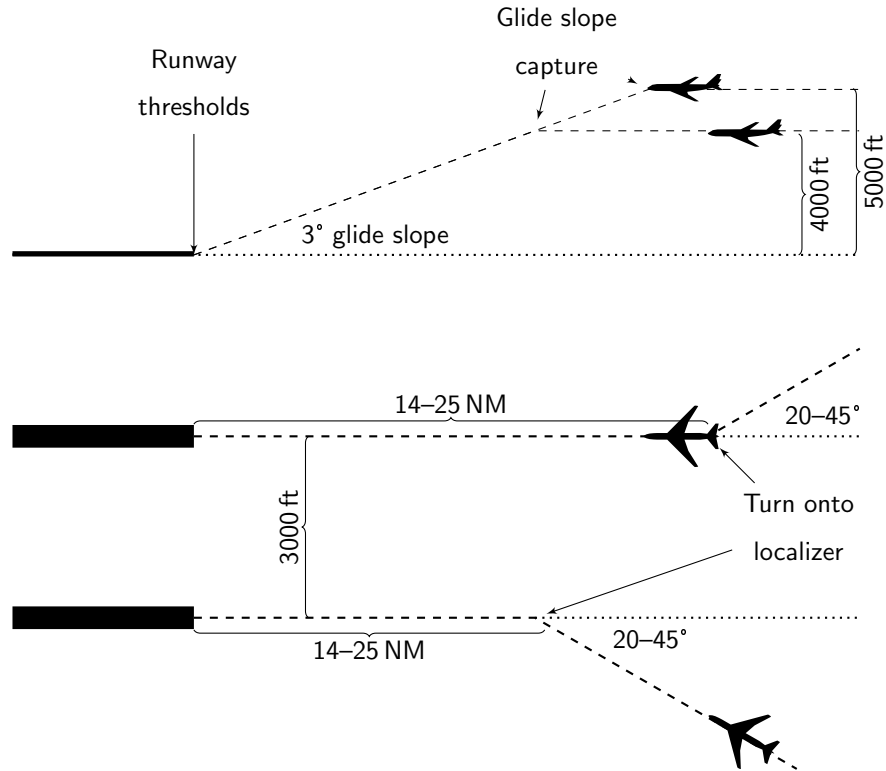


Figure 4-4: Generated encounter set configuration

The difference between the two generated encounter sets involves the prevalence of blunders. For stress-testing, an over-sampled blunder encounter set is desired. The blunder encounter set includes the following blunders and corresponding frequencies:

- Aircraft intercepting the localizer co-altitude (15 per 1000 aircraft approaches)
- Aircraft establishing the incorrect ILS approach (wrong runway) (3 per 1000 aircraft approaches)
- Aircraft breaking out toward the other runway (150 per 1000 aircraft approaches)

These blunder frequencies were chosen to reflect the tuning priorities for ACAS Xo.

Deviations from centerline during final approach are probabilistically distributed such that smaller deviations are much more likely than larger deviations. Aircraft

tracks that overshoot their localizer intercept or experience large centerline deviations during final approach are over-sampled from the original probability distribution by a factor of up to 1000, compared to a trajectory with little deviation. This over-sampling results in a higher proportion of trajectories that experience large deviations. The purpose of this blunder encounter set is to stress-test TCAS and ACAS Xo, so that their performance can be observed and differentiated during blunder scenarios. For the nominal generated encounter set, break-outs, incorrect ILS captures, and co-altitude join-ups are ignored.

Because ACAS Xo is tuned against an historical blunder encounter set, testing against the generated encounter sets is necessary to ensure no adverse effects exist from the original tuning. The historical encounter set was used for tuning since it effectively stress-tests the logic to resolve severe blunders, while also ensuring appropriate alert rates. The generated encounter sets engender more representative aircraft behaviors during CSPO by allowing the designer to specify the desired dynamics, instead of the historical trajectories dictating the aircraft behavior. Tables 4.4 and 4.5 compare the performance of ACAS Xo and TCAS for the generated blunder encounter set and the generated nominal (lateral deviations only) encounter set, respectively.

Table 4.4: Metric Results for the Generated Blunder Encounter Set

System	NMACs			Risk Ratio	Nuisance Rate	Strengthenings	Reversals
	Unresolved	Induced	Missed Alerts				
ACAS Xo	0	0	1	0.0011	0.022	158	27
TCAS	0	0	2	0.0023	0.133	111	224

Although ACAS Xo is not tuned for the two generated encounter sets, it still matches or outperforms TCAS in terms of both safety and alerting in both encounter sets. For the generated blunder encounter set, as shown in Table 4.4, TCAS misses two alerts and resolves the remainder of the 877 NMACs, and ACAS Xo misses one alert. However, each of the missed alerts for both systems is due to the slow convergence occurring below the alert inhibition altitude. By eliminating the missed alerts from consideration, both systems effectively resolve all NMACs in the encounter set.

Even though their safety levels are considered equivalent, ACAS Xo has a lower nuisance alert rate of 2.2% compared to 13.3% with TCAS. Furthermore, ACAS Xo only issues reversals in 27 encounters whereas TCAS reverses in 224 encounters. Although strengthenings should be minimized when possible, the slightly increased frequency with ACAS Xo is of no concern, especially considering the dramatic improvement in nuisance alert rate.

Table 4.5: Metric Results for the Generated Nominal Encounter Set

System	NMACs			Risk Ratio	Nuisance Rate	Strengthenings	Reversals
	Unresolved	Induced	Missed Alerts				
ACAS Xo	2	0	1	0.0192	0.042	107	38
TCAS	10	6	1	0.1090	0.181	376	925

For the nominal generated encounter set, as shown in Table 4.5, TCAS is unable to resolve ten of the 156 NMACs, misses one alert, and induces six NMACs that would not have occurred otherwise. ACAS Xo, on the other hand, is only unable to resolve two NMACs and misses one alert. Similar to the historical encounter set, the missed alert experienced by both systems is due to slow closure below the alert inhibition altitude. The nuisance alert rate is 4.2% with ACAS Xo compared to 18.1% with TCAS. ACAS Xo also experiences a lower number of encounters with strengthenings and reversals (107 and 38, respectively) compared to TCAS (376 and 925, respectively).

ACAS Xo outperforms TCAS with respect to safety and operational suitability for all three encounter sets. ACAS Xo is able to consistently prevent NMACs while also maintaining a nuisance alert rate significantly lower than TCAS. Reversals are much less frequent with ACAS Xo for each encounter set, and strengthenings are, at worst, only slightly higher than TCAS to offset the lower frequency of reversals and alerts.

4.4 Operational Performance Analysis

An operational performance model has been previously developed that allows for CSPO performance analysis at different runway and aircraft configurations. This model allows the designer to sweep through various configurations and quickly analyze the effects of varying different aspects. The model also allows for analysis of TCAS and ACAS Xo performance when presented with different navigational errors. The study in this section compares the operational performance of TCAS and the ACAS Xo logic during normal operations when no NMACs occur during the approach. Because no NMACs occur nominally in these encounters, any alerts issued by TCAS or ACAS Xo are classified as nuisance alerts.

The operational performance of the system is evaluated at a single point during final approach. The positions and velocities of the two aircraft at the evaluation point are defined by various parameters. These parameters include runway spacing and threshold stagger, threshold elevation and crossing height, aircraft distance from threshold, glide slope, ground speed, and navigational mode [4]. The left aircraft, hereafter referred to as the own aircraft, is simulated as starting 10 NM from the runway threshold at a 3° glide slope and 140 kt, resulting in an altitude of about 3200 ft. The right aircraft, hereafter referred to as the intruder, is initiated with intrail distance (stagger behind the own aircraft) varying between 0 and 1 NM at 0.1 NM increments, and ground speed differential (relative to the own aircraft) varying between 0 and 80 kt at 10 kt increments. The lateral runway separation is set at 3000 ft with no runway stagger, both runways at 0 ft elevation, and both aircraft crossing 50 ft above the threshold.

Four aircraft navigational modes are applied in the evaluations to reflect various navigational deviations [4]. Two of the modes sample distributions defined in the Collision Risk Model developed by the International Civil Aviation Organization [14]:

- **Perfect:** There is no lateral or vertical deviation from the intended course.
- **Required Navigation Performance (RNP):** A lateral deviation is introduced as an initial centerline bias. The deviation is sampled from a zero-mean




Figure 4-5: Probability (in percent) of alert

normal distribution per the RNP 0.1 standards [16, 26].

- **ILS CAT II Autopilot:** Lateral and vertical deviations are sampled from the Collision Risk Model distributions for the Category (CAT) II Autopilot mode.
- **ILS CAT II Flight Director:** Lateral and vertical deviations are sampled from the Collision Risk Model distributions for the CAT II Flight Director mode.

RNP defines performance standards related to an aircraft's ability to navigate precisely along a desired flight path. RNP utilizes ground-based and space-based navigation aids, as well as on-board systems. RNP is unique in that it requires an on-board performance monitoring and alerting system that monitors the actual navigation performance achieved during flight and alerts the pilot if standards are not met [6]. In the Autopilot mode, the pilot is not needed to control the aircraft. The autopilot uses the pre-loaded ILS approach specifications to guide the aircraft along its assigned flight path. In the Flight Director mode, the pilot is in control of the aircraft's trajectory. The on-board flight director displays the required pitch and bank angles on the attitude indicator for the pilot to follow.

The evaluations are performed using Monte Carlo simulation for 10,000 encounters at each combination of intrail distance and speed differential. The encounter trajectories are simulated for 15 s prior to the evaluation point so that the own aircraft can build a track of the intruder based on sensor measurements of range, bearing, and altitude at each time step.

The first set of plots in Figure 4-5 shows the alert probabilities for TCAS and ACAS Xo when both aircraft have perfect navigation. The region of maximum alert probability for TCAS is where the intrail distance is between 0 and 0.3 NM and the speed differential is between 50 and 80 kt. Speed differentials greater than 50 kt can be significant, especially considering the relatively slow speeds flown during approach. The areas of high alert probability indicate where the predicted time until closest approach is close to its minimum. As the intrail distance increases or the overtake speed decreases, the time until closest approach increases and TCAS will not alert

as often. In contrast, ACAS Xo alerts the most for intrail distances less than 0.1 NM at varying overtake speeds. Interestingly, ACAS Xo is not as sensitive to overtake speed as TCAS since the ACAS Xo contour runs parallel to the vertical axis (speed differential), whereas TCAS contours are crossed more regularly when varying overtake speed. With perfect navigation, albeit unrealistic, the alert rate with ACAS Xo is less than 2% at all points.

The second set of plots in Figure 4-5 shows the alert probabilities for TCAS and ACAS Xo when both aircraft use the RNP navigational mode. The regions of maximum alert probability for both TCAS and ACAS Xo expand due to the increased aircraft track uncertainty. For TCAS, the alert probability tends to be more sensitive to intrail distance than overtake speed at all but the smallest intrail distances, indicated by the more vertically oriented contour lines. However, for intrail distances less than 0.2 NM the contours become more horizontal and thus more sensitive to overtake speed. The ACAS Xo alert probability is almost completely dependent on intrail distance and insensitive to overtake speed due to the near-vertical contours. Again, there is a significantly smaller alert probability with ACAS Xo compared to TCAS. Even in the regions of greatest alert probability with ACAS Xo, just above 6%, the same regions with TCAS experience probabilities ranging from 8% to 35% depending on the overtake speed.

The third set of plots in Figure 4-5 shows the alert probabilities for TCAS and ACAS Xo when both aircraft are using the CAT II Autopilot navigational mode. The regions of maximum alert probability for both TCAS and ACAS Xo shrink compared to the RNP mode due to more precise navigation. The trends for TCAS and ACAS Xo are similar to those with RNP, though the alert probability of ACAS Xo approaches zero for overtake speeds less than 20 kt. The ACAS Xo alerting behavior also becomes slightly more dependent on overtake speed as the contours angle more toward the horizontal axis. The maximum alert probability with TCAS reaches 30%, whereas ACAS Xo slightly exceeds 2%.

The last set of plots in Figure 4-5 shows the alert probabilities for TCAS and ACAS Xo when both aircraft are using the CAT II Flight Director navigational mode.

The regions of maximum alert probability for both TCAS and ACAS Xo expand slightly from the CAT II Autopilot mode due to the lack of autopilot precision. The trends mirror those observed with the CAT II Autopilot mode. The maximum alert probability with TCAS approaches 35% and ACAS Xo approaches 3%.

The results illustrated in Figure 4-5 indicate that, when using TCAS, the greatest alert probabilities occur in regions of small intrail distance and high overtake speed. Either decreasing the intrail distance or increasing the overtake speed will monotonically increase the alert probability for any navigational mode. ACAS Xo also tends to experience the greatest alert probability in regions of small intrail distance and high overtake speed. However, small intrail distances with small overtake speeds sometimes experience similarly high alert probabilities, especially with less accurate navigational modes. This reduced sensitivity to overtake speed reflects the fact that ACAS Xo is tuned for sensitivity to slow-closure scenarios and does not necessarily require a high closure rate to issue an RA.

Along the same lines, the alert probability of ACAS Xo does not necessarily monotonically increase with overtake speed. Overall, ACAS Xo experiences much lower alert probabilities than TCAS. ACAS Xo reaches a maximum 6% alert probability with RNP navigation; otherwise, the nuisance alert rate with ACAS Xo does not exceed 3%. These ACAS Xo alert probabilities pose an exceptional advantage over TCAS which reaches an alert probability of at least 30% with all imperfect navigational modes.

ACAS Xo dramatically outperforms TCAS with respect to the operational performance model, which was developed independently of the model used in the tuning process. Therefore, the tuning process not only results in a solution that satisfies the original encounter set, but one that satisfies a variety of encounter models developed independently for CSPO.

4.5 Policy Examples

TCAS and ACAS Xo policies for generating advisories can be visualized for a given encounter at various vertical configurations. Likewise, alert probabilities can be determined for each system at various horizontal configurations. Understanding the specific alerting behavior of the two systems provides valuable insight into the strengths and weaknesses of each system. A no-blunder CSPO encounter is generated where two aircraft are separated by 1000 ft laterally on a side-by-side 3° glide slope, as illustrated in Figure 4-6. Separation of 1000 ft is chosen since both TCAS and ACAS Xo are likely to alert at that configuration. Aircraft 2 (not equipped with a collision avoidance system) is initiated at an altitude of 2500 ft while the initial altitude of Aircraft 1 is varied between 1500 and 3500 ft in 10 ft bins. Figure 4-7 shows the policies of Aircraft 1 for TCAS and ACAS Xo at this approach configuration. The policies displayed are based on Aircraft 1 continuing on its nominal trajectory (i.e., no pilot response). Removing pilot response allows for a better understanding of how vertical stagger affects each system's alerting behavior.




Figure 4-6: No-blunder CSPO configuration (vertical sweep)

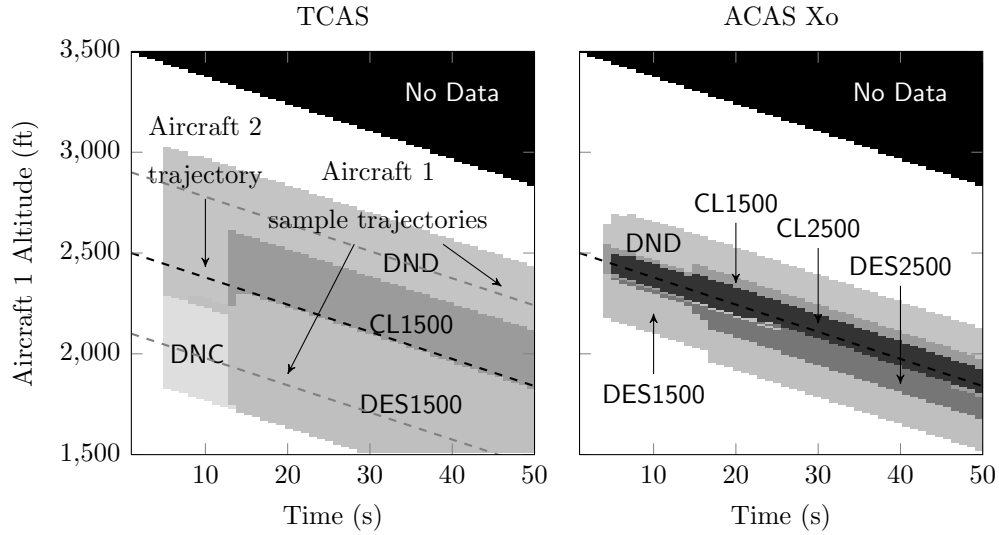


Figure 4-7: Example policy plots for Aircraft 1

The first notable difference between TCAS and ACAS Xo is the range of vertical separations at which an RA is issued. TCAS issues a preventive RA (e.g., Do Not Climb, Do Not Descend) to Aircraft 1 if Aircraft 2 is within 600 ft of Aircraft 1's altitude, and will issue a corrective RA (e.g., Climb, Descend) if Aircraft 2 remains between 300 ft below and 600 ft above Aircraft 1. ACAS Xo's alerting range is much smaller; Aircraft 2 must be between 200 ft below and 300 ft above for an RA to be issued to Aircraft 1.

The other main difference between the two systems is that ACAS Xo is likely to initially issue a corrective RA, whereas TCAS always begins with a preventive RA at this horizontal configuration. ACAS Xo is also quick to issue strengthenings so that separation is achieved without delay. Compared to TCAS, ACAS Xo shrinks the alerting range, indicated by the smaller shaded regions, while issuing more effective RAs (i.e., correctives and strengthenings) earlier in the encounter.

For varied horizontal configurations, the alert probabilities for TCAS and ACAS Xo are simulated for a no-blunder CSPO encounter where both aircraft are on a 3° glide slope beginning at an altitude of 2500 ft, as illustrated in Figure 4-8. Aircraft 1 is initiated at the origin for each simulation. The relative initial north and east positions of Aircraft 2 are both varied between -5000 ft and 5000 ft in 250 ft bins. Each bin

is randomly sampled and simulated with five different random noise generator seeds. The alert probability at a given point in Figure 4-9 represents the alert probability of Aircraft 1 when Aircraft 2 is initiated at that horizontal position.

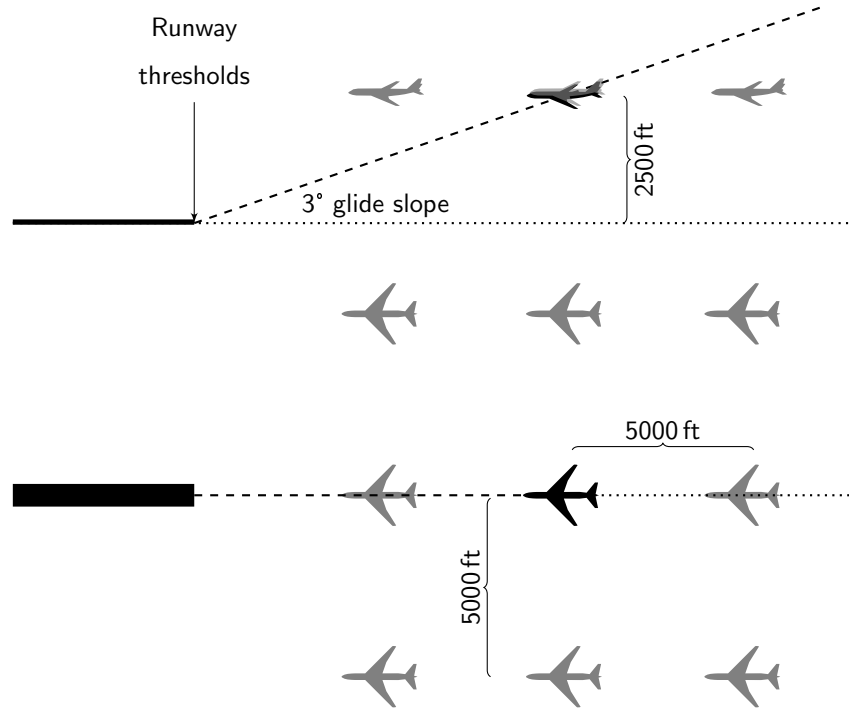


Figure 4-8: No-blunder CSPO configuration (horizontal sweep)

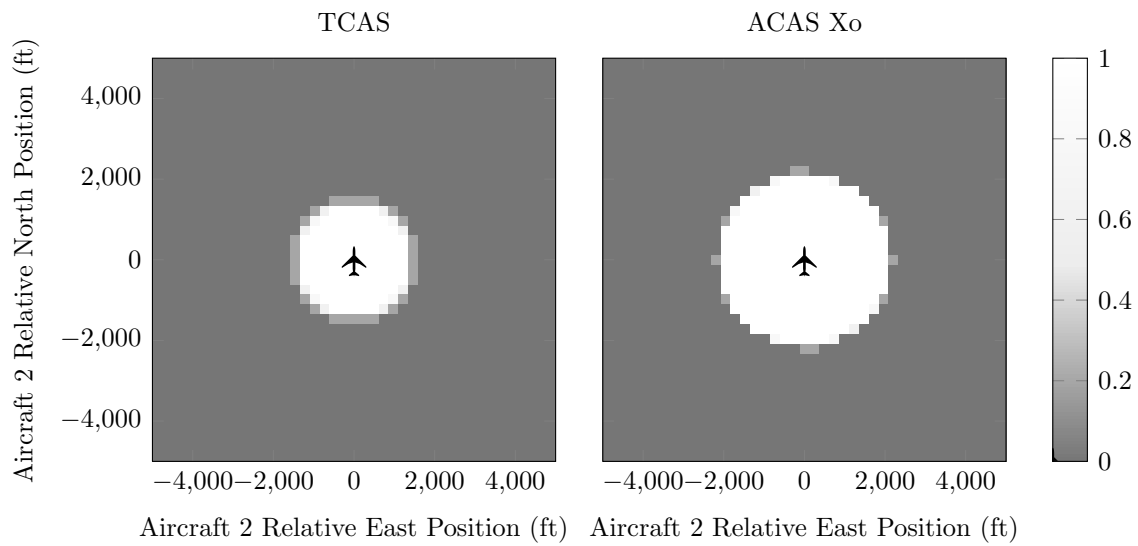


Figure 4-9: Example horizontal alert probabilities for Aircraft 1

The alerting range of TCAS at a co-altitude configuration is smaller than that of ACAS Xo, as shown in Figure 4-9. With co-altitude aircraft, TCAS maintains an alerting radius of 1375 ft in all directions, whereas ACAS Xo's alerting radius is 2125 ft. The alerting range of TCAS is determined by the distance modification heuristic. ACAS Xo, on the other hand, is tuned for nominal CSPO at 3000 ft runway separations.


The larger alerting radius of ACAS Xo also reflects the fact that the encounter set used in the tuning process over-samples blunders and essentially tunes ACAS Xo for stress-testing. Tuning for frequent blunders results in an alerting radius that is smaller than the nominal runway separation such that nuisance alerts are reduced, but large enough to successfully resolve blunders during stress-testing. As indicated in Figure 4-7, the alerting behavior of ACAS Xo is much more sensitive to vertical separation. This heightened sensitivity causes the alerting radius of ACAS Xo to decrease more quickly than TCAS with increasing vertical separation.

4.6 Application of Worst-Case Analysis to ACAS Xo Development


To illustrate the behavior of the ACAS Xo logic with respect to the worst-case analysis of Chapter 2, four sample encounters with 3000 ft runway separations are provided in Figures 4-10 through 4-13. Each blunder is initiated at the 15-second mark so that the own aircraft can build a track of the intruder based on sensor measurements at each time step.

The encounter in Figure 4-10 illustrates two aircraft on a collision course, with the intruder executing a constant-drift blunder at a 30° drift angle in the worst-case configuration. ACAS Xo alerts 3 s after the blunder begins (at the 18-second mark). The initial RA meets the required alert time of 9 s in Figure 2-5, and ACAS Xo is able to resolve the NMAC. Similarly, the intruder in Figure 4-11 blunders at 45°, and ACAS Xo alerts 3 s after the blunder begins. According to Figure 2-5, the available

alert time is 4 s, so the NMAC is successfully avoided.




(a) Horizontal Profile




(b) ACAS Xo Vertical Profile

Figure 4-10: Sample encounter



(a) Horizontal Profile




(b) ACAS Xo Vertical Profile


Figure 4-11: Sample encounter

The encounter in Figure 4-12 illustrates two aircraft on a collision course, with the intruder executing a constant-turn blunder at 1 deg/s. The blunder is initiated at the 15-second mark and, as shown in Figure 4-12(b), ACAS Xo alerts 20s after the blunder begins. This meets the required alert time of 22s in Figure 2-7. ACAS Xo is able to resolve the NMAC by strengthening the Climb command 3s after the initial RA.

Figure 4-13 illustrates a similar encounter with the intruder executing a constant-turn blunder at 3 deg/s. Again the blunder is initiated at the 15-second mark. As shown in Figure 4-13(b), ACAS Xo alerts 10s after the blunder begins. This exceeds the required alert time of 9s in Figure 2-7. ACAS Xo is thus unable to resolve the NMAC in time, resulting in a vertical miss distance of 63 ft when horizontal conflict first occurs.




(a) Horizontal Profile




(b) ACAS Xo Vertical Profile

Figure 4-12: Sample encounter



(a) Horizontal Profile



(b) ACAS Xo Vertical Profile

Figure 4-13: Sample encounter

Figure 4-14 compares the constant-drift available alert times with the actual initial alert times of ACAS Xo when simulated on each drift angle's worst-case scenario with 3000 ft runway separation, a 5 s pilot response delay, and a 2500 ft/min avoidance maneuver vertical rate. Similarly, Figure 4-15 compares the available and ACAS Xo alert times for the constant-turn blunder. As expected, the actual ACAS Xo initial alert time decreases as the blunder severity increases with drift angle or turn rate. ACAS Xo becomes unable to avoid an NMAC in the worst-case scenario once the drift angle exceeds 45° or, for the constant-turn case, the turn rate exceeds 2 deg/s.

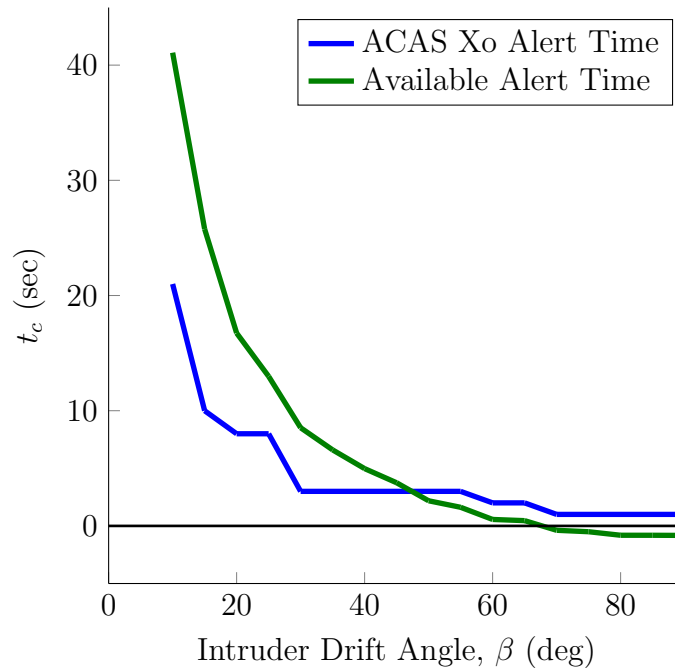


Figure 4-14: Worst-case available time to alert (t_c) vs. drift angle; Actual ACAS Xo alert time vs. drift angle

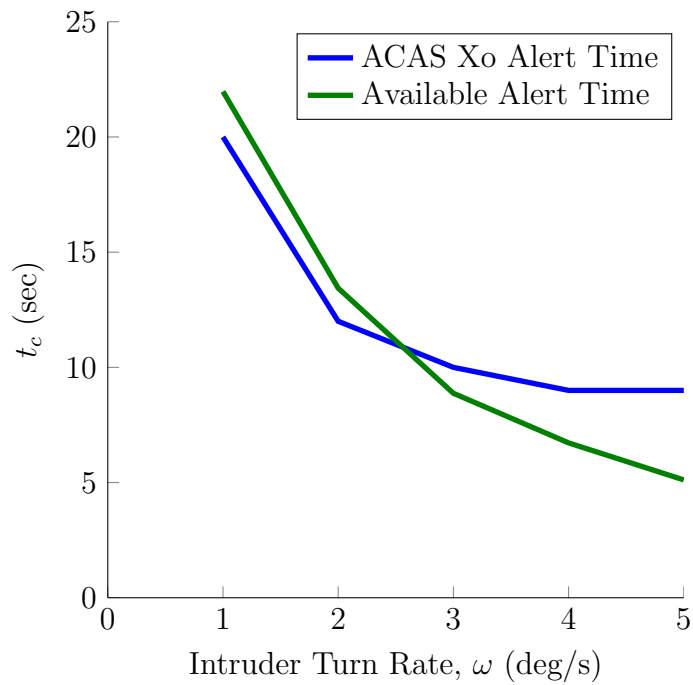


Figure 4-15: Worst-case available time to alert (t_c) vs. turn rate; Actual ACAS Xo alert time vs. turn rate

There are several implications for collision avoidance that arise from the results of the worst-case analysis and examples shown here. First, constant-turn blunders should be of more concern when attempting to extract the worst-case scenarios. As observed in Figure 4-15, constant-turn blunders can cause situations where the required time of alert is on the order of 15 s or less. Although the required alert times are smaller for some constant-drift scenarios, it is much more difficult for the ACAS Xo surveillance to quickly identify a turn blunder. The initial slow closure between the aircraft does not immediately prompt an alert and the final miss distances tend to be much smaller than for constant-drift blunders.

Another noteworthy implication is that state uncertainty and delayed pilot responses limit the ability to alert in time for certain worst-case scenarios. As state uncertainty and pilot response delay grow, the time available to alert in order to avoid an NMAC will decrease significantly. This becomes problematic in situations where the available time to alert is already small.

4.7 Discussion

This chapter analyzed the performance of ACAS Xo on various encounters after it was tuned on the historical encounter set in Chapter 3. Section 4.1 confirms that ACAS Xo outperforms TCAS on the historical encounter set in terms of both safety and alerting behavior. Fewer NMACs, nuisance alerts, strengthenings, and reversals result with ACAS Xo than TCAS. The effect of each individual parameter was analyzed as well. Without each of the five design parameters set at their respective values from the tuning procedure, either safety or alerting behavior is adversely affected. The ACAS Xa cycles setting has the most dramatic effect on safety by increasing the risk ratio by a factor of five. The slow closure rho threshold has a similarly negative effect on nuisance alert rate and the number of strengthenings, which increase by over a factor of five and four, respectively. The complexity of simultaneously tuning multiple design parameters to achieve desirable alerting behavior is made evident by the results of Section 4.2. Such complexity supports the need for the automated tuning

process applied in Chapter 3.

Section 4.3 uncovered the flexibility afforded by ACAS Xo in that one combination of design parameter settings can perform well on independent encounter sets. The tuning process results in design parameter settings that lead ACAS Xo to outperform TCAS on both generated encounter sets. Furthermore, the performance of ACAS Xo with respect to the operational performance model of Section 4.4 reinforces the flexibility of ACAS Xo. Even when varying navigational error and uncertainty is introduced in yet another independent encounter set, ACAS Xo still maintains consistently lower alert rates than TCAS at almost all combinations of overtake speed and intrail distance.

The alerting regions associated with ACAS Xo and TCAS in Section 4.5 reflect the foundations of each logic. As opposed to TCAS, the ACAS Xo horizontal and vertical alerting regions reflect the results of the tuning process. ACAS Xo shrinks the vertical alerting region compared to TCAS to account for the smaller vertical separations associated with CSPO. However, the horizontal alerting region is larger than that of TCAS due to the fact that ACAS Xo is specifically tuned for encounters with 3000 ft runway separations. The TCAS logic, on the other hand, cannot be similarly tuned for specific procedures and configurations such as CSPO with 3000 ft runway separations.

Finally, the ACAS Xo performance in Section 4.6 validates the results of Chapter 2 in that ACAS Xo is only able to avoid an NMAC in the worst-case scenarios when it alerts within the available time calculated in Chapter 2. ACAS Xo is able to resolve each NMAC in the example encounters of Section 4.6 except for the 3 deg/s constant-turn blunder. This failure indicates that more emphasis should be placed on constant-turn blunders since they are more difficult for the system to immediately detect than constant-drift blunders.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Conclusions and Further Work

The surrogate modeling process, paired with Max EI search, results in a tuned ACAS Xo logic that outperforms TCAS in terms of both safety and operational suitability. ACAS Xo results in fewer NMACs, nuisance alerts, strengthenings, and reversals than TCAS for the historical encounter set. ACAS Xo also outperforms TCAS with respect to the generated encounter sets and the operational performance model, which were developed independently of the historical encounter model for which ACAS Xo was tuned. For the operational performance model, the nuisance alert rate of ACAS Xo does not exceed 3% except during RNP navigation where the nuisance alert rate reaches 6%. TCAS reaches nuisance alert rates over 30% for all imperfect navigational modes.

The ability of ACAS Xo to provide protection with minimal alerting will increase the efficiency of CSPO, especially during IMC, and help achieve NextGen goals including increased throughput at high-volume airports. Furthermore, tuning ACAS Xo for specific aircraft behaviors and approach configurations may eliminate the need for additional collision avoidance systems such as ground-based PRM.

As for the worst-case analysis of collision avoidance, the results indicate that pilot response delay is an important factor in the limits of collision avoidance, and removing a 5 s delay from the equation can make a significant difference in many encounters. Auto-alerting is one possible means of completely eliminating pilot response delays. As a continuation of the worst-case analysis in Chapter 2, there are a number of

additional goals to accomplish. First, it would be useful to analyze how coordinated alerts would affect the previous analysis. There are currently systems in development that have the ability to coordinate collision avoidance alerts between multiple aircraft. This could potentially resolve many additional conflicts by providing alerts to the blundering aircraft as well as the non-blundering aircraft and increasing vertical separation at close range. This analysis also assumed that only Climb and Descend commands were issued by the system. Incorporating the ability to simultaneously command horizontal and vertical maneuvers could potentially provide greater separation and collision avoidance capability.

Additional work remains that can further enhance the performance of ACAS Xo. First, ACAS Xa and ACAS Xo functionality must be combined so that system operators may easily transition from one mode to another in flight. As an additional application, the use of ACAS Xo during parallel departures and corresponding special procedures will be relevant to many airports in addition to parallel approaches. Another area of interest is how an autopilot response model would affect the alerting logic, as opposed to the human pilot response model. Finally, additional data obtained from target aircraft, such as bank angle or autopilot status, may be included in future versions of ACAS Xo to enhance the surveillance capabilities of the system.

There are also improvements that can be made to the ACAS Xo tuning process. In this study the objective value weightings were determined by performing a weighting sweep alongside a hand-ranking of sample points to validate the selected weightings. This process can be improved by implementing methods that incorporate preference elicitation by learning from expert preferences. Furthermore, only one encounter set was used to directly tune ACAS Xo, and that encounter set does not necessarily represent a realistic probability distribution of encounter types. The encounter set used for tuning also only included encounters with 3000 ft runway separations. The wide range of possible encounter scenarios calls for further analysis on other encounter models that may represent CSPO encounters differently.

Regardless, the satisfactory performance of ACAS Xo on the generated encounter sets and operational performance model indicates that the tuning process can yield a

solution that performs well on independent CSPO encounter models not considered during optimization. The tuning process used here is not limited to ACAS Xo and can be applied to a variety of research efforts. Successful application requires that the user is equipped with a means of simulating results, able to reduce the scope of the problem to a small number of variables, and can meaningfully combine multiple output metrics into one representative performance metric.

The contributions of this thesis include: worst-case alerting requirements at various parallel approach configurations, proving optimization via surrogate modeling as a viable means of tuning collision avoidance systems, and achieving ACAS Xo performance superior to that of TCAS during CSPO. Tuning ACAS Xo for CSPO, and comparing TCAS performance with respect to various encounter models, displays the efficiency and effectiveness of the tuning process for an airborne collision avoidance system. The time required to complete the tuning process pales in comparison to the time it may take a human designer to determine an optimal combination of parameter settings. Even then the ideal balance of safety and operational suitability is unlikely to be achieved without the use of an automated surrogate modeling and tuning process.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A

Objective Value Weighting Sweep

A weighting sweep was performed to test various linear combinations of the primary performance metrics (risk ratio and nuisance alert rate) and to select the objective value weighting that most effectively tunes ACAS Xo for the historical CSPO encounter set.

Tables 3.6 and 3.7 in Section 3.8 contain the infill results for the 95%/5% weighting; therefore, the tables are not repeated here. Tables A.1 and A.2 show the eighty infill points and their simulation results when a 97.5%/2.5% weighting is applied in calculating the objective value. Similarly, Tables A.3 through A.10 show the infill results for weighting combinations of 92.5%/7.5%, 90%/10%, 85%/15%, and 80%/20%.

The lowest objective value using the 97.5%/2.5% weighting is achieved at data point 115 in Table A.2. Compared to the best performance achieved with the 95%/5% weighting (0.0014 risk ratio and 12.9% nuisance alert rate), data point 115 results in only one NMAC (0.0007 risk ratio) and a 17.0% nuisance alert rate. The increased weighting toward risk ratio results in an unacceptably high nuisance alert rate in return for one fewer NMAC.

Similarly, the lowest objective values achieved with the 92.5%/7.5% and 90%/10% weightings represent points that result in one additional NMAC and only slight changes in nuisance alert rate from 12.9%. The lowest objective value using the 92.5%/7.5% weighting is achieved at data point 94 in Table A.4, with a 0.0021 risk ratio and 13.1% nuisance alert rate.

The 90%/10% weighting minimizes the objective value at data point 84 in Table A.6, with a 0.0021 risk ratio and 12.8% nuisance alert rate. The 0.1% reduction in nuisance alert rate cannot be concluded to warrant any subsequent detriment to safety. The sample point which minimizes nuisance alert rate, while matching the safety level achieved with the 95%/5% weighting (0.0014 risk ratio), is only able to reduce the nuisance alert to 14.0%. The priorities communicated by the heavier weighting of nuisance alert rate lead the lowest objective values to represent risk ratios no lower than 0.0021, which is undesirable for the purposes of ACAS Xo.

The lowest objective values achieved with the 85%/15% and 80%/20% weightings represent points that result in additional NMACs and slight decreases in nuisance alert rate from 12.9%. The lowest objective value using the 85%/15% weighting is achieved at data point 86 in Table A.8, with a 0.0028 risk ratio and 11.6% nuisance alert rate. The 80%/20% weighting minimizes the objective value at data point 66 in Table A.9, with a 0.0056 risk ratio and 10.3% nuisance alert rate.

Again, as the weightings shift more in favor of nuisance alert rate, the supposedly optimal sample points result in unacceptably high risk ratios and marginally reduced nuisance alert rates, as shown in Figure A-1. For these reasons, the 95%/5% weighting is determined to best represent the priorities of ACAS Xo performance for the historical CSPO encounter set and yield the most desirable behavior during the infill process.

Table A.1: Evaluation Results for the First Forty Infill Points (97.5%/2.5%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
41	-0.0025	7	200	5	1,875		0.0028	0.134	18.3308
42	-0.0025	5	312	75	1,500		0.0021	0.158	6.7043
43	-0.0025	4	400	20	2,323		0.0028	0.199	16.0938
44	-0.0050	4	400	0	1,835		0.3912	0.083	2.1058
45	-0.0036	4	400	5	1,500		0.0035	0.127	5.7239
46	-0.0048	7	400	5	2,128		0.0021	0.152	27.9734
47	-0.0047	7	400	11	2,000		0.0021	0.145	21.2414
48	-0.0050	6	400	4	2,500		0.0014	0.215	56.4351
49	-0.0050	4	400	8	2,500		0.0021	0.194	34.5396
50	-0.0050	7	400	45	1,500		0.0028	0.128	9.8495
51	-0.0049	4	201	57	2,500		0.0049	0.178	43.0989
52	-0.0032	6	294	4	2,500		0.0014	0.227	51.0968
53	-0.0025	4	400	91	2,500		0.0028	0.245	22.2137
54	-0.0050	7	200	97	1,500		0.0133	0.086	16.1840
55	-0.0050	7	397	14	1,500		0.0028	0.128	9.8244
56	-0.0025	7	200	50	1,500		0.0035	0.124	9.6211
57	-0.0050	4	200	8	1,500		0.0168	0.072	10.9948
58	-0.0025	4	400	3	1,500		0.0042	0.160	5.2062
59	-0.0050	4	200	41	2,500		0.0049	0.178	43.0733
60	-0.0050	6	393	4	1,500		0.0021	0.121	9.3934
61	-0.0050	7	400	8	2,500		0.0021	0.222	61.8660
62	-0.0025	7	200	67	2,498		0.0021	0.232	59.9509
63	-0.0050	7	400	45	2,500		0.0021	0.222	61.9161
64	-0.0050	4	400	87	2,500		0.0021	0.194	34.5396
65	-0.0041	5	238	50	2,250		0.0028	0.136	35.7827
66	-0.0025	7	200	38	1,500		0.0035	0.124	9.6461
67	-0.0050	4	400	71	2,500		0.0021	0.194	34.5895
68	-0.0040	7	200	100	2,250		0.0021	0.140	47.3864
69	-0.0050	4	400	94	1,500		0.0077	0.103	8.0000
70	-0.0025	4	200	4	1,500		0.0056	0.103	7.2254
71	-0.0040	7	201	100	2,211		0.0035	0.135	44.6812
72	-0.0050	7	200	21	1,503		0.0133	0.086	16.2341
73	-0.0043	6	355	49	2,094		0.0014	0.152	22.2986
74	-0.0037	5	338	13	1,672		0.0042	0.126	9.8480
75	-0.0049	5	400	0	1,522		0.2582	0.099	2.5965
76	-0.0025	4	200	27	2,500		0.0021	0.210	32.4047
77	-0.0050	5	280	57	1,934		0.0049	0.107	22.6539
78	-0.0036	7	200	63	2,078		0.0042	0.126	34.9476
79	-0.0050	7	363	0	1,500		0.2134	0.114	3.0615
80	-0.0025	7	200	83	1,500		0.0035	0.124	9.6211

Table A.2: Evaluation Results for the Second Forty Infill Points (97.5%/2.5%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
81	-0.0025	4	400	4	1,897		0.0035	0.168	7.6384
82	-0.0043	4	304	69	2,259		0.0021	0.140	22.2115
83	-0.0050	4	399	71	1,500		0.0070	0.103	8.1752
84	-0.0025	4	400	4	1,938		0.0035	0.169	8.1144
85	-0.0050	4	400	53	1,500		0.0077	0.103	8.0001
86	-0.0025	7	200	38	2,500		0.0021	0.233	60.4519
87	-0.0045	5	295	50	1,934		0.0035	0.112	20.1347
88	-0.0025	7	398	10	2,500		0.0021	0.281	39.2992
89	-0.0050	4	200	62	1,500		0.0168	0.072	10.9948
90	-0.0025	7	400	29	1,500		0.0021	0.201	7.0964
91	-0.0025	4	200	17	1,500		0.0056	0.103	7.2504
92	-0.0025	4	328	1	1,650		0.2568	0.128	2.0496
93	-0.0036	5	375	48	2,370		0.0007	0.191	30.5360
94	-0.0048	5	326	88	1,702		0.0063	0.110	12.7577
95	-0.0050	7	400	59	2,500		0.0021	0.222	61.9160
96	-0.0050	6	244	34	2,094		0.0049	0.116	35.4879
97	-0.0050	7	400	74	2,500		0.0021	0.222	61.9161
98	-0.0050	7	341	100	1,617		0.0042	0.118	12.4396
99	-0.0044	6	259	29	2,140		0.0028	0.134	31.8553
100	-0.0050	5	288	52	1,969		0.0056	0.108	23.7558
101	-0.0050	7	400	25	2,500		0.0021	0.222	61.9160
102	-0.0033	6	313	87	1,673		0.0021	0.147	9.4428
103	-0.0027	5	400	53	1,875		0.0021	0.174	9.8196
104	-0.0025	4	399	6	1,521		0.0035	0.160	5.2310
105	-0.0031	6	238	19	2,407		0.0007	0.188	46.9080
106	-0.0033	6	277	27	1,969		0.0021	0.140	17.8360
107	-0.0050	4	200	85	1,500		0.0168	0.072	10.9948
108	-0.0031	7	384	50	2,000		0.0021	0.185	14.0308
109	-0.0025	4	388	98	2,125		0.0035	0.178	10.6235
110	-0.0041	7	393	28	2,061		0.0021	0.162	21.4076
111	-0.0025	4	400	80	1,500		0.0035	0.160	5.1812
112	-0.0047	6	344	49	1,677		0.0035	0.118	13.3398
113	-0.0025	4	400	35	1,500		0.0035	0.160	5.1560
114	-0.0031	7	388	100	1,624		0.0021	0.176	8.3961
115	-0.0037	4	382	46	2,344		0.0007	0.170	21.7904
116	-0.0025	7	399	4	1,930		0.0021	0.209	10.9039
117	-0.0050	6	275	1	1,750		0.3737	0.081	2.3536
118	-0.0029	6	300	75	1,937		0.0028	0.147	15.1183
119	-0.0050	5	336	0	1,672		0.3457	0.085	2.3328
120	-0.0038	6	346	57	1,676		0.0021	0.134	10.5306

Table A.3: Evaluation Results for the First Forty Infill Points (92.5%/7.5%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
41	-0.0029	4	323	50	1,500		0.0042	0.129	0.0136
42	-0.0050	4	307	0	2,489		0.5521	0.066	0.5156
43	-0.0025	7	201	5	1,500		0.0035	0.124	0.0126
44	-0.0041	6	200	7	1,500		0.0084	0.092	0.0147
45	-0.0025	7	208	0	1,750		0.2841	0.124	0.2721
46	-0.0025	4	200	40	2,500		0.0021	0.210	0.0177
47	-0.0050	7	200	10	1,500		0.0133	0.086	0.0188
48	-0.0050	4	400	53	1,500		0.0077	0.103	0.0148
49	-0.0025	7	209	4	2,063		0.0021	0.158	0.0138
50	-0.0039	7	399	4	1,500		0.0021	0.151	0.0132
51	-0.0025	4	400	91	2,500		0.0028	0.245	0.0209
52	-0.0025	4	200	27	2,500		0.0021	0.210	0.0177
53	-0.0025	4	200	4	1,500		0.0056	0.103	0.0129
54	-0.0050	4	200	8	2,500		0.0049	0.178	0.0179
55	-0.0050	7	293	0	2,416		0.5038	0.078	0.4719
56	-0.0025	7	202	81	1,500		0.0021	0.140	0.0125
57	-0.0050	7	397	61	1,500		0.0028	0.128	0.0122
58	-0.0025	7	202	3	1,500		0.0021	0.140	0.0125
59	-0.0025	7	200	17	1,500		0.0035	0.124	0.0126
60	-0.0025	4	200	37	1,500		0.0056	0.103	0.0129
61	-0.0050	7	200	44	2,500		0.0042	0.200	0.0189
62	-0.0050	4	400	88	1,500		0.0077	0.103	0.0148
63	-0.0025	4	204	29	1,500		0.0056	0.115	0.0138
64	-0.0050	7	400	73	2,500		0.0021	0.222	0.0186
65	-0.0050	7	291	64	2,474		0.0021	0.197	0.0167
66	-0.0025	7	200	66	2,500		0.0021	0.233	0.0194
67	-0.0025	7	400	21	1,500		0.0021	0.202	0.0171
68	-0.0037	6	326	75	1,644		0.0021	0.134	0.0120
69	-0.0050	7	399	25	2,500		0.0021	0.222	0.0186
70	-0.0050	4	400	12	1,500		0.0077	0.103	0.0148
71	-0.0025	4	400	7	1,500		0.0042	0.160	0.0159
72	-0.0025	7	400	58	2,500		0.0021	0.281	0.0230
73	-0.0050	7	200	85	1,500		0.0133	0.086	0.0188
74	-0.0050	4	201	98	1,500		0.0168	0.072	0.0209
75	-0.0050	4	200	56	2,500		0.0049	0.178	0.0179
76	-0.0047	7	399	0	1,521		0.2099	0.121	0.2032
77	-0.0041	6	341	68	1,659		0.0028	0.127	0.0121
78	-0.0025	4	374	0	1,500		0.1882	0.147	0.1852
79	-0.0044	4	200	100	1,703		0.0126	0.083	0.0179
80	-0.0025	7	305	4	1,500		0.0021	0.178	0.0153

Table A.4: Evaluation Results for the Second Forty Infill Points (92.5%/7.5%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
81	-0.0050	4	314	98	1,521		0.0070	0.094	0.0135
82	-0.0045	4	306	68	2,256		0.0021	0.138	0.0123
83	-0.0050	4	400	71	1,500		0.0077	0.103	0.0148
84	-0.0050	7	400	94	2,500		0.0021	0.221	0.0186
85	-0.0050	7	400	34	1,500		0.0028	0.128	0.0122
86	-0.0050	7	293	50	2,461		0.0021	0.192	0.0163
87	-0.0041	5	259	45	1,535		0.0056	0.105	0.0131
88	-0.0050	7	200	48	1,500		0.0133	0.086	0.0188
89	-0.0025	7	200	18	2,500		0.0021	0.233	0.0194
90	-0.0050	4	356	100	1,665		0.0070	0.106	0.0144
91	-0.0050	7	200	65	1,500		0.0133	0.086	0.0188
92	-0.0028	6	224	51	2,052		0.0014	0.143	0.0121
93	-0.0025	7	224	6	2,000		0.0021	0.154	0.0135
94	-0.0044	5	255	25	2,166		0.0021	0.131	0.0117
95	-0.0050	7	200	4	1,687		0.0112	0.091	0.0172
96	-0.0048	4	400	5	2,380		0.0021	0.162	0.0141
97	-0.0050	7	288	0	1,947		0.4241	0.084	0.3986
98	-0.0025	6	281	6	1,904		0.0028	0.159	0.0145
99	-0.0045	4	233	50	2,021		0.0084	0.103	0.0155
100	-0.0025	4	310	100	1,937		0.0028	0.153	0.0140
101	-0.0050	7	203	0	1,656		0.3688	0.076	0.3468
102	-0.0044	5	385	34	2,375		0.0014	0.177	0.0146
103	-0.0034	6	350	50	1,688		0.0028	0.142	0.0132
104	-0.0025	4	400	99	1,500		0.0042	0.160	0.0159
105	-0.0038	5	372	52	2,367		0.0007	0.186	0.0146
106	-0.0050	7	400	3	2,153		0.0021	0.151	0.0133
107	-0.0025	7	360	0	2,121		0.3058	0.183	0.2966
108	-0.0025	7	200	14	1,500		0.0035	0.124	0.0126
109	-0.0050	7	200	89	2,500		0.0042	0.200	0.0189
110	-0.0043	4	400	28	1,500		0.0049	0.113	0.0130
111	-0.0050	5	341	100	2,105		0.0028	0.125	0.0120
112	-0.0029	6	212	22	2,035		0.0014	0.142	0.0119
113	-0.0041	7	400	4	1,934		0.0021	0.155	0.0136
114	-0.0050	4	353	100	2,172		0.0035	0.130	0.0130
115	-0.0050	4	400	42	2,500		0.0021	0.194	0.0165
116	-0.0025	4	310	100	1,969		0.0028	0.154	0.0141
117	-0.0048	4	394	38	2,353		0.0021	0.157	0.0137
118	-0.0039	5	242	33	2,172		0.0035	0.129	0.0129
119	-0.0025	4	268	0	1,555		0.2505	0.112	0.2402
120	-0.0031	6	328	23	1,673		0.0028	0.150	0.0139

Table A.5: Evaluation Results for the First Forty Infill Points (90%/10%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
41	-0.0029	4	323	50	1,500		0.0042	0.129	0.0167
42	-0.0050	4	306	0	2,469		0.5500	0.066	0.5016
43	-0.0033	4	200	20	2,500		0.0035	0.198	0.0229
44	-0.0025	7	200	5	1,500		0.0028	0.124	0.0149
45	-0.0025	6	400	5	2,500		0.0014	0.272	0.0285
46	-0.0032	6	200	0	1,606		0.3310	0.086	0.3065
47	-0.0050	4	400	4	1,500		0.0070	0.103	0.0166
48	-0.0050	4	397	11	1,500		0.0070	0.103	0.0166
49	-0.0030	7	400	4	1,500		0.0014	0.179	0.0192
50	-0.0037	4	200	5	2,500		0.0042	0.191	0.0229
51	-0.0045	6	200	5	2,000		0.0056	0.106	0.0156
52	-0.0025	4	203	40	2,500		0.0021	0.217	0.0236
53	-0.0025	7	400	58	2,500		0.0014	0.281	0.0294
54	-0.0050	4	400	87	2,500		0.0021	0.194	0.0213
55	-0.0050	7	200	45	2,500		0.0042	0.200	0.0238
56	-0.0025	7	200	68	2,500		0.0014	0.233	0.0245
57	-0.0025	7	200	8	1,500		0.0028	0.124	0.0149
58	-0.0050	7	200	33	1,500		0.0133	0.086	0.0206
59	-0.0025	7	400	21	1,500		0.0014	0.202	0.0214
60	-0.0025	7	200	92	2,500		0.0014	0.233	0.0245
61	-0.0025	4	400	29	1,500		0.0035	0.160	0.0192
62	-0.0050	4	200	8	1,500		0.0168	0.072	0.0223
63	-0.0050	7	292	44	2,469		0.0014	0.195	0.0207
64	-0.0040	5	200	100	1,788		0.0077	0.096	0.0165
65	-0.0038	4	353	45	1,875		0.0021	0.129	0.0148
66	-0.0050	7	400	13	1,500		0.0021	0.128	0.0147
67	-0.0025	7	200	81	1,500		0.0028	0.124	0.0149
68	-0.0033	6	275	50	1,733		0.0028	0.131	0.0156
69	-0.0050	4	200	98	1,500		0.0168	0.072	0.0223
70	-0.0050	4	200	61	1,500		0.0168	0.072	0.0223
71	-0.0050	7	400	55	1,500		0.0021	0.128	0.0147
72	-0.0025	4	392	98	1,500		0.0035	0.160	0.0192
73	-0.0037	6	326	51	1,653		0.0021	0.133	0.0152
74	-0.0031	6	380	0	1,500		0.1728	0.153	0.1709
75	-0.0050	5	285	0	1,969		0.4605	0.071	0.4215
76	-0.0050	5	295	0	2,109		0.4955	0.070	0.4529
77	-0.0050	4	400	71	1,500		0.0070	0.103	0.0166
78	-0.0048	5	287	100	2,067		0.0042	0.117	0.0155
79	-0.0025	7	200	65	1,500		0.0028	0.124	0.0149
80	-0.0025	4	275	4	1,912		0.0028	0.137	0.0163

Table A.6: Evaluation Results for the Second Forty Infill Points (90%/10%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
81	-0.0025	4	200	36	1,500		0.0056	0.103	0.0153
82	-0.0045	5	225	87	1,508		0.0084	0.091	0.0166
83	-0.0032	5	234	53	1,500		0.0042	0.110	0.0148
84	-0.0050	7	400	96	1,500		0.0021	0.128	0.0147
85	-0.0047	5	337	0	1,500		0.2722	0.091	0.2541
86	-0.0037	5	231	62	1,531		0.0070	0.102	0.0165
87	-0.0048	5	400	100	2,500		0.0007	0.210	0.0216
88	-0.0050	4	200	26	2,500		0.0056	0.178	0.0228
89	-0.0025	4	200	29	1,500		0.0056	0.103	0.0153
90	-0.0048	5	250	99	2,168		0.0056	0.120	0.0170
91	-0.0025	7	282	9	1,898		0.0014	0.167	0.0179
92	-0.0025	7	400	71	1,500		0.0014	0.202	0.0214
93	-0.0025	7	265	2	1,552		0.0014	0.159	0.0172
94	-0.0037	5	235	47	2,017		0.0035	0.119	0.0151
95	-0.0025	4	400	72	2,500		0.0021	0.245	0.0264
96	-0.0045	5	209	100	1,531		0.0091	0.092	0.0174
97	-0.0025	7	400	2	1,500		0.0014	0.202	0.0214
98	-0.0050	7	200	75	2,500		0.0042	0.200	0.0238
99	-0.0025	4	200	90	1,500		0.0056	0.103	0.0153
100	-0.0042	5	203	0	2,163		0.5059	0.069	0.4623
101	-0.0025	7	207	3	1,500		0.0014	0.140	0.0153
102	-0.0047	5	297	86	2,406		0.0014	0.167	0.0179
103	-0.0049	7	295	43	2,445		0.0014	0.187	0.0199
104	-0.0050	5	299	0	2,168		0.5073	0.069	0.4635
105	-0.0034	5	218	49	1,578		0.0049	0.107	0.0151
106	-0.0049	7	200	6	2,500		0.0035	0.200	0.0232
107	-0.0025	4	316	100	1,500		0.0021	0.143	0.0162
108	-0.0025	7	299	3	1,500		0.0014	0.158	0.0171
109	-0.0048	7	275	75	2,500		0.0021	0.211	0.0229
110	-0.0050	7	200	56	2,500		0.0042	0.200	0.0238
111	-0.0047	5	359	69	2,500		0.0007	0.212	0.0218
112	-0.0048	6	296	100	2,062		0.0028	0.122	0.0147
113	-0.0039	5	200	94	1,757		0.0077	0.096	0.0165
114	-0.0048	5	275	100	2,500		0.0014	0.199	0.0211
115	-0.0045	5	275	36	1,500		0.0056	0.100	0.0151
116	-0.0025	7	400	6	1,500		0.0014	0.202	0.0214
117	-0.0025	4	400	12	2,000		0.0028	0.171	0.0196
118	-0.0047	5	266	100	1,500		0.0070	0.097	0.0160
119	-0.0047	6	325	63	2,000		0.0021	0.129	0.0148
120	-0.0047	6	399	24	2,500		0.0021	0.219	0.0238

Table A.7: Evaluation Results for the First Forty Infill Points (85%/15%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
41	-0.0025	7	200	5	2,500		0.0014	0.233	0.0361
42	-0.0030	4	325	36	1,500		0.0042	0.128	0.0228
43	-0.0025	5	200	0	1,813		0.3779	0.093	0.3351
44	-0.0050	4	400	9	1,502		0.0070	0.103	0.0213
45	-0.0050	4	400	4	1,500		0.0070	0.103	0.0213
46	-0.0042	5	213	87	1,568		0.0091	0.096	0.0221
47	-0.0025	7	200	81	1,500		0.0028	0.124	0.0210
48	-0.0025	7	250	4	1,500		0.0014	0.140	0.0222
49	-0.0025	4	200	92	2,387		0.0021	0.170	0.0273
50	-0.0041	7	200	4	2,500		0.0028	0.209	0.0337
51	-0.0025	4	400	4	2,499		0.0021	0.244	0.0385
52	-0.0050	7	399	57	2,500		0.0014	0.222	0.0344
53	-0.0050	4	304	0	2,279		0.5297	0.068	0.4605
54	-0.0025	4	200	39	2,500		0.0021	0.210	0.0333
55	-0.0025	7	400	30	1,500		0.0014	0.202	0.0315
56	-0.0045	7	300	13	1,954		0.0042	0.124	0.0221
57	-0.0025	4	200	50	1,500		0.0056	0.103	0.0202
58	-0.0037	7	400	5	1,500		0.0014	0.156	0.0247
59	-0.0025	4	400	5	2,074		0.0028	0.175	0.0286
60	-0.0050	7	204	33	1,500		0.0091	0.095	0.0220
61	-0.0050	4	400	54	2,496		0.0021	0.193	0.0307
62	-0.0050	7	200	44	2,500		0.0042	0.200	0.0336
63	-0.0050	7	200	60	1,500		0.0133	0.086	0.0242
64	-0.0050	4	400	12	1,500		0.0070	0.103	0.0213
65	-0.0050	4	400	14	1,750		0.0056	0.108	0.0210
66	-0.0050	4	400	72	1,500		0.0070	0.103	0.0213
67	-0.0025	4	397	69	2,500		0.0021	0.245	0.0385
68	-0.0050	7	400	74	1,500		0.0021	0.128	0.0209
69	-0.0050	7	294	0	1,995		0.4360	0.083	0.3830
70	-0.0050	7	399	49	1,500		0.0021	0.128	0.0209
71	-0.0050	7	251	5	1,904		0.0049	0.107	0.0202
72	-0.0050	7	200	67	1,500		0.0133	0.086	0.0242
73	-0.0050	7	400	8	2,500		0.0014	0.222	0.0344
74	-0.0025	7	200	15	1,500		0.0028	0.124	0.0210
75	-0.0050	4	201	29	1,500		0.0161	0.078	0.0254
76	-0.0050	5	200	5	1,500		0.0154	0.078	0.0248
77	-0.0044	5	200	100	1,683		0.0119	0.089	0.0234
78	-0.0042	6	327	50	2,120		0.0014	0.144	0.0227
79	-0.0025	4	400	7	2,500		0.0021	0.245	0.0385
80	-0.0038	4	313	100	2,266		0.0021	0.148	0.0241

Table A.8: Evaluation Results for the Second Forty Infill Points (85%/15%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
81	-0.0050	4	201	97	1,500		0.0161	0.078	0.0254
82	-0.0046	7	335	88	1,750		0.0021	0.127	0.0208
83	-0.0050	4	200	100	2,500		0.0056	0.178	0.0314
84	-0.0037	4	300	100	2,239		0.0028	0.138	0.0231
85	-0.0025	7	200	22	1,500		0.0028	0.124	0.0210
86	-0.0033	6	242	50	1,565		0.0028	0.116	0.0197
87	-0.0044	6	331	50	2,110		0.0014	0.140	0.0223
88	-0.0050	4	373	42	1,500		0.0070	0.103	0.0213
89	-0.0050	7	400	93	2,500		0.0014	0.222	0.0344
90	-0.0044	6	400	0	2,082		0.3870	0.112	0.3457
91	-0.0050	7	400	24	1,500		0.0021	0.128	0.0209
92	-0.0043	5	306	56	1,538		0.0049	0.113	0.0211
93	-0.0026	7	400	3	1,621		0.0014	0.200	0.0312
94	-0.0025	4	200	65	1,500		0.0056	0.103	0.0202
95	-0.0045	7	333	100	1,609		0.0028	0.125	0.0211
96	-0.0038	4	300	5	1,998		0.0042	0.117	0.0211
97	-0.0042	5	250	22	2,156		0.0042	0.124	0.0222
98	-0.0043	6	338	50	2,148		0.0014	0.145	0.0229
99	-0.0050	4	212	69	1,502		0.0161	0.078	0.0254
100	-0.0025	4	200	56	2,500		0.0021	0.210	0.0333
101	-0.0045	6	350	25	1,982		0.0014	0.131	0.0208
102	-0.0050	7	204	47	2,126		0.0042	0.122	0.0219
103	-0.0030	7	400	0	1,859		0.2645	0.164	0.2495
104	-0.0042	5	294	50	1,531		0.0056	0.104	0.0204
105	-0.0031	4	200	100	1,906		0.0063	0.105	0.0210
106	-0.0043	5	336	0	1,578		0.2834	0.096	0.2553
107	-0.0041	6	200	100	1,562		0.0091	0.093	0.0217
108	-0.0046	5	212	100	1,563		0.0091	0.091	0.0213
109	-0.0044	7	200	19	1,500		0.0091	0.092	0.0216
110	-0.0050	7	200	74	2,500		0.0042	0.200	0.0336
111	-0.0041	7	400	100	2,135		0.0014	0.165	0.0259
112	-0.0025	4	200	97	1,500		0.0056	0.103	0.0202
113	-0.0025	7	400	72	1,500		0.0014	0.202	0.0315
114	-0.0046	6	262	79	1,594		0.0063	0.107	0.0214
115	-0.0043	6	291	61	1,516		0.0049	0.109	0.0206
116	-0.0034	6	253	60	2,000		0.0021	0.139	0.0226
117	-0.0039	5	313	96	1,705		0.0042	0.124	0.0222
118	-0.0050	4	400	41	2,250		0.0028	0.138	0.0231
119	-0.0032	7	242	4	1,940		0.0021	0.134	0.0219
120	-0.0050	4	400	63	2,234		0.0028	0.136	0.0228

Table A.9: Evaluation Results for the First Forty Infill Points (80%/20%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos.	Rho	Risk Ratio	Nuisance Rate	Obj Value
41	-0.0025	7	200	6	2,469		0.0014	0.217	0.0446
42	-0.0050	4	306	0	2,437		0.5451	0.066	0.4493
43	-0.0032	7	200	0	1,648		0.3233	0.091	0.2769
44	-0.0025	7	400	5	1,500		0.0014	0.202	0.0415
45	-0.0050	4	400	10	2,500		0.0021	0.194	0.0406
46	-0.0025	4	400	91	2,500		0.0021	0.245	0.0507
47	-0.0050	4	200	5	2,125		0.0133	0.101	0.0309
48	-0.0050	4	200	30	1,500		0.0168	0.072	0.0277
49	-0.0025	4	200	40	2,500		0.0021	0.210	0.0437
50	-0.0050	4	200	86	1,500		0.0168	0.072	0.0277
51	-0.0050	7	200	32	1,500		0.0133	0.086	0.0279
52	-0.0025	7	201	66	2,500		0.0014	0.233	0.0477
53	-0.0025	7	400	58	2,500		0.0014	0.281	0.0574
54	-0.0050	7	398	45	1,500		0.0021	0.128	0.0272
55	-0.0025	7	400	49	1,500		0.0014	0.202	0.0415
56	-0.0050	7	200	16	1,500		0.0133	0.086	0.0279
57	-0.0050	4	200	55	1,500		0.0168	0.072	0.0277
58	-0.0050	4	200	8	1,862		0.0133	0.084	0.0274
59	-0.0044	4	300	63	2,266		0.0035	0.134	0.0296
60	-0.0050	7	399	21	1,500		0.0021	0.128	0.0272
61	-0.0050	4	200	61	1,500		0.0168	0.072	0.0277
62	-0.0034	5	324	63	1,633		0.0042	0.131	0.0296
63	-0.0025	7	400	81	1,500		0.0014	0.202	0.0414
64	-0.0050	7	292	50	2,483		0.0021	0.201	0.0418
65	-0.0046	4	313	84	2,250		0.0028	0.135	0.0293
66	-0.0025	4	201	35	1,500		0.0056	0.103	0.0250
67	-0.0045	5	339	24	1,739		0.0056	0.115	0.0275
68	-0.0050	7	400	94	1,500		0.0021	0.128	0.0272
69	-0.0036	4	205	99	1,743		0.0091	0.101	0.0274
70	-0.0047	4	200	98	1,500		0.0147	0.074	0.0266
71	-0.0050	4	200	100	1,500		0.0168	0.072	0.0277
72	-0.0047	4	220	74	1,586		0.0119	0.083	0.0262
73	-0.0031	6	262	81	1,524		0.0028	0.132	0.0285
74	-0.0035	4	200	5	1,500		0.0105	0.086	0.0257
75	-0.0038	4	200	5	2,500		0.0042	0.191	0.0416
76	-0.0044	4	353	49	1,719		0.0042	0.116	0.0266
77	-0.0040	4	213	41	1,875		0.0070	0.100	0.0255
78	-0.0044	7	331	74	1,659		0.0021	0.128	0.0273
79	-0.0025	4	400	70	1,500		0.0028	0.160	0.0343
80	-0.0031	6	239	75	2,012		0.0014	0.135	0.0281

Table A.10: Evaluation Results for the Second Forty Infill Points (80%/20%)

Data Pt	R Alert	Cycles	HCD	VCD	Slow Clos. Rho	Risk Ratio	Nuisance Rate	Obj Value
81	-0.0050	7	200	21	2,500	0.0042	0.200	0.0434
82	-0.0050	4	300	0	2,209	0.5437	0.060	0.4470
83	-0.0036	4	387	97	1,619	0.0035	0.129	0.0286
84	-0.0044	5	341	39	1,539	0.0049	0.111	0.0262
85	-0.0040	5	344	59	1,695	0.0042	0.122	0.0278
86	-0.0050	6	200	5	2,000	0.0084	0.102	0.0272
87	-0.0045	6	350	75	1,664	0.0035	0.120	0.0267
88	-0.0025	4	267	86	2,210	0.0035	0.158	0.0343
89	-0.0050	4	200	44	2,500	0.0056	0.178	0.0400
90	-0.0050	5	400	4	1,500	0.0063	0.113	0.0277
91	-0.0025	5	400	5	1,949	0.0014	0.186	0.0383
92	-0.0050	4	200	91	2,500	0.0056	0.178	0.0400
93	-0.0050	5	299	0	2,121	0.4976	0.069	0.4119
94	-0.0025	4	200	27	2,500	0.0021	0.210	0.0437
95	-0.0025	7	200	18	1,500	0.0028	0.124	0.0271
96	-0.0045	6	365	99	1,500	0.0021	0.130	0.0277
97	-0.0046	7	356	99	1,508	0.0021	0.136	0.0288
98	-0.0050	7	200	94	2,125	0.0077	0.116	0.0293
99	-0.0050	7	200	3	1,500	0.0133	0.086	0.0279
100	-0.0050	4	376	0	1,518	0.2834	0.087	0.2441
101	-0.0047	4	300	99	2,257	0.0042	0.129	0.0292
102	-0.0040	4	332	100	2,500	0.0014	0.202	0.0416
103	-0.0025	7	200	92	2,000	0.0021	0.140	0.0297
104	-0.0025	4	400	31	1,704	0.0035	0.163	0.0354
105	-0.0045	6	350	13	1,998	0.0014	0.132	0.0275
106	-0.0050	6	300	0	2,160	0.4878	0.076	0.4053
107	-0.0050	4	200	80	1,656	0.0147	0.076	0.0270
108	-0.0050	7	200	69	1,500	0.0133	0.086	0.0279
109	-0.0044	4	200	8	1,500	0.0154	0.077	0.0278
110	-0.0050	7	200	28	1,938	0.0084	0.101	0.0270
111	-0.0045	5	250	99	1,500	0.0091	0.091	0.0255
112	-0.0050	4	200	83	1,625	0.0147	0.075	0.0268
113	-0.0025	4	200	38	1,875	0.0042	0.114	0.0261
114	-0.0025	6	204	29	1,500	0.0014	0.134	0.0279
115	-0.0050	4	200	35	1,719	0.0140	0.078	0.0269
116	-0.0041	4	338	88	2,116	0.0028	0.128	0.0279
117	-0.0025	7	400	13	1,500	0.0014	0.202	0.0415
118	-0.0047	5	323	30	1,797	0.0063	0.113	0.0277
119	-0.0031	5	250	67	2,000	0.0028	0.127	0.0277
120	-0.0025	7	400	66	1,500	0.0014	0.202	0.0415




Figure A-1: Risk ratio (blue) and nuisance alert rate (red) at the optimal point found for each weighting combination

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

- [1] D. Asmar and M. Kochenderfer. Robust airborne collision avoidance through dynamic programming. Technical Report ATC-398, MIT Lincoln Laboratory, 2013.
- [2] B. Carpenter. A probability-based alerting logic for aircraft on parallel approach. Master’s thesis, Massachusetts Institute of Technology, 1996.
- [3] B. Carpenter and J. Kuchar. Probability-based collision alerting logic for closely-spaced parallel approach. In *AIAA Meeting Papers on Disc*, 1997.
- [4] J. Chryssanthacopoulos, M. Kochenderfer, and W. Olson. Assessment of the alert frequency of the traffic alert and collision avoidance system during closely spaced parallel operations. Technical Report 42PM-TCAS-0087, MIT Lincoln Laboratory, 2011.
- [5] R. Dawes and B. Corrigan. Linear models in decision making. *Psychological Bulletin*, 81(2):95–106, Feb 1974.
- [6] A. Duquette. Nextgen goal: Performance-based navigation. Fact Sheet, Mar 2010.
- [7] Y. Ebrahimi. Parallel runway requirement analysis study: Volume 1 - the analysis. NASA Contractor Report 191549 Contract NAS1-18027, Boeing Commercial Airplane Group, Seattle, Dec 1993.
- [8] W. Edwards and F. Barron. Smarts and smarter: Improved simple methods for multiattribute utility measurement. *Organizational Behavior and Human Decision Processes*, 60:306–325, 1994.
- [9] Federal Aviation Administration. *Introduction to TCAS II Version 7.1*, February 2011.
- [10] A. Forrester, A. Keane, and N. Bressloff. Design and analysis of noisy computer experiments. *AIAA Journal*, 44(10):2331–2339, Oct 2006.
- [11] A. Forrester, A. Sobester, and A. Keane. *Engineering Design via Surrogate Modelling: A Practical Guide*. American Institute of Aeronautics and Astronautics, 2008.

- [12] D. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, 1989.
- [13] P. Hennig and C. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13:1809–1837, 2012.
- [14] ICAO 9274. *Manual On the Use of the Collision Risk Model (CRM) for ILS Operations*, Jan 1997.
- [15] R. Iman, J. Helton, and J. Campbell. An approach to sensitivity analysis of computer models, i - introduction, input variable selection and preliminary variable assessment. *Journal of Quality Technology*, 13:174–183, July 1981.
- [16] International Civil Aviation Organization. *Required Navigation Performance Authorization Required (RNP AR) Procedure Design Manual*, 1st edition, 2009. Doc 9905.
- [17] Joint Planning and Development Office. *Concept of Operations for the Next Generation Air Transportation System*. Version 3.2.
- [18] D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.
- [19] S. Kirkpatrick, Jr. C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [20] G. Kjellstrom. On the efficiency of gaussian adaptation. *Journal of Optimization Theory and Applications*, 71(3):589–597, 1991.
- [21] M. Kochenderfer and J. Chryssanthacopoulos. Robust airborne collision avoidance through dynamic programming. Technical Report ATC-371, MIT Lincoln Laboratory, 2011.
- [22] D. Krige. A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Chemical, Metallurgical and Mining Society of South Africa*, 52(6):119–139, 1951.
- [23] J. Kuchar and A. Drumm. The traffic alert and collision avoidance system. *Lincoln Laboratory Journal*, 16:277–296, 2007.
- [24] J. Lee and N. Moray. Trust, control strategies and allocation of function in human-machine systems. *Ergonomics*, 35:1243–1270, 1992.
- [25] S. Massimini. Simultaneous independent and dependent parallel instrument approaches. Technical report, MITRE, 2006.
- [26] T. Meyer and J. Bradley. The evolution from area navigation (rnav), required navigation performance (rnp), to rnp rnav. Information Paper, Oct 2001. Presented by J. Williams at the ICAO GNSS IP11 Meeting.

- [27] B. Muir. Trust between humans and machines, and the design of decision aids. *International Journal of Man-Machine Studies*, 27:527–539, 1987.
- [28] W. Olson, J. Olszta, and L. Javits. Topa report 7: October 2010 - march 2011. Technical report, MIT Lincoln Laboratory, 2011.
- [29] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 3rd edition, 2007.
- [30] A. Pritchett, B. Carpenter, K. Asari, J. Kuchar, and R. Hansman. Issues in airborne systems for closely-spaced parallel runway operations. In *Digital Avionics Systems Conference*, 1995.
- [31] A. Pritchett and R. Hansman. Pilot non-conformance to alerting system commands during closely spaced parallel approaches. In *Digital Avionics Systems Conference*, volume 2, pages 9.1–1–9.1–8, 1997.
- [32] R. Regis and C. Shoemaker. Constrained global optimization of expensive black box functions using radial basis functions. *Journal of Global Optimization*, 31:153171, 2005.
- [33] E. Rovira, K. McGarry, and R. Parasuraman. Effects of imperfect automation on decision making in a simulated command and control task. *Human Factors*, 49:76–87, 2007.
- [34] RTCA. *Minimum Operational Performance Standards for Traffic Alert and Collision Avoidance System II (TCAS II) Version 7.1: Attachment A (Syntax of Pseudocode)*, June 2008.
- [35] J. Sacks, W. Welch, T. Mitchell, and H. Wynn. Design and analysis of computer experiments. *Statistical Science*, 4(4):409–423, Nov 1989.
- [36] M. Schonlau. *Computer Experiments and Global Optimization*. PhD thesis, University of Waterloo, 1997.
- [37] G. Slayton and C. Manberg. Patent us 2011/0282582 a1: Dynamic collision avoidance systems and methods, 2011.
- [38] W. Stillwell, D. Seaver, and W. Edwards. A comparison of weight approximation techniques in multiattribute utility decision making. *Organizational Behavior and Human Performance*, 28:62–77, 1981.
- [39] M. Toussaint. *The Bayesian Search Game*. Springer, 2012.
- [40] A. Vela. *Understanding Conflict-Resolution Taskload: Implementing Advisory Conflict-Detection and Resolution Algorithms in an Airspace*. PhD thesis, Georgia Institute of Technology, 2011.
- [41] A. Vela, K. Smith, and W. Olson. A closely spaced parallel operations encounter model: Version 1.2. Technical report, MIT Lincoln Laboratory, 2013.

- [42] D. Whitley. A genetic algorithm tutorial. *Statistics and Computing*, 4:65–85, 1994.
- [43] T. Yechout. *Introduction to Aircraft Flight Mechanics*. American Institute of Aeronautics and Astronautics, 2003.