ECAI 2025
I. Lynce et al. (Eds.)

doi:10.3233/FAIA251466

Supporting Stability Analysis of Aircraft Flight Control Laws

Luca Bonalumi[a,*], **Carmelo Raciti**[b] **and Francesco Amigoni**[a]

[a]Politecnico di Milano
[b]Leonardo SpA - Aeronautics Division

Abstract. Traditionally, flight control laws are verified, among other analyses, by demonstrating robust stability in the flight envelope (namely, in the region of operation of aircraft). Stability requirements are given in terms of gain and phase margins. These margins give an indication of the robustness of the control laws in presence of unmodeled effects and uncertainty. This paper proposes an AI-based approach to support the stability analysis tasks by identifying a structured partition of the flight envelope, where each region of the partition exhibits locally homogeneous stability characteristics. Building on recursive partitioning methods from Machine Learning, the proposed approach leverages the interpretability of tree-based models to facilitate human expert validation and usage of the obtained results.

1 Introduction

Modern high-performance aircraft are often designed to be naturally unstable, as this enables substantially enhanced maneuverability [14]. However, without a flight control system, the aircraft would not be able to maintain a stable and controlled flight. Through a fly-by-wire architecture, the pilot inputs are not directly linked to the aircraft's control surfaces. Instead, control algorithms take over certain aspects of handling in real time, sending commands to the control surfaces' actuators to compensate for the aircraft's inherent instability, to reduce pilot workload, and to optimize aircraft performance. In this context, the *flight control laws* define how pilot inputs are translated into deflections of the control surfaces, based on various factors, such as the context of flight and the aircraft configuration.

A critical step in the design process of aircraft – in terms of effort and expense – involves the *clearance* of the flight control laws. The clearance process aims at verifying that the control laws allow the aircraft to meet safety and handling requirements. The complexity of this process lies primarily in the breadth of flight conditions to be analyzed. An exhaustive validation must be carried out throughout the whole intended region of operation (called *flight envelope*). Interpreting the results of these activities requires significant expertise. It is the responsibility of human domain experts to delineate regions of the flight envelope where safe operation of the aircraft is ensured, and to identify areas where flight restrictions must be imposed to mitigate risks [7].

Given the scale and complexity of the clearance process, there is a growing motivation to explore methods that can assist human domain experts in analyzing results and making informed decisions. Artificial Intelligence (AI) techniques have been demonstrated successful in various engineering domains, particularly in tasks involving large-scale data analysis, pattern recognition, and decision support. However, to date, the application of AI methods to the flight control law clearance process remains largely unexplored. Because the clearance process ultimately leads to flight certification, the aim is not for AI to replace human expert judgment through a fully automated solution, but rather to support human expert decision-making. In this sense, any supporting tool must be not only effective but also interpretable. In a safety-critical domain such as aviation, the ability of human experts to understand, trust, and justify the output of an AI system is essential.

To understand where AI methods can provide the most value, it is useful to break down the clearance process into its constituent tasks, each aimed at verifying compliance with a selected collection of criteria. These tasks can be grouped into three categories: (i) stability analysis, which evaluates the behavior of the closed-loop system; (ii) handling qualities, which assess whether the aircraft responds to pilot inputs in an acceptable manner; and (iii) manned simulations, which involve pilot-in-the-loop testing to confirm pilot acceptance.

This paper proposes an AI-based approach to support the stability analysis task, with the specific objective of identifying regions of the flight envelope where the aircraft exhibits coherent behavior in terms of stability robustness. Traditionally, robustness is assessed in terms of gain and phase margins. The clearance criterion is expressed as a minimum threshold for the gain and phase margin. A common strategy to ensure exhaustive coverage of the flight envelope consists of evaluating the system's response on a grid of flight points. The resulting values are then analyzed by human domain experts, who look for emerging patterns or similarities in the robustness metrics. The final output is a partition of the flight envelope space into homogeneous regions, each associated with a qualitative summary of its stability robustness.

Specifically, here we propose a tree-based space partitioning method able to assist human domain experts in the identification of the regions in the flight envelope having recognizably distinct robustness trends. Given the critical importance of capturing local worst cases in the stability indicators, we introduce a tree construction approach that guarantees no overestimation of the robustness metrics. In the absence of a benchmark dataset, the validity of our results is assessed through close collaboration with human domain experts, who provide supervision and judgment based on their operational knowledge and experience. The application of this method to two synthetic datasets – designed in collaboration with domain experts – suggests that it has the potential to provide significant time savings

* Corresponding Author. Email: luca.bonalumi@polimi.it.

in the manual analyses typically involved in the clearance process, by automatically highlighting coherent trends and areas of concert.

This paper is organized as follows. Section 2 surveys the relevant related work. Section 3 outlines the process of stability analysis in flight control laws clearance. Section 4 formulates the specific stability analysis problem we address, while Section 5 describes our tree-based solving approach. Section 6 discusses the experimental results. Finally, Section 7 concludes the paper.

2 Related Work

Partitioning or segmentation of spaces, like feature spaces, is used to divide a spatial domain into non-overlapping subsets that exhibit local similarity. The partitioning process is often hierarchical, applying recursive binary splits to divide a region of space iteratively. Decision tree methods are the simplest approaches for producing a recursive partitioning of the space. The Classification and Regression Tree (CART) algorithm [4] is widely used. In CART, each leaf node approximates the target output locally by using either the majority class – for a classification task – or the mean target value – for a regression task. At each step, the algorithm selects the binary split that maximizes the reduction in an impurity metric, such as the Gini index or mean squared error.

Tree-based methods offer unique properties that make them a natural fit for problems requiring interpretability [3], [8], [12]. On the one hand, they can model non-linear relationships between predictors and targets without requiring variable transformations. Their recursive partitioning captures complex patterns through the composition of simple local approximations, while preserving the correspondence to the original feature space. In addition, the hierarchical structure of the decision paths offers a clear insight into which features drive decisions at each node, particularly when using axis-aligned splits and for moderate depth of the tree structure.

Maintaining a compact tree structure is indeed crucial to interpretability, as excessive branching can overwhelm the human readers and obscure the underlying decision logic. Several works have stemmed from the observation that a piecewise constant fit of the data can produce excessive fragmentation in the partition structure, thus hindering interpretability. These studies explored the use of more complex parametric models at the tree nodes, with the idea that a more expressive approximator could help maintain a more compact tree structure. Decision trees of this kind are generally referred to as *model trees* or *hybrid trees* in the Machine Learning community, with M5 [15] being one of the earliest and most influential examples. More recently, a general framework for model-based recursive partitioning using model trees has been proposed [16], in which a parametric model is used to fit a region of data. At each node, the model parameters are tested for instability over a set of partitioning variables, and the model is split with respect to the variable associated with the highest instability.

Other studies have focused on the use of nonparametric models in the tree leaves, such as Gaussian Processes. The method of [9] partitions the space as in the CART algorithm to delineate regions that are assumed to follow a stationary Gaussian Process. This idea has been further developed in [11] using covariance approximations for a more efficient computation with large datasets. Both studies follow a Bayesian approach to tree construction and allow for an explicit estimation of the predictive uncertainty.

Bayesian methods for optimizing the tree structure [5] have been explored to address a major limitation of the CART algorithm – namely, its greedy construction approach that does not guarantee a globally optimal tree. Other works have addressed this limitation through alternative formulations. In [1], the problem of finding the decision tree with a globally optimal misclassification error is formulated as a mixed-integer optimization problem. With this method, an optimal decision tree can be found in a single optimization step, with a performance comparable to that of state-of-the-art ensemble methods such as random forests and gradient boosted trees. The advantage of this approach, with respect to ensemble methods, is that the interpretability is preserved, while the major drawback is the limited size of the datasets for which the optimal decision tree problem is practically solvable.

Beyond tree-based models, alternative partitioning strategies have been proposed based on geometric tessellations of the space, aimed at achieving local similarity. Successful examples can be found adopting Voronoi tessellations, particularly for the segmentation of geographical datasets. With this method, the space is divided into regions, where each region is associated with a center and composed of the points in the space that are closer to its center than to any other. In the approach proposed by [10], a spatial model for soil permeability is constructed using piecewise Gaussian Processes, with a Voronoi tiling defining the partition. Compared to tree-based methods, in which the region boundaries are defined by simple splitting rules, Voronoi tessellations create regions based on distance. As a result, it is generally harder for humans to explain and visualize the decision boundaries, especially in high-dimensional spaces.

Additionally, clustering methods such as k-means or DBSCAN [6] induce a partition in the input space that groups data points into regions of local similarity. In this setting, the problem is framed as unsupervised and does not directly account for a supervised target variable, requiring the target variable to be treated as an additional dimension of the feature space. Moreover, traditional clustering algorithms have not been designed with the explicit goal of producing an interpretable spatial partition, and typically offer little insight into the final cluster memberships. To address this, recent studies have explored ad hoc techniques for interpretable clustering, as well as ex-post techniques that try to explain the rationale behind the partitioning induced by traditional clustering methods. Many of these approaches leverage a tree structure to improve interpretability. For instance, the work of [2] presents an interpretable clustering approach based on decision trees. Similarly, [13] proposes the Iterative Mistake Minimization algorithm, which approximates a given k-means clustering by means of a decision tree with k leaves, each one representing a cluster.

In our work, we exploit a tree-based approach to preserve the interpretability of the space partitioning in the context of stability analysis in flight control laws clearance.

3 Background

In this section, we better outline the process of stability analysis that our system supports. For further details, please refer to [7].

The primary objective of stability analysis is to assess the robustness of the stability of the aircraft throughout the entire flight envelope. Stability margins can be computed from the open-loop frequency response of the linearized system. To ensure that all critical operating conditions are covered, a large number of flight points are selected using a fine grid spanning the entire flight envelope. These flight points are defined as coordinates along the dimensions of the physical flight parameters characterizing the flight envelope: Mach number (M), altitude, airspeed (CAS), angle of attack (AoA), and normal load factor (Nz). Specifically, the Mach number indicates the

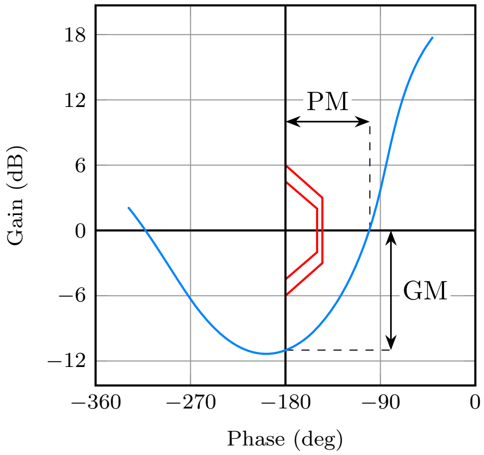

Figure 1. A Nichols diagram visualizing the open-loop frequency response (in blue) of a generic system. The gain and phase margins are denoted as GM and PM, respectively.

aircraft's speed relative to the speed of sound; airspeed refers to the velocity of airflow around the aircraft, which is critical for performance and handling; angle of attack is the angle between the aircraft's wing and the oncoming air; and the normal load factor measures the stress experienced by the aircraft due to acceleration, expressed in g's.

The open-loop frequency response is typically plotted on a *Nichols diagram*, where gain and phase margins can be evaluated. An example is shown in Figure 1: the requirements on the minimum acceptable gain and phase margins are represented as an exclusion region (in red) that must not be crossed by the frequency response curve. Two boundary red lines are visible in the diagram: the outer one defines the requirements that hold for the nominal model parameter values, while the inner one holds when tolerance thresholds are applied to the nominal values, to take uncertainty into account.

The task of the human domain expert is to identify the regions of flight points where the requirements are not met, as well as the local worst-case values for the stability margins. In case no violation of the requirements is detected, the aircraft configuration is cleared without limitations. Instead, in the case of a degradation of the robustness indicators below the acceptability threshold, a flight restriction may be issued for the affected region of the flight envelope. When issuing such limitations, human experts rely on a comprehensive view of the clearance results, seeking correlations across the evaluations and integrating the assessment with additional analyses when necessary. Ultimately, decisions regarding the severity and necessity of flight restrictions are guided by human expert judgment and accumulated experience.

A complete assessment of a single aircraft configuration typically involves evaluating thousands of flight points. To support a comprehensive understanding of the system's robustness across the entire flight envelope, experts also rely on visualizations that display the variation of a selected stability margin across the grid of flight points. These plots present one flight envelope dimension at a time on the x-axis, while the y-axis shows the corresponding stability margin at each grid point. By examining multiple plots – each one focusing on a different envelope dimension – experts can more easily identify patterns, trends, or potential regions of concern. An example of this visualization is shown in Figure 2, where the variation of the gain margin on the full flight envelope is plotted separately as a function of the Mach number, the airspeed, and the angle of attack.

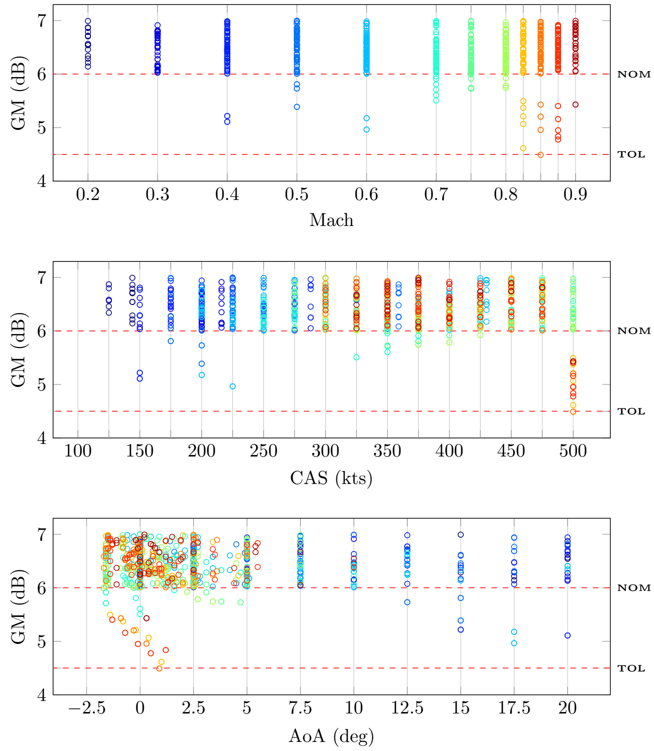

Figure 2. Visualization of the gain margin (GM) with respect to the Mach number, CAS, and AoA dimensions. The acceptable gain margin thresholds for the nominal conditions (NOM) and in the case of applied tolerances (TOL) are shown as red dashed lines. For visual reference, all flight points with the same Mach number are displayed using the same color in all the plots.

The final output of the stability analysis is a comprehensive textual report in which human domain experts aggregate the findings of their evaluations. The purpose of this report is to delineate distinct regions of the flight envelope where system robustness is observed to be homogeneous and to provide a reasoned judgment on the acceptability of the stability margins. In doing so, human experts effectively divide the flight envelope into non-overlapping regions based on local similarity in robustness indicators and identified trends. This human expert-driven partitioning forms a natural foundation for framing the task as a space partitioning problem, which we explore in the next section.

4 Problem Formulation

In this section, we provide the formal definition for the task of segmenting the flight envelope into locally homogeneous regions as a space partitioning problem.

Let $X \subseteq \mathbb{R}^d$ denote the flight envelope, defined as the domain of all admissible flight conditions. Each point $x \in X$ represents a unique combination of d physical flight parameters. In our case, $d = 4$: Mach number, altitude, airspeed, angle of attack, and normal load factor.

Let $g : X \to \mathbb{R}$ be a stability robustness indicator defined as a function of the flight envelope, such that $g(x)$ denotes the value of the robustness measure (e.g., gain margin, phase margin) at flight condition x. Typically, g is evaluated on a discrete set of N flight points sampled from X, resulting in the dataset $\mathcal{D} = \{\langle x_i, g(x_i)\rangle\}_{i=1}^N$.

This dataset reflects the computed stability margins across the grid used in the flight clearance activities.

A solution consists of a partition $\mathcal{R}$ of the flight envelope into non-overlapping regions that cover the entire domain X. Formally, a partition $\mathcal{R} = \{R_1, R_2, \ldots, R_k\}$ is such that $R_i \subseteq X$, $\bigcup_{i=1}^{k} R_i = X$, and $R_i \cap R_j = \emptyset$ for $i \neq j$. Additionally, in order to reflect the practices and expectations of human domain experts, we impose the following requirements on partitions:

1. *Hyperrectangular structure*: Each region R_i must be representable as an axis-aligned hyperrectangle in $\mathbb{R}^d$, i.e.,
$$R_i = [a_1^{(i)}, b_1^{(i)}] \times \cdots \times [a_d^{(i)}, b_d^{(i)}].$$
2. *Local homogeneity*: Within each region R_i, the robustness indicator g should exhibit limited variability; that is, the values of $g(x)$ for $x \in R_i$ should be consistent.
3. *Interpretability*: The partitioning must be intelligible and actionable for human domain experts. This includes a preference for coarse regions over highly fragmented ones.

The requirement for hyperrectangular partition regions reflects the practices commonly adopted in the clearance process. For the sake of readability of the individual regions, human domain experts typically define them as the Cartesian product of closed intervals along the dimensions of the flight envelope.

While this requirement is formally specified, the remaining two – local homogeneity and interpretability – are currently expressed only informally. To meaningfully contribute to the automatic identification of a possible solution, they require a more principled and operational formulation. In particular, we address local homogeneity by introducing a cost function $c(R)$ that quantifies the internal coherence of each region R as a function of the values $g(x)$ of its flight points $x \in R$. This allows us to pose the space partitioning task as an optimization problem in which the optimal solution is obtained by minimizing the total cost across all regions. In the next section, we propose a way to quantify the homogeneity of regions in the spirit of capturing the preferences of human domain experts.

The requirement of interpretability is harder to express formally. While the hyperrectangle constraint contributes towards this goal, excessively fragmented partitions can be difficult to interpret or justify, despite each individual region being simple in shape (and possibly coherent). As such, we defer the explicit treatment of interpretability to the next section.

5 Tree-Based Model

In this section, we present our approach for solving the problem of flight envelope partitioning in support of the flight clearance activities. We propose a tree-based recursive partitioning method, specifically designed to produce region splits that are both locally coherent and globally interpretable to domain experts.

As outlined in the previous section, the problem is posed as an optimization task, where the goal is to minimize a suitable cost function c that quantifies the degree of incoherence within the partitioned regions. The proposed method adopts a greedy, top-down construction approach, inspired by the classical CART algorithm. At each node, the flight envelope domain is split along the axis-aligned cut that yields the greatest reduction in incoherence, as measured by the designed cost function c.

In the remainder of this section, we first explore the formulation for the coherence cost function c, reflecting the types of trends and structures typically sought by flight control laws clearance human experts. We then address how interpretability can be promoted in our method, particularly with respect to the readability and compactness of the final partition.

5.1 Local Coherence of Regions

In the clearance process, human domain experts do not simply look for uniform values of the robustness measure g within a region, but for recognizable patterns that well describe the local trends of the system stability indicators. Standard statistical measures of dispersion – such as variance or standard deviation – are insufficient to capture such structured behavior, as they penalize both noise and meaningful trends indiscriminately. Instead, we define coherence in terms of how well a region's robustness values $g(x)$ can be approximated by a simple, interpretable model – specifically, a local linear fit. This aligns with the visual analysis practices of experts, who often judge the system robustness on plots in which flight envelope dimensions are displayed separately (see Figure 2).

In practice, we propose a linear regression model to fit the values $g(x)$ of a region R of flight points, and define the coherence measure $c(R)$ of that region as the Mean Squared Error (MSE) of the flight points in R with respect to their best-fitting local linear model. Notably, the MSE serves a dual role in our approach: both as the loss function used to determine the best-fitting local model and as the coherence metric c used to evaluate the quality of a region. This formulation enables a recursive, greedy construction of the partitioning tree that is equivalent to fitting a model tree with linear regression models at each node, with minor differences.

In detail, the tree is built recursively using the following procedure:

1. Fit a linear regression model to approximate the values $g(x)$ for the flight points x in the current region R (initially, $R = X$) by minimizing the MSE.
2. For each possible candidate split of R (according to a grid of values for each dimension; remember that the splits are axis-aligned), divide the data into two subsets R' and R'' and fit a linear regression model to each.
3. Choose the split R_a, R_b that minimizes the maximum of the coherence costs of the two resulting subsets:
$R_a, R_b \leftarrow \arg\min_{R', R''} \max\{c(R'), c(R'')\}$
(i.e., that minimizes the maximum MSE over the corresponding R' and R'').
4. If the reduction of c (namely, $|c(R) - \max\{c(R_a), c(R_b)\}|$) is above a given threshold, then split the region R into two subregions R_a and R_b.

The main difference with traditional model tree regression algorithms, such as M5, is the evaluation of the greedy-best split. In fact, while model tree regression often compares the impurity of the parent node with the sum of the impurities of the children nodes, here we choose to compare the parent node impurity with the maximum value of impurity between children nodes. This favors splits in which both subregions are individually coherent and, as a consequence, ensures the highest aggregated impurity reduction.

In addition, in the context of flight control clearance, particular emphasis is placed on identifying worst-case behaviors, which can be critical for subsequent decisions. For this reason, a classical ordinary least squares (OLS) regression may not be fully appropriate, as it minimizes symmetric residuals and may fail to acknowledge the

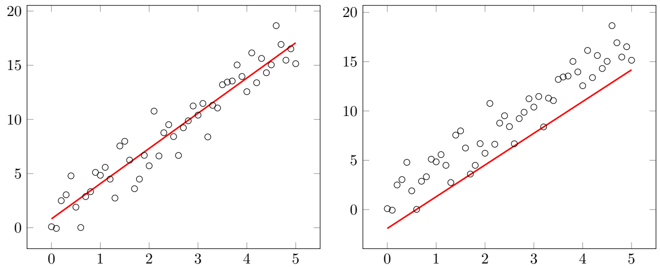

Figure 3. Comparison between the ordinary least-squares solution (on the left) and the constrained least-squares solution (on the right).

presence of outliers in case they are symmetrically scattered around the major linear trend of points. We therefore employ a modified regression approach that penalizes the underestimation of the robustness metric g. Specifically, we use a constrained least squares formulation that enforces the residuals to be non-negative, thereby ensuring the model surface does not overestimate the robustness indicator. Figure 3 compares the models fitted with ordinary least squares and with our constrained least squares formulation that enforces non-negativity of the residuals. This modification yields a more conservative and domain-aligned interpretation of the worst-case behavior while preserving the mathematical structure of the MSE-based coherence measure.

5.2 *Interpretability and Structural Constraints*

Interpretability plays a central role in determining the utility of a generated partition. In our framework, interpretability is encouraged through three additional optional strategies:

1. *Regularization of regression models*. We apply Lasso and Ridge regularization to the linear models at each leaf, promoting simplicity in the fitted surfaces. Lasso, in particular, has the added benefit of performing feature selection, often yielding models that depend on a small subset of the input dimensions. This not only reduces overfitting, but also enhances the readability of each region's behavior.
2. *Single-variable regression fits*. To further improve clarity, we introduce an optional constraint that restricts each region's model to depend on only one input variable. This has a direct benefit: the resulting fit can be immediately interpreted within the context of the standard visualizations used in the clearance process, which plot robustness indicators against one envelope dimension at a time (e.g., Mach number or angle of attack, see Figure 2).
3. *Tree pruning and structure control*. While the coherence objective encourages fine-grained splitting, this may result in excessive fragmentation – undermining interpretability. To counteract this, we enforce limits on the maximum tree depth and minimum impurity gain required for a split. These structural constraints act as a regularization mechanism, ensuring that the resulting partition remains intelligible and aligned with the granularity used by human analysts.

6 Experimental Results

In this section, we evaluate the proposed tree-based partitioning method using two synthetic datasets developed in collaboration with human domain experts. This choice was motivated, in part, by the absence of publicly available annotated datasets for flight envelope

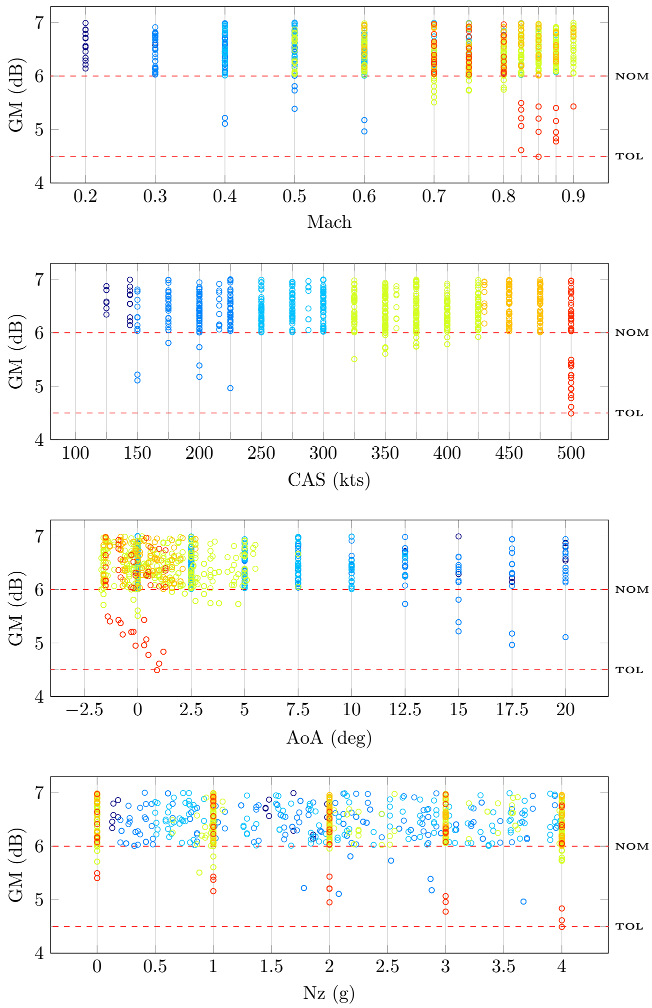

Figure 4. Visualization of the identified partition regions for the flight points in the first dataset. Points in the same region are plotted with the same color.

partitioning, and the high cost and complexity of conducting full-scale clearance campaigns for early-stage validation. The datasets are designed to reflect typical trends and challenges encountered during real-world stability analysis tasks. The results are assessed qualitatively by human experts based on the clarity and relevance of the identified partitions.

Each dataset consists of approximately 1,000 flight points, sampled over a non-uniform grid in a four-dimensional flight envelope defined by Mach number, airspeed, angle of attack, and normal load factor. The robustness metric analyzed is the gain margin.

6.1 *Second Dataset*

The second dataset exhibits a more widespread and severe deterioration of the robustness indicator:

1. A progressively worsening trend with increasing angle of attack, representing a key driver of margin degradation.
2. A lighter deterioration in gain margin at high airspeeds.

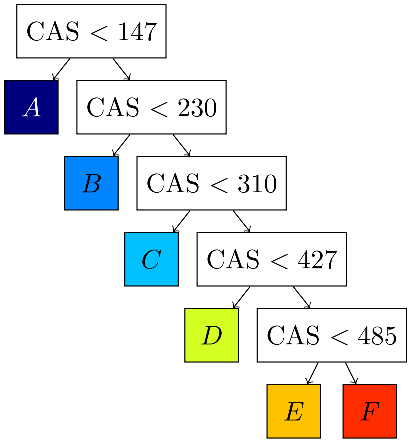

Figure 5. Decision tree defining the partition structure for the first dataset. Regions are denoted with the letters A to F, whose colors correspond to those in Figure 4.

6.2 First Dataset

The first dataset features three distinct regions of degraded robustness:

1. A mild degradation of approximately 1 dB below the nominal threshold, occurring at low airspeeds and Mach number in the interval $[0.4, 0.6]$.
2. A lighter degradation at medium airspeed, centered around Mach number 0.7 and 0.8.
3. A severe degradation for Mach number greater than 0.8, particularly at high airspeeds.

The resulting partition analyzed here is produced using the constrained least squares approach, combined with single-variable regression fits at each tree node. Tree expansion continues as long as the impurity reduction exceeds a threshold of 100.

The partition is visualized in Figure 4, where each subplot displays the gain margin as a function of a single flight envelope variable. Points are colored according to their assigned partition region, enabling experts to assess how well the model captures robustness trends and whether region boundaries align with their expectations. The underlying decision tree structure is shown in Figure 5.

The analysis of the generated partition reveals that all splits occur along the airspeed (CAS) dimension. This is considered a positive result, as the most prominent trends in the gain margin are indeed observed along CAS. Moreover, the resulting decision tree structure is very simple and intuitively clear to a human expert. Notably, for this dataset, alternative model fitting choices and regularization settings discussed in the previous section do not lead to substantial differences in the outcome.

The resulting partition that we analyze here is produced using the constrained least squares approach, in conjunction with the single-variable regression fits. The expansion of the tree is stopped if the impurity reduction is less than 100.

The partition is visualized in Figure 6, and the underlying decision tree is shown in Figure 7. Notice that the tree structure that originates in this setting is actually more fragmented than the one we show in Figure 7. Indeed, the original tree, shown in Figure 8, is pruned in order to reduce the number of regions to 4.

Overall, the generated partition shows a more complex structure than that of the previous dataset. The severe gain margin reduction happening at low airspeeds is captured in the partition (region B),

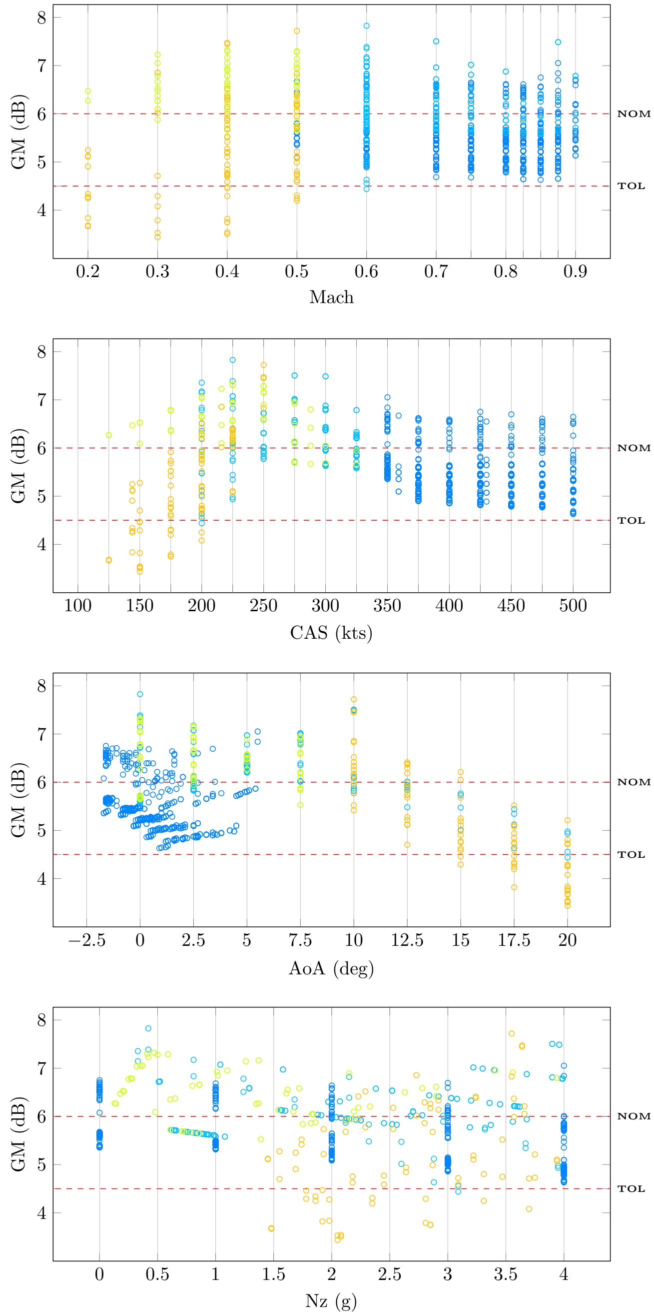

Figure 6. Visualization of the identified partition regions for the flight points in the second dataset. Points in the same region are plotted with the same color.

as well as the lighter deterioration at higher airspeeds (region D). In the end, pruning the initial decision tree structure was beneficial for improving the explainability of the solution, while maintaining a clear identification of the robustness trends.

7 Conclusion

In this work, we introduced a tree-based recursive space partitioning method tailored for the analysis of stability margins in flight control laws clearance. Our approach is driven by the dual objective of maintaining local coherence – interpreted in terms of visually and

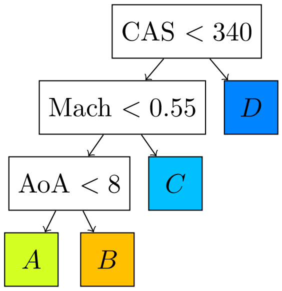

Figure 7. Pruned decision tree defining the partition structure for the second dataset. Regions are denoted with the letters A to D, whose colors correspond to those in Figure 6.

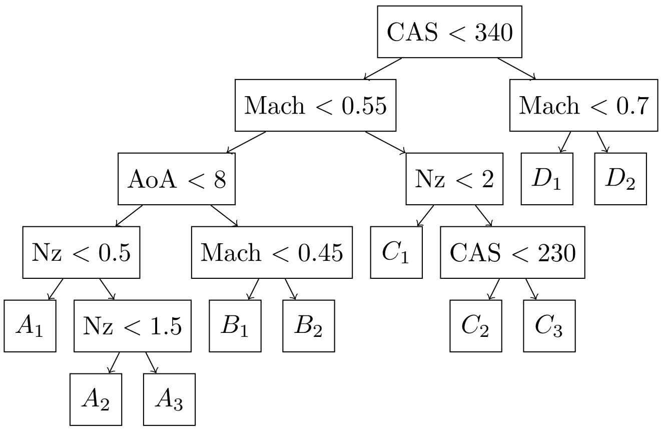

Figure 8. Decision tree defining the partition structure for the second dataset. Regions are denoted with A_i, B_i, C_i, and D_i.

structurally meaningful trends in the robustness metrics – and ensuring global interpretability of the resulting partition structure. We proposed a novel formulation of the coherence cost function, based on constrained linear regression models that emphasize local worst-case behavior. We further addressed the interpretability of the solution by adopting strategies such as single-variable regression fits and regularization, which promote clarity and alignment with domain-specific practices. The method was tested on two synthetic datasets, carefully designed with the support of human domain experts to reflect realistic degradation patterns in the gain margin. The results, evaluated qualitatively by the experts, demonstrate the capability of our method to capture critical trends and produce interpretable region boundaries within the flight envelope.

Future work will extend the breadth of this approach to other tasks of the flight clearance activities, attempting to frame the clearance problem on its full scale. Moreover, expanding the validation to real-world datasets remains a key step toward operational deployment.

Acknowledgements

We gratefully acknowledge the Flight Mechanics and Control Laws team at Leonardo SpA - Aeronautics Division for their invaluable expertise and continuous support throughout this work. This project was carried out under their supervision and was made possible thanks to the support and insights of their specialists.

References

[1] D. Bertsimas and J. Dunn. Optimal classification trees. *Machine Learning*, 106:1039–1082, 2017.

[2] D. Bertsimas, A. Orfanoudaki, and H. Wiberg. Interpretable clustering: An optimization approach. *Machine Learning*, 110(1):89–138, 2021.

[3] L. Breiman. Statistical modeling: The two cultures. *Statistical science*, 16(3):199–231, 2001.

[4] L. Breiman, J. Friedman, R. Olshen, and C. Stone. Classification and regression trees. *Biometrics*, 40:874, 1984.

[5] H. Chipman, E. George, and R. McCulloch. Bayesian cart model search. *Journal of the American Statistical Association*, 93(443):935–948, 1998.

[6] M. Ester et al. A density-based algorithm for discovering clusters in large spatial databases with noise. *Knowledge Discovery and Data Mining*, 96(34):226–231, 1996.

[7] C. Fielding, A. Varga, S. Bennani, and M. Selier. *Advanced techniques for clearance of flight control laws*, volume 283. Springer, 2003.

[8] A. Freitas. Comprehensible classification models: a position paper. *ACM SIGKDD Explorations Newsletter*, 15(1):1–10, 2014.

[9] R. Gramacy and H. Lee. Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*, 103(483):1119–1130, 2008.

[10] H.-M. Kim, B. Mallick, and C. Holmes. Analyzing nonstationary spatial data using piecewise Gaussian processes. *Journal of the American Statistical Association*, 100(470):653–668, 2005.

[11] B. Konomi, H. Sang, and B. Mallick. Adaptive Bayesian nonstationary modeling for large spatial datasets using covariance approximations. *Journal of Computational and Graphical Statistics*, 23(3):802–829, 2014.

[12] C. Molnar. *Interpretable Machine Learning*. 3rd edition, 2025. URL https://christophm.github.io/interpretable-ml-book.

[13] M. Moshkovitz, S. Dasgupta, C. Rashtchian, and N. Frost. Explainable k-means and k-medians clustering. In *Proceedings of the 37th International Conference on Machine Learning*, pages 7055–7065, 2020.

[14] W. Pinsker. The flying qualities of aircraft with actively augmented pitch and yaw stability. Technical report, Royal Aircraft Establishment, 1979.

[15] J. Quinlan. Learning with continuous classes. In *Proceedings of the 5th Australian Joint Conference on Artificial Intelligence*, pages 343–348, 1992.

[16] A. Zeileis, T. Hothorn, and K. Hornik. Model-based recursive partitioning. *Journal of Computational and Graphical Statistics*, 17(2):492–514, 2008.