

Mission profile clustering for usage-based health modeling of flight control actuators applied to a fleet of advanced jet trainers

Original

Mission profile clustering for usage-based health modeling of flight control actuators applied to a fleet of advanced jet trainers / Baldo, L., De Martin, A., Terner, M., Jacazio, G., Orchard, M.E., Sorli, M.. - ELETTRONICO. - 17 (1):(2025). (17th Annual Conference of the Prognostics and Health Management Society Bellevue (USA) October 27 – 30, 2025) [10.36001/phmconf.2025.v17i1.4545].

Availability:

This version is available at: 11583/3005159 since: 2025-11-24T10:23:15Z

Publisher:

ANNUAL CONFERENCE OF THE PROGNOSTICS AND HEALTH MANAGEMENT SOCIETY

Published

DOI:10.36001/phmconf.2025.v17i1.4545

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Mission Profile Clustering for Usage-Based Health Modeling of Flight Control Actuators Applied to a Fleet of Advanced Jet Trainers

Leonardo Baldo¹, Andrea De Martin², Mathieu Turner³, Giovanni Jacazio⁴, Marcos E. Orchard⁵, and Massimo Sorli⁶

^{1,2,4,6} *Department of Mechanical and Aerospace Engineering,
Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino, Italy*

leonardo.baldo@polito.it

andrea.demartin@polito.it

giovanni.jacazio@formerfaculty.polito.it

massimo.sorli@polito.it

³ *Leonardo S.p.A. Aircraft Division, Strada per Malanthero 9, Caselle Torinese, Torino, Italy*

mathieu.turner@leonardo.com

⁵ *Department of Electrical Engineering,*

Faculty of Physical and Mathematical Sciences, University of Chile, Santiago, Chile

morchard@ing.uchile.cl

ABSTRACT

This work introduces a mission profile clustering pipeline aimed at supporting usage-based health modeling of electro hydraulic flight control actuators employed in a fleet of Advanced Jet Trainer (AJT) aircraft. The study is part of a broader, high-level, modular Prognostics and Health Management (PHM) framework developed to predict Unscheduled Removals (URs) of the AJT horizontal tail flight control actuators. Operating in an industrial setting, this PHM effort specifically addresses the challenge of extracting prognostic information from a legacy fleet already in service, leveraging existing operational data to improve asset availability. The overall project leverages an extensive real-world dataset that spans over ten years and more than 25000 flight hours accumulated by a fleet of as many as 20 aircraft. This paper specifically focuses on the Flight Clustering Module within the Data Processing layer of the PHM framework, which serves as a critical enabler for future feature projections. Through an in-depth analysis of the underlying principles and a detailed overview of the main system interfaces, this work proposes a practical solution for categorizing and classifying mission profiles while highlighting the challenges of working with real operational data. After a pre-processing pipeline, developed to standardize and align time-series flight data, the clean trends are then clustered via a Self-Organizing Map

(SOM). In this work, a systematic SOM hyperparameter tuning pipeline is also introduced. The tuning routine employs a grid search strategy to optimize the SOM hyperparameters by jointly evaluating the topographic error, the quantization error, and the percentage of grid utilization. The result of the application of the trained SOM on the dataset is a set of Clustered Mission Types (CMTs), each linked to specific statistical distributions of actuator usage increments. These clusters are integrated into the broader PHM framework to simulate future aircraft behavior and estimate component degradation. Placed in an operational industrial environment, this methodology effectively connects mission-specific usage patterns with predictive health modeling, improving the modeling ability of PHM systems, and laying the foundation for smarter usage-based maintenance planning in aviation operations.

1. INTRODUCTION

Prognostic and Health Management (PHM) strategies have been offering the possibility of monitoring equipment health status since their first introduction in the context of Joint Strike Fighter (JSF) development in the early 2000s (Smith, Schroeder, Navarro, & Haldeman, 1997; Hess & Fila, 2002). Since then, these methodologies have evolved from theoretical concepts to practical enablers of smarter and more efficient maintenance approaches in aerospace systems (Kordestani, Orchard, Khorasani, & Saif, 2023). By allowing a shift from reactive to predictive maintenance, PHM is actively contributing to extending component life, enhanc-

Leonardo Baldo et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ing safety, and improving mission availability across a wide range of diversified platforms. A PHM strategy is composed of a set of interconnected functional layers that integrate the functions of diagnosis, prognosis, and health management. If the diagnosis focuses on the analysis of the current time step, from a theoretical point of view, the prognosis projects a number of features in the future and compares their evolution with a deterministic or probabilistic threshold (Vachtsevanos, 2006; Acuña-Ureta, Orchard, & Wheeler, 2021; Kordestani, Saif, Orchard, Razavi-Far, & Khorasani, 2021). Finally, the health management function uses the information obtained in an operational perspective. A key distinction in implementation is whether the equipment is still in development or is already operational. When PHM is integrated during the design phase, there is full control over sensor selection, data acquisition, and system integration. Conversely, applying PHM to legacy systems often involves working with fixed hardware and pre-existing data protocols. In these cases, a significant challenge lies in utilizing historical data, commonly gathered for purposes such as Structural Health Monitoring (SHM), control, or diagnostics, to obtain meaningful insights, despite the data not being specifically structured for PHM (Piatti, Walker, Figueroa, & Underwood, 2021; Leao, Fitzgibbon, Puttini, & De Melo, 2008; Esperon-Míguez, Jennions, & John, 2015). This is a quite common scenario for an industry that wants to digitize its assets to include them in a PHM perspective, and, as such, represents a common problem which deserves attention. It is within this context that the Advanced Jet Trainer (AJT), shown in Figure 1, manufactured by Leonardo S.p.A., has been selected as a candidate platform for the development of a comprehensive PHM architecture, which is presented in Section 2 and explained in details in (Baldo, De Martin, Terner, Jacazio, & Sorli, 2025). As a starting point, the focus has been placed on the all-moving horizontal tail Electro Hydraulic Actuator (EHA), identified as a representative and critical component.

The aerospace industry has historically been the foundational sector for the development of PHM strategies and, as a result, a wide range of systems have been analyzed in a PHM perspective. In the context of the Flight Control System (FCS), some studies have approached the diagnosis and prognosis of flight control actuators focusing on specific damage modeling based on physics-based formulations. In current aircraft configurations, the state-of-the-art flight control actuators are represented by EHAs, which are hydraulically powered (Maré, 2017). Despite their ruggedness and resistance to disturbance, EHAs are subject to fatigue and wear that depend on the specific way an aircraft is flown. Static maintenance schedules fail to account for this variability, treating all flight hours as equal. This simplification may lead to either over-conservative replacements or unexpected flight No-Gos.

Several works have focused on actuator health modeling (Baldo, De Martin, Jacazio, & Sorli, 2025). Physics-based



Figure 1. Photo of the AJT considered for the PHM framework development.

models provide a detailed understanding of failure mechanisms, but are often difficult to scale or calibrate on large fleets (De Martin, Jacazio, & Sorli, 2022; Shahkar & Khorasani, 2022). On the other hand, data-driven approaches rely on patterns learned from operational data (Lu, Yuan, & Ma, 2018). These methods can be highly effective, provided that the data is properly structured and the variability in usage is adequately represented. Unfortunately, most studies in the literature address constrained scenarios in which data is obtained through simulation models or test benches, thus omitting the upstream challenges linked to data availability and representativeness in the first place (Schoenmakers, 2020).

This paper presents the details of one of the most important modules within the AJT PHM framework, the Flight Clustering Module (FCM), which is responsible for the clustering of flight profiles in order to provide a set of propagation parameters to enable feature propagation in the future.

2. THE PHM FRAMEWORK

As already stated in Section 1, one of the most critical differences that affects how PHM strategies are implemented is whether the equipment is still in the design phase or is an already operational asset. This fundamental difference has profound implications on how PHM can be effectively integrated and the strategies that can be used (Rodrigues, Yoneyama, & Nascimento Jr, 2012). For already operational equipment, PHM solutions must be developed around existing datasets, many of which were not originally intended for health monitoring but were collected for other purposes. A key challenge in these cases lies in the ability to infer, extract, and reinterpret operational insights from data streams that were never structured with PHM in mind (Lukens, Rousis, Baer, Lujan, & Smith, 2022). Against this background, the developed PHM system is distinguished by several key properties: data-driven, scalable, robust, grounded, and transferable.

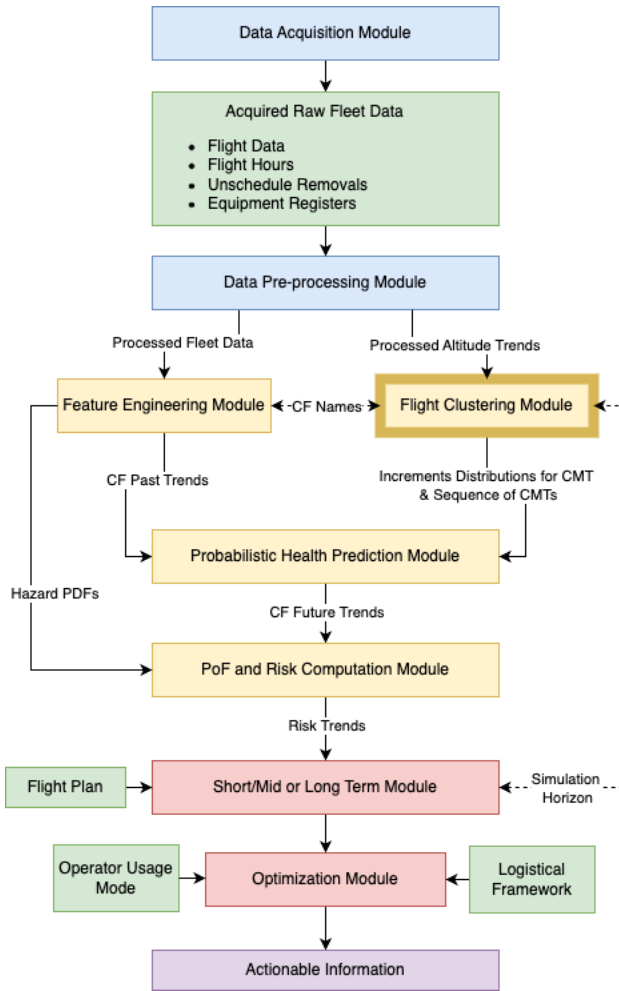


Figure 2. Flowchart of the PHM framework. Each color represents a functional layer: blue represents Layer 1 (Data Management, ETL & Pre-processing), yellow Layer 2 (Data Processing, Features & Prognosis), and red Layer 3 (Health Management, Strategies & Support). The location of the Flight Clustering Module is highlighted by thick edges. Green identifies data sources (internal or external), while purple highlights the framework output.

Furthermore, it utilizes aircraft-level data and is designed to be "as simple as possible and as complex as necessary". The goal has been to design and validate a Minimum Viable Product (MVP) around a case study that can serve as a scalable foundation for extending PHM capabilities to additional systems and subsystems across the aircraft. The case study, selected both for the consistent literature research gap and as a result of internal criticality analysis, is the AJT horizontal tail EHA, essential for the aircraft control (Baldo, 2024). Figure 2 reports the modular framework with a color code that highlights the functional layers of the framework. The key challenges of the project are related to the availability of data, which drastically limited the approaches that could have been followed. The AJT uses a Health Usage Management

System (HUMS) that was designed for SHM and, as a result, logs aircraft-level data in an exceedance event-trigger fashion (Kwakye, Jennions, & Ezhilarasu, 2024; Kappas & Frith, 2017). The result is a set of timeseries which do not include actuator control data and which are saved with a variable frequency. These challenges led to the choice of a more traditional strategy that employs the first four Statistical Moments (SM) (i.e., mean, variance, skewness, and kurtosis) (Zhenya, Lu, Ma, Yuan, & Chen, 2015; Soualhi et al., 2018; Baghli, Delpha, Diallo, & Hallouche, 2018). This more traditional approach enhances the explainability and scalability qualities of the proposed framework. Further details on the overall framework can be found in (Baldo, De Martin, et al., 2025).

After raw data are acquired by the Data Acquisition Module, each data source is pre-processed according to data-specific procedures. The acquired data consists of flight data time series, flight hours, Unschedule Removals (URs) information, and equipment registers. The Data Pre-processing Module leads to two different outputs: fleet data is sent to the Feature Engineering Module, while clustering data (which are univariate time series data) are the input for the Flight Clustering Module, where they are first treated with a customized additional pre-process pipeline, as explained later in Section 3.

In the Feature Engineering Module, the four SM lump each flight behavior into a compact set of parameters, enabling consistent performance monitoring even under low-quality or sparse data conditions. Over time, the moments accumulate into condition indexes (CIs), which represent the usage of the system (Wyłomańska & Zimroz, 2014). These CIs are ranked on the basis of their correlation with URs by analyzing the distribution of their magnitude increases between UR events. This ranking identifies the most relevant and informative Cumulative Features (CFs) and defines hazard Probability Distribution Functions (PDFs) for the later calculation of Probability of Failure (PoF) and risk assessment. Further details on the Feature Engineering Module can be found in (Baldo, De Martin, Terner, & Sorli, 2024). The FCM, which will be explained in detail in Section 3, includes three core functions: (i) processing altitude trends and applying a dedicated clustering algorithm to group flights with similar mission characteristics into Clustering Mission Types (CMTs), using flight path profiles as the discriminating feature. This classification serves as a foundation for both usage modeling and scenario forecasting. (ii) the module generates sequences of plausible future CMTs, either sampled from historical distributions or derived from mission-specific plans, enabling forward-looking simulations. (iii) it assigns characteristic feature increments (through a PDF) to each CMT, representing the expected contribution of that mission type to cumulative system degradation. All outputs from this module (future CMT sequences and associated feature increments) are passed to the Probabilistic Health Prediction Module, where they are com-

bined with features extracted from the Feature Engineering stage. This module employs Monte-Carlo simulations to forecast the evolution of key features over the predicted flight sequence. The resulting probabilistic CF future trends are then compared against pre-established hazard PDFs to estimate the PoF and generate actionable risk trends (Montoya, Valderama, Quintero, Pérez, & Orchard, 2020). The third layer is responsible for integrating risk trends into an existing industrial logistics and maintenance framework. The integration of risk-based outputs into operational planning remains subject to the specific needs and procedures of the end user and is still under definition. However, the ability to predict component health opens the door to both short-term and long-term decision support strategies (see Figure 5). In the short-term, predefined sequences of CMTs can be used to simulate the evolution of UR risk, providing insight into how specific mission sets impact the health of a given EHA installed on a particular aircraft. This enables mission-specific risk evaluation and supports data-driven fleet allocation strategies, where aircraft selection can be optimized to minimize the probability of UR for upcoming operations. From a long-term planning perspective, CMT sequences can be generated in an open-loop fashion by sampling historical mission distributions or applying Markov chain algorithms. This statistical approach enables the estimation of health degradation trends over extended horizons, supporting higher-level logistical processes such as supply chain optimization, maintenance planning, and spare parts provisioning. Together, these capabilities reinforce the strategic value of PHM systems by aligning health insights with real-world operational and logistical decisions.

3. THE FLIGHT CLUSTERING MODULE

To link actuator prognostics with aircraft use, it is important to accurately describe and quantify mission profiles (Kannemans & Jentink, 2002). An effective approach is to group similar flights into clusters and associate each cluster with a defined probabilistic pattern of expected actuator usage. Once this classification is available, it becomes possible to simulate damage evolution by assembling sequences of missions, whether for short-term planning or long-term projections. As a result, the clustering process becomes a central part of the broader PHM system, providing a connection between operational data and component-level health estimation. A breakdown of the FCM is reported in Figure 3, while a detailed view of the Custom Pre-processing block is reported in Figure 4.

As soon as the Clustering Data is inherited from the Pre-processing module, it is sent to the Custom Pre-processing block which performs additional pre-processing, specifically for clusterization, to the data. First, since clustering is performed on a carefully selected subset of the available dataset, comprising more than 3,000 flights from four aircraft, the relevant data subset is selected. This subset captures the diver-

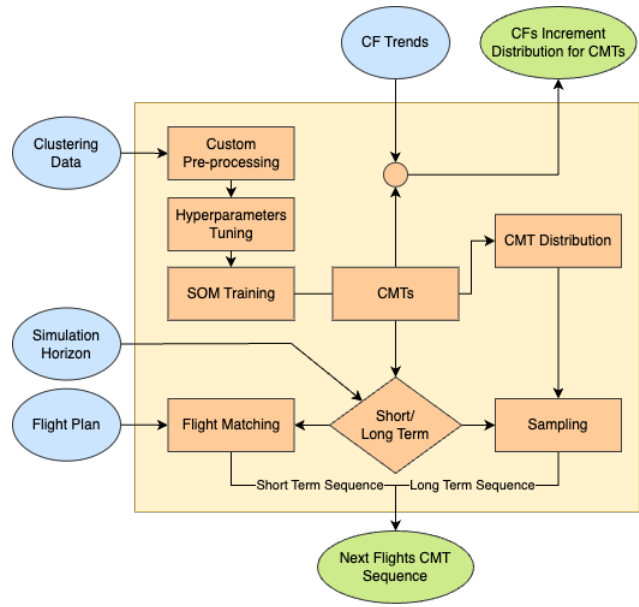


Figure 3. Flowchart of the FCM module. The blue circles represent module inputs whereas the green circles represent the module output.




Figure 4. Custom pre-processing flowchart: a detailed view of the "Custom Pre-processing" block of Figure 3.

sity of missions without overloading the clustering process with redundant patterns. Altitude was chosen as the key parameter for this analysis on the basis of three main considerations. First, variations in altitude closely reflect the behavior of the HT on the longitudinal axis, providing meaningful information on the dynamic performance of the aircraft. Although the AJT features an all-moving HT that combines the roles of both a stabilizer and an elevator, its effects are primarily observed in longitudinal flight mechanics. For this reason, altitude serves as a strong parameter for characterizing HT usage in relation to the broader motion of the aircraft. Second, altitude is a well-established way of representing mission profiles. It is commonly used in visual flight analyzes, allowing for consistent comparisons across different operational scenarios. Third, altitude trends are intuitive and easy to interpret, making them a practical choice for various stakeholders, including engineers, operators, and maintenance personnel, who rely on accessible data to understand and manage aircraft behavior.

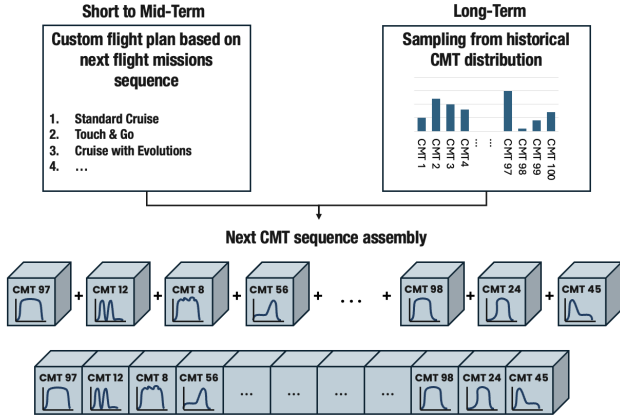


Figure 5. Next flight CMT sequence assembly procedure: short- to mid-term with flight plan information and long-term with historical sampling.

Each flight altitude time series has been initially resampled at a uniform rate of 1 Hz using linear interpolation. The resulting data have then been formatted into a custom Python time series structure. Following this, variance-based standardization has been applied to normalize the data across the dataset. Finally, each flight has been resampled again at a fixed length of 1000 data points, producing a consistent input shape of 1×1000 . This final resampling step is necessary to comply with the input requirements of the SOM algorithm, which requires a uniform time series length across samples. The choice of 1,000 points has been chosen considering the average flight duration (approximately 1 hour) and represents a balance between capturing sufficient temporal detail and maintaining computational efficiency. Once preprocessed, the altitude time series were passed into a SOM model (whose hyperparameters have been previously tuned) specifically adapted for time-series clustering to account for the visualization of temporal misalignments. The Python package MiniSom was employed in this project (Vettigli, 2018).

Once the map is trained, each CMT is associated with the statistical behavior of actuator usage observed in its flights. This allows the model to predict damage increments by associating each new flight with a CMT and retrieving the corresponding distribution. The Simulation Horizon is taken as an input and triggers two different paths. On the one hand, the simulation can predict a specific sequence of flight based on a flight matching procedure that compares the flight plan for the next flights with the set of CMTs identifying the most similar ones. On the other hand, in the long-term prediction, the model samples from the CMT distribution to generate future scenarios (see Figure 5). From a data perspective, flight data can be categorized as a multivariate time series dataset whose altitude trend represents a univariate time series. From the task point of view, given the lack of labels identifying the mission performed during flights, unsupervised clustering remains the

only viable option. The review of time series clustering methods reveals a limited range of options currently available in the literature (Aghabozorgi, Seyed Shirkorshidi, & Ying Wah, 2015). Among the most frequently used methods are time series k-means (Huang et al., 2016; Tang, Gu, Shen, & Chen, 2015), k-shape (Paparrizos & Gravano, 2015a), and Self-Organizing Maps (SOMs) (Javed, Rizzo, Lee, & Gramling, 2024). Additionally deep self-supervised encoders learn task-agnostic embeddings that can be clustered downstream. In the context of the presented PHM setting, which involves univariate, unlabeled data prioritizing interpretability, efficiency, and deployability, a SOM offers an optimal solution. SOM learns a topology-preserving two-dimensional grid, transforming the latent space into an interpretable map of operational clusters. Clusters naturally form as coherent regions in the U-matrix with physically meaningful prototypes that can be examined by engineers. This approach requires no extensive time alignment and, unlike deep learning architectures that often involve complex hyperparameter tuning and data augmentation, SOMs are computationally efficient, scale effectively to large datasets, and can be trained rapidly on standard hardware. Additionally, SOMs support incremental updates when new data become available, making them particularly suitable for fleet monitoring applications. Let us now consider n as the number of time series (≈ 3000), m as the length of each series (1000), k as the number of clusters (if known), G the number of SOM prototypes and E the number of training epochs. From a computational standpoint, k-Shape requires per-iteration complexity $O(\max\{nkm \log m, nm^2, km^3\})$ (Paparrizos & Gravano, 2015b). Soft-DTW k-means incurs quadratic cost in the sequence length m , yielding $O(nkm^2)$ per iteration plus barycenter updates (Cuturi & Blondel, 2017). In contrast, a SOM compares each series with G prototypes and updates neighbors, training in $O(nGmE)$ and inferring in $O(Gm)$ per series, without building full $n \times n$ distance matrices. For instance, when tested over a small subset (250) of the flight databases, the SOM run in 0.24 % of the k-shape running time and comparable time (around 4s) with respect to K-means. The execution of the full dataset or other benchmarking methodologies was not feasible, as the code run extended over several days, without converging. In the next section the development of the SOM model is analyzed in depth.

4. SOM DEVELOPMENT

SOMs are relatively simple neural networks made up of two main layers: an input layer and an output layer composed of interconnected units (Kohonen, 1982). The output layer is usually arranged in a two-dimensional grid (rectangular or hexagonal), enabling an intuitive visualization of the results. In practice, each input (in this case, a flight) is assigned to the grid position that best matches its features, known as the Best Matching Unit (BMU). Once the SOM hyperparameters

are set, the network is trained using the selected subset of altitude trends from the flight dataset. Due to its architecture, the SOM establishes a direct correspondence between each map cell and a specific cluster.

As with all Machine Learning (ML) methods, SOM models are directly influenced by a set of hyperparameters that affect model performance. Despite the popularity of application leveraging SOM high dimensionality reduction capabilities, the applications involving time series analysis are definitely scarcer. As a result, a hyperparameter tuning pipeline had to be designed. The first part of the tuning focused primarily on two key parameters: the total number of clusters and the aspect ratio of the SOM grid. Although common heuristics suggest estimating the number of clusters as $5 \cdot \sqrt{\text{Number of Inputs}}$, this formulation produced an impractical high number for operational use (216). We set the number of clusters to 100 after trial-and-error tuning to balance interpretability and resolution, incorporating feedback from operations personnel. The map aspect ratio, defined as the proportion between the number of neurons along the horizontal and vertical dimensions, was optimized by evaluating the topographic product, a metric designed to assess the preservation of topological relationships during mapping. This metric has been adopted in various SOM optimization studies to guide the selection of a suitable grid geometry (Forest, Lebbah, Azzag, & Lacaille, 2020).

Figure 6 reports the topographic product trend with varying aspect ratio, highlighting the best aspect ratio of 1 (square map of 100x100 neurons). To ensure statistical consistency and mitigate the effects of random initialization, each SOM configuration with a fixed aspect ratio was independently trained 10 times. The next step has involved the selection of the two main hyperparameters: sigma and the learning rate. In particular, an end-to-end hyperparameter tuning pipeline based on grid search is proposed as reported in the pseudo code 1, exploring a total of 504 parameter combinations. The scoring \mathbf{S} is defined in Eq. 1. After the selection of these last parameters, the network topology is fixed and the SOM architecture is ready to be trained and employed in the FCM. In fact, as already stated, SOMs operate in two distinct phases: training and mapping. In the training phase, the input dataset is used to build a lower-dimensional representation.

$$\mathbf{S} = QE + TE + 2 \cdot (1 - \mathbf{U}_{\text{rate}}) \quad (1)$$

where QE is the quantization error, TE is the topographical error, and \mathbf{U}_{rate} represents the utilization rate (i.e., the percentage of neurons that are the BMU for at least one flight). The tuning operation has been performed with a fixed seed to safeguard reproducibility. The solution space is reported in the tridimensional graph in Figure 7: the final values of sigma and learning rate are 1.3 and 0.4, respectively.

Algorithm 1 Enhanced SOM Hyperparameter Optimization

Require: Dataset \mathbf{X} , patience p , threshold τ

Ensure: Optimal SOM parameters θ^* , trained SOM \mathcal{M}^*

```

1: Split  $\mathbf{X} \rightarrow (\mathbf{X}_{\text{train}}, \mathbf{X}_{\text{val}})$ 
2: Initialize parameter grid  $\mathcal{G}$ 
3:  $\mathcal{S}^* \leftarrow \infty$ 
4: for  $(\sigma, \alpha, T) \in \mathcal{G}$  do
5:   Create SOM  $\mathcal{M}$  with parameters  $(\sigma, \alpha)$ 
6:    $(T_{\text{actual}}, \text{conv}) \leftarrow \text{TRAINSON}(\mathcal{M}, \mathbf{X}_{\text{train}}, T, p, \tau)$ 
7:   Compute errors:  $QE_{\text{val}}, TE_{\text{val}}$ 
8:   Compute utilization rate:  $\mathcal{U}_{\text{rate}}$ 
9:    $\mathcal{S} \leftarrow QE_{\text{val}} + TE_{\text{val}} + 2(1 - \mathcal{U}_{\text{rate}})$ 
10:  if  $\mathcal{S} < \mathcal{S}^*$  then
11:     $\theta^* \leftarrow (\sigma, \alpha, T_{\text{actual}}, \text{conv})$ 
12:     $\mathcal{S}^* \leftarrow \mathcal{S}, \mathcal{M}^* \leftarrow \mathcal{M}$ 
13:  end if
14: end for
15: return  $\theta^*, \mathcal{M}^*$ 

16: function TRAINSON( $\mathcal{M}, \mathbf{X}, T_{\text{max}}, p, \tau$ )
17:    $QE_{\text{prev}} \leftarrow \infty, c \leftarrow 0$ 
18:   for  $t = 1$  to  $T_{\text{max}}$  do
19:     Train  $\mathcal{M}$  for 1 iteration
20:     if  $t \bmod 50 = 0$  then
21:        $QE_{\text{curr}} \leftarrow$  quantization error of  $\mathcal{M}$ 
22:       if  $|QE_{\text{prev}} - QE_{\text{curr}}| < \tau$  then
23:          $c \leftarrow c + 1$ 
24:         if  $c \geq p$  then return  $(t, \text{True})$ 
25:       end if
26:     else
27:        $c \leftarrow 0$ 
28:     end if
29:      $QE_{\text{prev}} \leftarrow QE_{\text{curr}}$ 
30:   end if
31: end for
32: return  $(T_{\text{max}}, \text{False})$ 
33: end function

```

During this phase, the SOM learns to organize the input data by adjusting the weights of its neurons based on their similarity to the input vectors. Each input is compared to all neurons on the map, and the one with the closest BMU, is identified. The BMU and its neighboring neurons are then updated to more closely resemble the input, gradually shaping the map to reflect the underlying structure of the dataset, in a procedure called competitive training. Throughout training, various classical SOM metrics are used to evaluate the quality of the clustering and guide the tuning process. The training lasts 100 epochs as already after 50 the errors stabilize and further epochs do not increase the accuracy of the model. Once training is complete, the mapping phase begins. In this phase, new input data can be projected onto the trained map. Each new input is assigned to its BMU, effectively

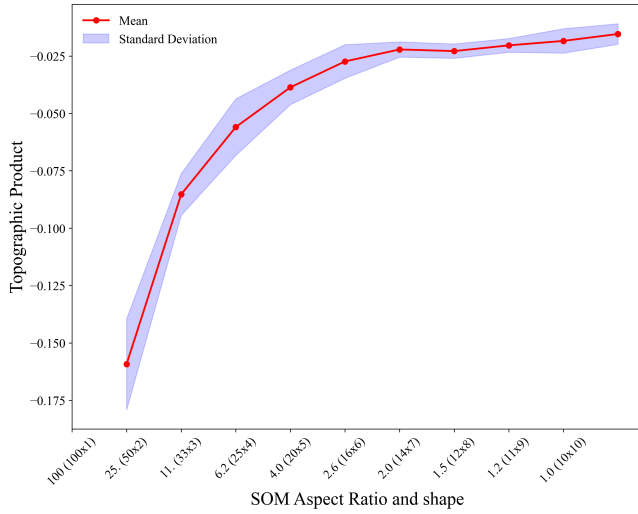


Figure 6. Topographic product trend in function of the aspect ratio along with the standard deviation band. In this instance, the SOM was trained considering the other parameters as default values.




Figure 7. SOM hyperparameters 3D space: quantization error, topographical error and utilization rate.

classifying it based on the structure learned during training. This process enables intuitive visualization and interpretation of complex, high-dimensional data by preserving topological relationships, data points that are similar in the input space remain close together in the map space.

5. RESULTS

The main advantage that made SOM popular is the extremely high explainability of the SOM grid, which places the identified similar neurons close to each other (Figure 8). A commonly used visualization method involves applying a color scale to the U-matrix. The U-matrix (Unified Distance Matrix) illustrates the Euclidean distances between neighboring neurons within the weight space and is reported in Figure 9. This approach effectively highlights cluster boundaries and topological relationships in the trained SOM by representing inter-neuron distances as a color-coded landscape. Darker areas indicate clusters of similar neurons, while lighter regions delineate the boundaries between clusters. This creates a landscape-like visualization in which valleys correspond to clusters of similar data points, and ridges or peaks indicate boundaries between different clusters. As a lead-in fighter training platform, the AJT operates within a structured syllabus of recurring mission types, yet individual flights exhibit substantial variability even within identical mission categories. A comprehensive analysis should incorporate both Figure 8 and Figure 9, as the information they present is interrelated. For instance, the U-matrix highlights a mega cluster of similar CMTs in the center of the grid; this indication is backed up by Figure 9 which tells us that that mega cluster represents common cruise flights. Other mega clusters can be identified: sustained high-altitude segments may indicate either basic navigation training, while descending altitude profiles typically correspond to instrument approach and landing practice. Conversely, low-altitude flight segments often represent air-to-ground training or tactical maneuvering exercises. Flights characterized by multiple altitude transitions generally reflect formation flying or complex tactical scenarios involving integrated mission elements or air-to-air combat scenarios. Modern pilot training emphasizes mission integration, where individual sorties frequently incorporate multiple training objectives, creating hybrid profiles that resist simple categorical classification. Finally, the similarities observed among certain CMTs should not be considered an error. Rather, they suggest that, when utilizing the 10-by-10 optimized network topology, these profiles are frequently occurring and involve multiple CMTs with only slight variations. SOMs offer intuitive and interpretable results, but assessing their performance is challenging, especially without labeled data for direct validation. To address this, various evaluation metrics are used for hyperparameter tuning and clustering assessment. These metrics are classified as internal, which use




Figure 8. CMT clusters with associated altitude trends. The gray lines represent the combined trends within each CMT, while the red line indicates the designated trend for that specific CMT.

only data characteristics, and external, which require label information.




Figure 9. U-matrix for the trained SOM: dark regions indicate similar clusters while lighter regions indicate non-similar clusters. The average number of flights per cluster is 32.5.

Common internal metrics include topographic error (0.033) (Kiviluoto, 1996), quantization error (11.827), combined error (Kaski & Lagus, 1996), silhouette index (0.047), Davies-

Bouldin index (2.390593), and topographic product (0.1023) (Bauer & Pawelzik, 1992). A quick implementation of these metrics is provided in a Python package developed by (Forest et al., 2020). The excellent topographic error and product indicate strong topology preservation, demonstrating that the SOM effectively maintains neighborhood relationships from the original high-dimensional space to the two-dimensional map. Conversely, the low silhouette and Davies-Bouldin indices suggest room for improvement, highlighting the challenges in identifying well-separated, compact clusters due to the complexity of the data.

6. CONCLUSIONS AND FURTHER DEVELOPMENTS

This study introduces a mission-aware approach to flight clustering for PHM applications. By analyzing and grouping altitude profiles from a subset of more than 3000 flights, a set of CMT that represents the most common operational behaviors in a fleet of AJTs was obtained. The use of SOMs allowed us to build an interpretable structure that connects real flight data to actuator usage. Each CMT has been associated with a distribution of actuator damage increments, enabling the modeling of health evolution over time. These results are integrated into an existing PHM framework, where they are used to simulate both short-term and long-term actuator degradation. The method supports various operational scenarios, from forecasting the impact of known future missions to generating random mission sequences for planning purposes. Moreover, new flights can be classified into existing CMTs, keeping the system responsive and up to date. The main limitations of the proposed approach are the need for a large amount of flight data and the assumption of recurring mission profiles, which may not always hold; however, Advanced Jet Trainers represent a peculiar case where recurring training missions make the framework applicable. The current research establishes a foundation for future analyses. For example, integrating multiple signals could potentially improve CMT identification accuracy: for instance the G-Force trend could provide additional stress information. However, the introduction of multivariate time series would necessitate the use of multidimensional SOMs or alternative machine learning architectures. In this context, the application of Graph Neural Networks (GNNs) for time series analysis has demonstrated promising results in existing literature and may further enhance the detection of distinctive patterns. Additionally, acquiring a labeled dataset of flight missions would facilitate the application of supervised machine learning techniques for classification purposes, thereby increasing the robustness and accuracy of the analysis.

ACKNOWLEDGMENT

This publication is part of the project PNRR-NGEU which has received funding from the MUR – DM 352/2022. This research is co-funded by Leonardo SpA.

REFERENCES

- Acuña-Ureta, D. E., Orchard, M. E., & Wheeler, P. (2021). Computation of time probability distributions for the occurrence of uncertain future events. *Mechanical Systems and Signal Processing*, *150*, 107332. doi: <https://doi.org/10.1016/j.ymssp.2020.107332>
- Aghabozorgi, S., Seyed Shirkhorshidi, A., & Ying Wah, T. (2015). Time-series clustering – a decade review. *Information Systems*, *53*, 16–38. doi: <https://doi.org/10.1016/j.is.2015.04.007>
- Baghli, M., Delpha, C., Diallo, D., & Hallouche, A. (2018). Three-Level Inverter Fault Detection and Diagnosis Using Current-Based Statistical Analysis. In *2018 Prognostics and System Health Management Conference (PHM-Chongqing)* (pp. 686–691). doi: 10.1109/PHM-Chongqing.2018.00123
- Baldo, L. (2024, October). Development of a data-driven condition-based maintenance methodology framework for an advanced jet trainer. In *PHM society european conference* (Vol. 8, pp. 5–5). doi: 10.36001/phme.2024.v8i1.3943
- Baldo, L., De Martin, A., Ternier, M., Jacazio, G., & Sorli, M. (2025). Implementing phm for legacy flight control actuators through operational aircraft data: Approach and lessons learned. *Results in Engineering*, *28*, 107214. doi: <https://doi.org/10.1016/j.rineng.2025.107214>
- Baldo, L., De Martin, A., Jacazio, G., & Sorli, M. (2025). A systematic literature review on phm strategies for (hydraulic) primary flight control actuation systems. *Actuators*, *14*(8). doi: 10.3390/act14080382
- Baldo, L., De Martin, A., Ternier, M., & Sorli, M. (2024). The journey towards condition-based maintenance: A framework for the horizontal tail actuator of an advanced jet trainer aircraft. In *34th congress of the international council of the aeronautical sciences (icas)*.
- Bauer, H.-U., & Pawelzik, K. (1992). Quantifying the neighborhood preservation of self-organizing feature maps. *IEEE Transactions on Neural Networks*, *3*(4), 570–579. doi: 10.1109/72.143371
- Cuturi, M., & Blondel, M. (2017). Soft-DTW: a differentiable loss function for time-series. In *Proceedings of the 34th international conference on machine learning* (Vol. 70, pp. 894–903). PMLR.
- De Martin, A., Jacazio, G., & Sorli, M. (2022, October). Performance Evaluation of Different PHM Strategies on the Performances of a Prognostic Framework for Electro-Hydraulic Actuators for Stability Control Augmentation Systems. In *Annual Conference of the PHM Society* (Vol. 14). doi: 10.36001/phm-conf.2022.v14i1.3289
- Esperon-Miguez, M., Jennions, I. K., & John, P. (2015). Implementing ivhm on legacy aircraft: Progress towards identifying an optimal combination of technologies. In P. W. Tse, J. Mathew, K. Wong, R. Lam, & C. Ko (Eds.), *Engineering asset management - systems, professional practices and certification* (pp. 799–812). Cham: Springer International Publishing.
- Forest, F., Lebbah, M., Azzag, H., & Lacaille, J. (2020). A survey and implementation of performance metrics for self-organized maps. *arXiv preprint arXiv:2011.05847*. Retrieved from <https://arxiv.org/abs/2011.05847>
- Hess, A., & Fila, L. (2002). The Joint Strike Fighter (JSF) PHM concept: Potential impact on aging aircraft problems. In *Proceedings, IEEE Aerospace Conference* (Vol. 6, pp. 6–6).
- Huang, X., Ye, Y., Xiong, L., Lau, R. Y., Jiang, N., & Wang, S. (2016). Time series k-means: A new k-means type smooth subspace clustering for time series data. *Information Sciences*, *367*, 1–13. doi: 10.1016/j.ins.2016.05.040
- Javed, A., Rizzo, D. M., Lee, B. S., & Gramling, R. (2024). Sometimes: self organizing maps for time series clustering and its application to serious illness conversations. *Data Mining and Knowledge Discovery*, *38*(3), 813–839. doi: 10.1007/s10618-023-00979-9
- Kannemans, H., & Jentink, H. W. (2002). A Method to Derive the Usage of Hydraulic Actuators From Flight Data. In *ICAS 2002 CONGRESS*.
- Kappas, J., & Frith, P. (2017). From hums to phm: Are we there yet. In *Proceedings of the 17th australian international aerospace congress, melbourne, vic, australia* (pp. 26–28).
- Kaski, S., & Lagus, K. (1996). Comparing self-organizing maps. In C. von der Malsburg, W. von Seelen, J. C. Vorbrüggen, & B. Sendhoff (Eds.), *International conference on artificial neural networks* (Vol. 1112, pp. 809–814). Berlin, Heidelberg: Springer.
- Kiviluoto, K. (1996). Topology preservation in self-organizing maps. In *Proceedings of international conference on neural networks (icnn)* (pp. 294–299).
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, *43*(1), 59–69.
- Kordestani, M., Orchard, M. E., Khorasani, K., & Saif, M. (2023). An Overview of the State of the Art in Aircraft Prognostic and Health Management Strategies. *IEEE Transactions on Instrumentation and Measurement*, *72*, 1–15. doi: 10.1109/TIM.2023.3236342
- Kordestani, M., Saif, M., Orchard, M. E., Razavi-Far, R., & Khorasani, K. (2021). Failure prognosis and applications—a survey of recent literature. *IEEE Transactions on Reliability*, *70*(2), 728–748. doi: 10.1109/TR.2019.2930195
- Kwakye, A. D., Jennions, I. K., & Ezhilarasu, C. M. (2024). Platform health management for aircraft maintenance – a review. *Proceedings of the Institution of Mechanical*

- Engineers, Part G: Journal of Aerospace Engineering*, 238(3), 267–283. doi: 10.1177/09544100231219736
- Leao, B. P., Fitzgibbon, K. T., Puttini, L. C., & De Melo, G. P. B. (2008, March). Cost-Benefit Analysis Methodology for PHM Applied to Legacy Commercial Aircraft. In *2008 IEEE Aerospace Conference* (pp. 1–13). Big Sky, MT, USA: IEEE. doi: 10.1109/AERO.2008.4526599
- Lu, C., Yuan, H., & Ma, J. (2018, December). Fault detection, diagnosis, and performance assessment scheme for multiple redundancy aileron actuator. *Mechanical Systems and Signal Processing*, 113, 199–221. doi: 10.1016/j.ymssp.2016.10.012
- Lukens, S., Rousis, D., Baer, T., Lujan, M., & Smith, M. (2022, October). A Data Quality Scorecard for Assessing the Suitability of Asset Condition Data for Prognostics Modeling. *Annual Conference of the PHM Society*, 14(1). doi: 10.36001/phmconf.2022.v14i1.3188
- Maré, J.-C. (2017). *Aerospace Actuators 2: Signal-by-Wire and Power-by-Wire*. John Wiley & Sons.
- Montoya, F. R. J., Valderrama, M., Quintero, V., Pérez, A., & Orchard, M. (2020). Time-of-failure probability mass function computation using the first-passage-time method applied to particle filter-based prognostics. In *Annual conference of the phm society* (pp. 11–11). doi: 10.36001/phmconf.2020.v12i1.1299
- Paparrizos, J., & Gravano, L. (2015a). k-shape: Efficient and accurate clustering of time series. In *Proceedings of the 2015 acm sigmod international conference on management of data* (p. 1855–1870). New York, NY, USA: Association for Computing Machinery. doi: 10.1145/2723372.2737793
- Paparrizos, J., & Gravano, L. (2015b). k-shape: Efficient and accurate clustering of time series. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data* (pp. 1855–1870).
- Piatti, F., Walker, M., Figueroa, F., & Underwood, L. (2021). A case study on the challenges and opportunities for the deployment of phm capabilities in existing engineering systems. In *2021 IEEE Aerospace Conference (50100)* (p. 1–8). doi: 10.1109/AERO50100.2021.9438426
- Rodrigues, L., Yoneyama, T., & Nascimento Jr, C. (2012). How aircraft operators can benefit from PHM techniques. In *Ieee aerospace conference proceedings*. doi: 10.1109/AERO.2012.6187376
- Schoenmakers, L. (2020). *Condition-based Maintenance for the RNLA F C-130H(-30) Hercules* (Master's Thesis). Eindhoven University of Technology.
- Shahkar, S., & Khorasani, K. (2022, May). A Multidimensional Bayesian Methodology for Diagnosis, Prognosis, and Health Monitoring of Electrohydraulic Servo Valves. *IEEE Transactions on Control Systems Technology*, 30(3), 931–943. doi: 10.1109/TCST.2021.3079198
- Smith, G., Schroeder, J., Navarro, S., & Haldeman, D. (1997). Development of a prognostics and health management capability for the Joint Strike Fighter. In *1997 IEEE Autotestcon Proceedings AUTOTESTCON '97. IEEE Systems Readiness Technology Conference* (pp. 676–682). doi: 10.1109/AUTEST.1997.643994
- Soualhi, A., Hawwari, Y., Medjaher, K., Clerc, G., Hubert, R., & Guillet, F. (2018). Phm survey: Implementation of signal processing methods for monitoring bearings and gearboxes. *International Journal of Prognostics and Health Management*, 9(2), 0.
- Tang, X., Gu, J., Shen, Z., & Chen, P. (2015). A flight profile clustering method combining twed with k-means algorithm for 4d trajectory prediction. In *2015 integrated communication, navigation and surveillance conference (icns)* (p. S3-1-S3-9). doi: 10.1109/ICNSURV.2015.7121260
- Vachtsevanos, G. (2006). *Intelligent fault diagnosis and prognosis for engineering systems*. John Wiley & Sons, Ltd. doi: 10.1002/9780470117842
- Vettigli, G. (2018). *Minisom: minimalistic and numpy-based implementation of the self organizing map*. Retrieved from <https://github.com/JustGlowing/minisom/>
- Wyłomańska, A., & Zimroz, R. (2014). Signal segmentation for operational regimes detection of heavy duty mining mobile machines—a statistical approach. *Diagnostyka*, 15(2), 33–42.
- Zhenya, W., Lu, C., Ma, J., Yuan, H., & Chen, Z. (2015). Novel method for performance degradation assessment and prediction of hydraulic servo system. *Scientia Iranica*, 22(4), 1604–1615.