AUTONOMOUS FAULT DIAGNOSIS: STATE OF THE ART AND AERONAUTICAL BENCHMARK

Julien Marzat[1,2], Hélène Piet-Lahanier[1], Frédéric Damongeot[1], Eric Walter[2]

1. ONERA, 29 av. de la Division Leclerc - BP 72, 92322 Châtillon, France. <FirstName>.<LastName>@onera.fr
2. Laboratoire des Signaux et Systèmes, CNRS-SUPELEC-Univ Paris-Sud, 91192 Gif-Sur-Yvette, France. <FirstName>.<LastName>@lss.supelec.fr

Abstract

This paper briefly reviews the state of the art in Fault Detection and Isolation (FDI), and describes a test-case for in-flight fault diagnosis. The main concepts are recalled, the links between the approaches are indicated and the most promising methods are highlighted. The nonlinear models of the system, its sensors and actuators are described in the second part of the article, along with fault scenarios. Constraints restrict the approaches that are applicable on the test-case and suggest quantitative indices that can be used to evaluate fault diagnosis strategies in an aeronautical context.

Index Terms

aircraft, analytical redundancy, fault detection and isolation, model-based FDI, nonlinear systems, pattern recognition

1. Introduction

Reliability and safety of aeronautical systems have always been a major concern for aircraft manufacturers. There is an absolute necessity to identify early unexpected changes in the system (referred as *faults*) before they lead to a complete breakdown (*failure*). *Fault Detection and Isolation* (FDI) and *Fault Tolerant Control* (FTC) for aerial and space vehicles have already been addressed in a large number of papers (see, e.g., [1][2]). FDI comprises *fault detection* ("something is wrong"), *fault isolation* ("the fault pertains to this component") and even *fault identification* ("this is the nature of the fault"). Active FTC aims at designing a control law that will fulfill the required function while taking advantage of FDI, whereas passive FTC pursues the same objective using robust control without FDI. This paper will focus exclusively on FDI, leaving aside reconfiguration. Most often, on-line fault diagnosis is tackled via hardware redundancy – multiple sensors or actuators with the same function. The major drawbacks of this approach, however, are higher costs and lower autonomy because of the additional weight, volume and power required [2]. Two main groups of approaches attempt to circumvent these difficulties. The first one is *analytical redundancy*, which exploits the relations between measured or estimated variables in order to detect possible dysfunctions of the system. This set of methods is commonly called *model-based*, where model should be understood as a knowledge-based dynamical model. The behavior of the system can indeed be modeled as a set of differential equations, usually in state-space form. The second group of approaches is often referred to as *model-free*, although it actually uses the redundancy and correlations of the data in a hidden manner. It exploits measurements, acquired in real time or available in a previously constructed database, using a behavioral model of the signal. In what follows, we shall keep with the classical terminology distinguishing model-based and model-free approaches.

Section 2 presents a brief, but eclectic, overview of the main FDI methods. The approaches are classified according to the kind of knowledge they require. The underlying hypotheses and uncertainty modelling are also taken into account in this classification. Model-free approaches are displayed first, followed by the model-based methods. Finally, we show how fault isolation is performed independently of the class of method. Throughout Section 2, we have attempted to gather methods coming from different fields in generic groups. Details are not given for lack of space, but references are provided to complete the insight. We do not claim exhaustivity; for example discrete-event simulation (including fault trees, petri nets...) is beyond the scope of this paper.

Section 3 describes the proposed aeronautical benchmark. First, the class of aerial systems considered along with the hypotheses and knowledge available are detailed. The explicit model of the system, including its sensors and actuators, is then presented. The generic set of devices composing the vehicle is considered as a given and no hardware redundancy is allowed. The set of possible faults affecting the sensors and actuators is specified and simulation results of normal and faulty flights are shown. Finally, the most promising approaches for this generic case study are pointed out, and performance indices are proposed in accordance with robustness requirements.

2. Fault Diagnosis Approaches

2.1. Basic definitions

The main concepts in fault diagnosis are recalled here, to avoid ambiguity. Further standard definitions can be found in [3]. A *fault* is an unpermitted deviation of at least one characteristic property or parameter of the system. It may lead to a *failure*, which is a permanent interruption of the system ability to perform some required function. The main goal of FDI is then to identify incipient faults as soon as possible so as to prevent failures. A fundamental problem in this context is the achievement of a satisfactory trade-off between non-detection and false alarms.

Residuals are signals computed on the basis of measurements in operation. They should remain small as long as there is no fault, and become sufficiently large to be noticeable whenever a fault occurs. Residuals are most often involved in the model-based methods, but should also be considered within model-free approaches.

2.2. Model-free approaches

When no explicit dynamical model is available, the system knowledge boils down to real-time measurements, possibly completed by process history. With such data, two main strategies could be adopted. In a sense, both aim at interpolating the new measured point based on the available data. The first strategy is *classification*. It involves building classes from the database either in a supervised way (i.e., with the help of an expert) or in an unsupervised manner (i.e., collecting elements of the database that are close to one another). A classifier is then trained with respect to these classes to perform the classification of the newly measured variables as representative of a healthy or faulty behavior. The second strategy is *regression*. It builds a statistical model that uses the redundancy of the process history in order to predict the values of the new variables and generate residuals by comparing predictions to measured values.

2.2.1. Qualitative approaches: When no process history is available, the only exploitable information concerning the system monitored is the empirical knowledge of experts, which can be used to build expert systems. An expert system consists of a set of rules that aim to mimic human reasoning, by associating premises and conclusions to determine the logical chain of events. A fault is then reported if a forbidden sequence of events is detected. The major drawbacks of this approach are its lack of generality and its inability to handle situations that have not been explicitly taken into account in the design of the knowledge base [4].

Another class of methods, close to the previous one, is known as *qualitative trend analysis* (QTA). Its purpose is to decompose a measured signal into a sequence of known primitives (e.g "stable", "increase" or "decrease"). This recognition can be achieved either by analyzing the sign of the successive derivatives and using them in a rule base, or by matching patterns with a database containing samples of known primitives [5]. Both techniques imply the cautious design of heuristic rules. Faults are identified in the same manner as with expert systems.

2.2.2. Early tendencies in pattern recognition for fault diagnosis: When a process history is available, diagnosis can be viewed as a pattern recognition task where newly acquired measurements are to be classified in predetermined modes. The main tasks to be achieved are precisely described in [6] and summarized in this section. The prior knowledge is a database comprising d observations of the n variables. First, two off-line operations have to be carried out. The data must be clustered into classes and a decision rule should be trained. M classes are thus defined and each vector of the database is assigned to one of them. For diagnosis, the modes to be considered are the healthy one and all of the possible faulty ones. Labeling can be performed by an expert, if available, or with an algorithm like k-means clustering [7]. If the database contains only non-faulty measurements, another solution is to perform one-class classification [8], although this will not make fault isolation practicable. Once the data have been labeled, a decision rule must be chosen and trained so as to classify the vectors in the proper classes. Parametric and non-parametric approaches are available for that purpose.

Parametric discrimination: If the probability density function of the observations is assumed fixed, it can be estimated from the data (e.g., with a maximum-likelihood method under hypotheses on the law). If moreover a prior probability is assigned to each class, then new measurement vectors can be classified using the Bayes rule. When there is no information concerning prior probability density functions of the observations and classes, only parametric discrimination by the mean of direct boundary computation is feasible. The simplest case is linear binary classification, on which most methods are built. Considering two classes, its aim is to find a hyperplane that splits the data into two parts with respect to the predefined labels. This separator is designed optimally according to some predefined cost function; a norm should be chosen to achieve data fitting, along with a regularization term to avoid overfitting. The solution is then obtained by minimizing this cost with, e.g., the Ho-Kashyap algorithm. For nonlinear problems where no linear separator exists, more complex functions (quadratic, cubic...) could be used but involve the tuning of a dangerously increasing set of parameters. A very popular solution to design separators for classification (or approximators for regression) has been to resort to *neural networks*. Actually, the design difficulty moves from choosing the parameters of the analytical separator to the selection of an activation function and the choice of the structure of the network, i.e., the number of layers and the number of neurons composing each of them. Minimizing the quadratic distance between the output of the network and the label of the class requires the tuning of the weights of the neurons, usually with the well-known back-propagation algorithm. Note that the convergence of this gradient algorithm to a global minimizer of the cost function is not guaranteed. These tools have been widely used in FDI for a large panel of applications [9].

Non-parametric classification methods: If the design of a separator is intractable, a distance combined with a voting scheme can complete classification. Considering the labeled data, a new point is classified in conformity with its neighborhood. The best-known method is the k-nearest neighbor algorithm, which gives its value to the new point according to the majority of the labels of the k-nearest points. Of course, a distance should be chosen to determine which points are the “nearest”. A histogram or a grid could also replace the distance to analyze the neighborhood influence on the point considered.

On-line classification is then carried out in conformity with the separation rule chosen. To take care of unexpected situations, *rejection* (both in distance and ambiguity) may also be considered. Distance rejection means that if the newly acquired measure is considered too far from the existing classes, then it is not taken into account. Ambiguity rejection concerns the points that are too close to the border of the classes, alternative solutions being the introduction of fuzzy borders or Bayesian priors. Rejected points may be used to add new classes to the classifier or to redefine the classification rule on-line.

2.2.3. Advanced pattern recognition – Kernel machines: Two key notions are used in modern pattern recognition, namely those of *kernel* and of *sparsity.* Kernel classifiers look for a function separating the classes. This function is built using training data. The main difficulty then is to represent the link between the samples and particularly to extract useful information from the large amount of redundant data. The *kernel trick* makes it possible to generalize linear methods, such as those described in Section 2.2.2, by mapping the data into a high-dimensional feature space. The output of a kernel machine could be expressed as $\mathbf{y}_{\mathrm{m}}(\mathbf{x}) = \sum_{i=1}^{d} \alpha_i \cdot k(\mathbf{x}, \mathbf{x}_i)$ where $\mathbf{x}$ is the new input point, the $\mathbf{x}_i$'s are training points, $k(.,.)$ is the kernel and the α_i's are weights to be tuned. The main advantage of this formulation is that it provides an easily computable kernel function that expresses the distance between two data points. There is also a need for sparsity as it would be computationally expensive to have significant weights on all the samples while all the points are not relevant. This is accomplished through an appropriate design of the cost function that will be minimized to find the weights α_i of the kernel machine. Mathematical details concerning statistical learning theory may be found in [10].

Vapnik's *Support Vector Machines* (SVM) have popularized these concepts [11]. The goal of SVM is to find a linear separator of the data in the higher-dimensional feature space. The linear separator is designed in order to achieve structural risk minimization (SRM). It aims at avoiding overfitting, which is the main problem of parametric discrimination approaches such as neural networks. The final function is expressed as a projection onto *support vectors*. Another interesting approach is the use of *Gaussian Processes* (GP), which generalize multivariate Gaussian distributions to infinite-dimensional spaces. Gaussian-process regression has been called *Kriging* by the geostatistical community. The covariance of the GP plays the role of the kernel, and sparsity facilitates computation for large-scale problems. These two methods have been shown to be closely related to other well-known regularization techniques, such as ridge regression or least-square classification [12].

The use of the tools described in this section is only at its beginning in the FDI field, which is surprising considering how straightforward it is to adapt these powerful new tools to the problems considered in Section 2.2.2. Very few applications of kernel machines to diagnosis have been reported so far [8], but it seems a promising way to perform or enhance fault detection. Moreover, the criteria used could be modified to perform regression. It then becomes possible to use the same formalism to create a black-box model that can generate residuals by comparing its outputs and the measurements on the system to detect the faults. Finally, it should be pointed out that the choice of the kernel and cost function is crucial and far from trivial, and that adequacy to the data must be carefully checked.

2.2.4. Principal Component Analysis (PCA): PCA projects the training data onto the eigenvectors of the covariance matrix and keeps only the most relevant ones, through a threshold on the eigenvalues, to achieve dimension reduction. The selected eigenvectors are then used to build a linear time-invariant statistical model that can serve to estimate the variables of the process in order to generate residuals. That could be seen as a kernel machine too, since a linear covariance corresponds to a possible choice of kernel. A recursive form exists for dynamical systems [13] and PCA can also be extended to the nonlinear case by the kernel-trick, as described previously [14]. An alternative method with similar extensions is *Partial Least Squares* (PLS) [5].

2.3. Model-based approaches

When the physics of the process is well known, it becomes possible to use an explicit knowledge-based dynamical model. Fault detection then amounts to checking whether the behavior of the monitored process is inconsistent with that of its model. This is done by exploiting the structure of the model and the existing input-output relationships.

2.3.1. Qualitative tools: If a model of the process is available but the confidence in its parameters and quantitative outputs is very low, a first approach could be to derive a qualitative model. Qualitative equations are then used to express the type of variation of the process variables. This qualitative physics has the same goal as the methods outlined in paragraph 2.2.1, i.e., to predict the evolution of the process in order to detect abnormal behaviors [15]. These causal links could also be modeled under the form of a *signed digraph* (SDG) [16]. Although complex systems cannot be efficiently supervised with these approaches, semi-qualitative methods may be of great help to monitor uncertain systems [17].

2.3.2. Parameter monitoring: For faults that may affect the values of some characteristic constants of the system itself, such as leaks entailing abnormal mass variations, parameter estimation techniques should be considered. The nominal value (or a set of admissible values in a bounded-error context [18]) of the monitored parameter vector is supposed known and FDI then boils down to estimating on-line the value of the parameters to generate residuals. Parameter estimation for models that are linear in their parameters can be achieved with, e.g., the recursive least-square algorithm, provided that the input signal is persistent. Nonlinear parameter estimation could be addressed in the same way but the techniques involved may be computationally much more expensive and not guaranteed to converge to an optimal solution.

When a system may suffer from abnormal oscillations or vibrations, diagnosis could be carried out by frequency analysis. Parametric signal models can be used to estimate the main frequencies and estimate possible changes in the system behavior [19].

2.3.3. State estimation: Estimating the state and output of the system makes it possible to create residuals by comparing the reconstructed signals with their real values. State estimators may be classified according to how they represent uncertainty. Considering linear systems first, *Luenberger* observers allow the reconstruction of the state variables under deterministic hypotheses, while *Kalman* filtering achieves the same operation in a stochastic context with assumptions on the distribution of the measurement noise and state perturbations. These now classical methods have been widely used for a large panel of applications [3]. A very useful extension for fault detection is the *Unknown-Input Observer* (UIO), which can be designed in deterministic and stochastic settings. The UIO aims at performing state estimation with minimal influence of the unknown inputs (i.e., exogenous disturbances).

Nonlinear state estimation is often addressed by linearizing the model around an operating point or along a trajectory in order to apply the previous techniques. This has given birth to the *Extended Luenberger Observer*

(ELO), the *Extended Kalman Filter* (EKF) and recently the *Extended Unknown Input Observer* (EUIO) [20]. These approaches lead to relatively easy computation, however linearizing implies losing information.

Contrary to the previous extensions of the Kalman filter, the *Unscented Kalman Filter* (UKF) does not linearize the model. This technique predicts the system behavior by using evaluations of the nonlinear model at a set of points approximating a Gaussian distribution of the state vector [21]. Based on a similar idea, sequential Monte Carlo methods such as *Particle Filtering* (PF) are a very promising approach to deal with nonlinearity. PF is now commonly used for tackling complex fault detection issues [22].

Many observer structures using the nonlinear model without linearizing it have been proposed, from the s*liding mode observer* [23] to nonlinear *adaptive observers* [24] and *high-gain observers* [25]. No general nonlinear structure can be defined and therefore tuning complexity is increasing. Recently a unifying theory of nonlinear observers has emerged [26], proposing the solving of a partial differential equation as the design methodology. If such a solution is reachable for aerospace applications, this could be an interesting framework to consider for observer-based fault diagnosis.

The methods presented so far either do not use an explicit uncertainty representation or assume a statistical distribution, most often Gaussian. An alternative approach is to use bounds on errors. This bounded-error approach can be used for linear and nonlinear models. In nonlinear state estimation, for example, interval analysis can be used to predict the evolution of the set of possible values for the state vector [27]. Part of the predicted set that are inconsistent with measurements can then be eliminated. Fault detection can then be performed by checking whether the resulting set is empty [28]. This uncertainty modelling is not specific to state estimation and can be used in parameter estimation or in the context of decoupling methods (see Section 2.3.4).

To avoid heavy computations, multiple-model strategies are also being investigated. They assume that the nonlinear model of the system can be approximated by interpolating between local linear models. This *Takagi-Sugeno* representation may be obtained analytically or by system identification. It is then possible to build a set of interpolating linear observers and achieve diagnosis operations [29].

To complete this overview, we would like to mention the very recent differential-algebra approach of [30], based on fast state-derivative estimations.

2.3.4. Decoupling methods : The *Fundamental Problem of Residual Generation* (FPRG) [31] is the maximization of sensitivity to faults while minimizing the influence of disturbances. Ideally, each generated residual should be sensitive to one particular fault only. Most classical FDI methods do not take perturbations into account and therefore do not address the FPRG. UIOs address the FPRG by considering the disturbances as unknown inputs, but can only be applied for specific perturbations structures.

Parity relations eliminate unknown state variables from model equations so as to produce residuals that only depend on the system inputs and outputs. This is done by exploiting the model structure and temporal redundancy on a short horizon. The methodology was initially developed for linear systems but it has been extended to bilinear and polynomial models (with the help of elimination theory for the latter). This method has also been applied to systems with unknown parameters [32]. Links between parity space methods and observers have been investigated in [33].

Polynomial approximation of nonlinear systems does not always allow the design of robust FDI filters. A nonlinear geometric approach has been proposed to deal with the FPRG for a larger class of nonlinear systems [31]. This method aims at finding a state and output coordinates transformation that leads to a new set of observable decoupled residuals. The algorithm has been successfully applied to an aircraft nonlinear model [34].

The FPRG could also be addressed through the H_∞ methodology. This multi-objective optimization approach provides a residual filter that maximizes (according to the H_∞ norm) the effect of the faults while minimizing some measure of the influence of the disturbances. One of the difficulties lies in transforming the system into standard form before applying the method [35].

2.4. Performing Fault Isolation and Residual analysis

If decoupling methods succeed in producing residuals that are sensitive to only one fault, fault isolation is directly achieved. Using observer-based techniques, two schemes for fault isolation have been considered, namely the dedicated observer scheme (DOS) and the generalized observer scheme (GOS). DOS is a bank of observers sensitive to only one fault while GOS is composed of observers sensitive to all faults except one. If such a decoupling is impossible, a fault-incidence matrix should be built. This Boolean table shows the influence of each fault on all the residuals. For fault isolation to be possible, each fault should affect a different set of residuals. Note that structural analysis could help a great deal to construct this incidence matrix [36].

When fault isolation has been shown to be possible, whatever the fault detection method employed, a residual analysis has to be performed. A test of the statistical hypotheses should be carried out to distinguish faulty from normal signal behavior. If a learning algorithm is considered, the test will determine whether the current operating point is located in a healthy class or in a faulty one, with the help of the preceding measurements. For a residual generated with a model-based approach, the test will identify a drift in the signal.

Direct fault diagnosis methods apply a fixed threshold as a detection test, but this ignores the previous values of the signal and it is very sensitive to outliers. More sophisticated tests are generally defined to detect a change in the mean of the signal. The Wald test (also called Sequential Probability Ratio Test – *SPRT*) and the Page-Hinkley test (also known as the cumulative sum algorithm – *CUSUM*), based on the maximum likelihood criterion, are most commonly encountered in residual processing. Remarkable on-line extensions of these basic tests have been developed [19].

3. Aeronautical Benchmark

3.1. Hypotheses and airframe dynamics

Test cases for fault diagnosis have already been defined for industrial applications (DAMADICS project [37]) or for specific problems affecting commercial planes such as oscillatory faults [38]. The aeronautical benchmark we aim at defining is centered on an autonomous guided aerial vehicle with typical sensors and actuating devices.

In this paper, the proposed case study involves a six-degree-of-freedom surface-to-air missile. Actuating is performed using usual flight control surfaces and propulsion regulation (either turboprop or ramjet). The main sensor is an Inertial Measurement Unit (IMU), comprising gyrometers and accelerometers. It could possibly be combined with a GPS, however this case is not considered here. These standard components are preset and there is *no hardware redundancy*. We consider a missile with body-of-revolution symmetry, thus the inertia matrix is $\mathbf{I} = \text{diag}(a, b, b)$. The position of the center of gravity on the missile axis is marked by x_{bG} , $[\varphi, \theta, \psi]$ are the Euler angles, $[v_{bx}, v_{by}, v_{bz}]$ is the speed in body coordinates, $[p, q, r]$ is the angular velocity, $[x, y, z]$ is the position in the inertial frame, $Q = \frac{1}{2} \cdot \rho \cdot (v_{bx}^2 + v_{by}^2 + v_{bz}^2)$ is the dynamic pressure, m is the aircraft mass, s_{ref} and l_{ref} are characteristic dimensions. The aerodynamic coefficients $c_{(.)}$ are piecewise constant, except c_{zb}, which is a function of the Mach number and the angle of attack. With this notation, the dynamics is described by the following state equations.

$$\begin{cases}
\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix} = \begin{bmatrix} \cos\psi\cos\theta & -\sin\psi\cos\varphi + \cos\psi\sin\theta\sin\varphi & \sin\psi\sin\varphi + \cos\psi sin\theta\cos\varphi \\ \sin\psi\cos\theta & \cos\psi\cos\varphi + \sin\psi\sin\theta\sin\varphi & -\cos\psi\sin\varphi + \sin\psi\sin\theta\cos\varphi \\ -\sin\theta & \cos\theta\sin\varphi & \cos\theta\cos\varphi \end{bmatrix} \cdot \begin{bmatrix} v_{bx} \\ v_{by} \\ v_{bz} \end{bmatrix} \\
\dot{v_{bx}} = r \cdot v_{by} - q \cdot v_{bz} - \sin\theta \cdot g + \frac{1}{m} \cdot [-Q \cdot s_{\text{ref}} \cdot c_x + f_x] \\
\dot{v_{by}} = p \cdot v_{bz} - r \cdot v_{bx} + \cos\theta \cdot \sin\varphi \cdot g + \frac{1}{m} \cdot \left[-Q \cdot s_{\text{ref}} \cdot c_{zb} \cdot \frac{v_{by}}{\sqrt{v_{by}^2 + v_{bz}^2}}\right] \\
\dot{v_{bz}} = q \cdot v_{bx} - p \cdot v_{by} + \cos\theta \cdot \cos\varphi \cdot g + \frac{1}{m} \cdot \left[-Q \cdot s_{\text{ref}} \cdot c_{zb} \cdot \frac{v_{bz}}{\sqrt{v_{by}^2 + v_{bz}^2}}\right] \\
\dot{p} = \frac{1}{a} \cdot Q \cdot s_{\text{ref}} \cdot l_{\text{ref}} \cdot c_{l\delta_l} \cdot \delta_l \\
\dot{q} = \frac{1}{b} \cdot \left[(b-a) \cdot p \cdot r + Q \cdot s_{\text{ref}} \cdot l_{\text{ref}} \cdot \delta_m - Q \cdot s_{\text{ref}} \cdot c_{zb} \cdot \frac{v_{bz}}{\sqrt{v_{by}^2 + v_{bz}^2}} \cdot (x_{bG} - x_{bG\text{ref}})\right] \\
\dot{r} = \frac{1}{b} \cdot \left[(a-b) \cdot p \cdot q + Q \cdot s_{\text{ref}} \cdot l_{\text{ref}} \cdot \delta_n + Q \cdot s_{\text{ref}} \cdot c_{zb} \cdot \frac{v_{by}}{\sqrt{v_{by}^2 + v_{bz}^2}} \cdot (x_{bG} - x_{bG\text{ref}})\right] \\
\dot{\varphi} = p + (q\sin\varphi + r\cos\varphi)\tan\theta \\
\dot{\theta} = q\cos\varphi - r\sin\varphi \\
\dot{\psi} = (q\sin\varphi + r\cos\varphi)\frac{1}{\cos\theta}
\end{cases}$$

The state vector is $\mathbf{x} = [x, y, z, v_{bx}, v_{by}, v_{bz}, p, q, r, \varphi, \theta, \psi]^{\mathsf{T}}$, the input vector is $\mathbf{u} = [\delta_l, \delta_m, \delta_n, f_x]^{\mathsf{T}}$ where the $\delta_{(.)}$'s are the deflection angles corresponding to the three axes and f_x is the propulsion rate. The IMU/INS system provides a measure of the entire state vector. IMU measurements are subject to errors such as biases, scale factors and noise. Considering, for example, a one-axis sensor $\delta\omega_p$ measuring the roll rate p, the measure is expressed as: $\delta\omega_p = k_p \cdot p + b_p + G(0, \sigma_p)$ where k_p is the scale factor, b_p the bias and σ_p the standard deviation of a zero-mean Gaussian white noise. These three parameters (for each sensor) are characteristic of the IMU and should fall within a set of values provided by the manufacturer.

The aircraft model belongs to the general class of nonlinear control-affine systems, which could be written as:

$$\begin{cases} \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x}) \cdot \mathbf{u} + \mathbf{H} \cdot \mathbf{w} \\ \mathbf{y} = \mathbf{C} \cdot \mathbf{x} + \mathbf{v} \end{cases}$$

where $\mathbf{f}$, $\mathbf{G}$, $\mathbf{H}$ and $\mathbf{C}$ are smooth vector field and matrices with appropriate dimensions, and $\mathbf{w}$ and $\mathbf{v}$ are the perturbation vector and measurement noise.

The aircraft has four flight control surfaces in a symmetric scheme with two elevators $(\delta_{m1}, \delta_{m2})$ and two rudders $(\delta_{n1}, \delta_{n2})$. Servomotors are modeled as second-order systems saturated in position and velocity. The equivalent deflection angles are computed as:

$$\begin{cases} \delta_l = \frac{1}{4} \cdot (-\delta_{m1} + \delta_{m2} - \delta_{n1} + \delta_{n2}) \\ \delta_m = \frac{1}{2} \cdot (\delta_{m1} + \delta_{m2}) \\ \delta_n = \frac{1}{2} \cdot (\delta_{n1} + \delta_{n2}) \end{cases}$$

These modelling and choice of sensors and actuators are representative of a large panel of aerospace vehicles [39].

3.2. *Flight scenario*

The surface-to-air missile aims at intercepting a target that is following a given path at a constant altitude. Taking off from the ground, the missile is guided toward its goal via a proportional navigation guidance law. A ground-based radar provides the missile-target line-of-sight rate to compute the guidance orders. Examples of target and interceptor trajectories are shown in Figure 1. The corresponding initial conditions are in Table I.

Initial conditions	Speed	Position	Departure time
Target	200 m/s	$x = 0$ m $y = 0$ m $z = 3000$ m	0 s
Missile	600 m/s	$x = -7000$ m $y = -5000$ m $z = 0$ m	5 s

Table I
INITIAL CONDITIONS

3.3. *Faults*

A mathematical model of actuator faults should be precisely defined. As in [37] and [39], we express the control input as $\mathbf{u} = \sigma_f \cdot k_f \cdot \mathbf{u_c} + (1 - \sigma_f) \cdot \bar{\mathbf{u}}$ where $\mathbf{u_c}$ is the control input that should be applied if no fault disturbed the system. The other parameters, depending on the type of fault taking place, are specified as follows:

$$\forall t > t_{\text{fault}}, \begin{cases} \text{No fault :} & \sigma_f = 1,\ k_f = 1 \\ \text{Loss of effectiveness :} & \sigma_f = 1,\ 0 < k_f < 1 \\ \text{Freezing :} & \sigma_f = 0,\ k_f = 1,\ \bar{\mathbf{u}} = \begin{cases} \mathbf{u}(t_{\text{fault}})\text{: Lock} - \text{in} - \text{place} \\ \mathbf{u}_{\text{im}} \text{ or } \mathbf{u}_{\text{iM}} \text{ (end stops) : hard} - \text{over} \end{cases} \end{cases}$$

Most often in FDI applications, sensor faults are expressed as additive biases on the measurements. Considering the form of the measurement equation, as explained in Section 3.1, it seems appropriate to consider faults that statistically affect the biases, scale factors or noise standard deviations of the sensor model. This actually represents an awry IMU and is more realistic than a bias appearing abruptly. Based on the possible faults influencing the components that have been defined, Table II illustrates faults that can be applied on the case study. Trajectories corresponding to Fault 2 (actuator fault) and to Fault 3 (sensor fault) are plotted in Figure 1.

Component affected	Fault	Interpretation	Tag
Propulsion	50% Loss of effectiveness	air intake problem	1
Rudder	Lock-in-place	power supply problem	2
z-accelerometer	Bias out of bounds	defective IMU	3
p-gyrometer	Scale factor out of bounds	defective IMU	4

Table II
FAULTS CONSIDERED

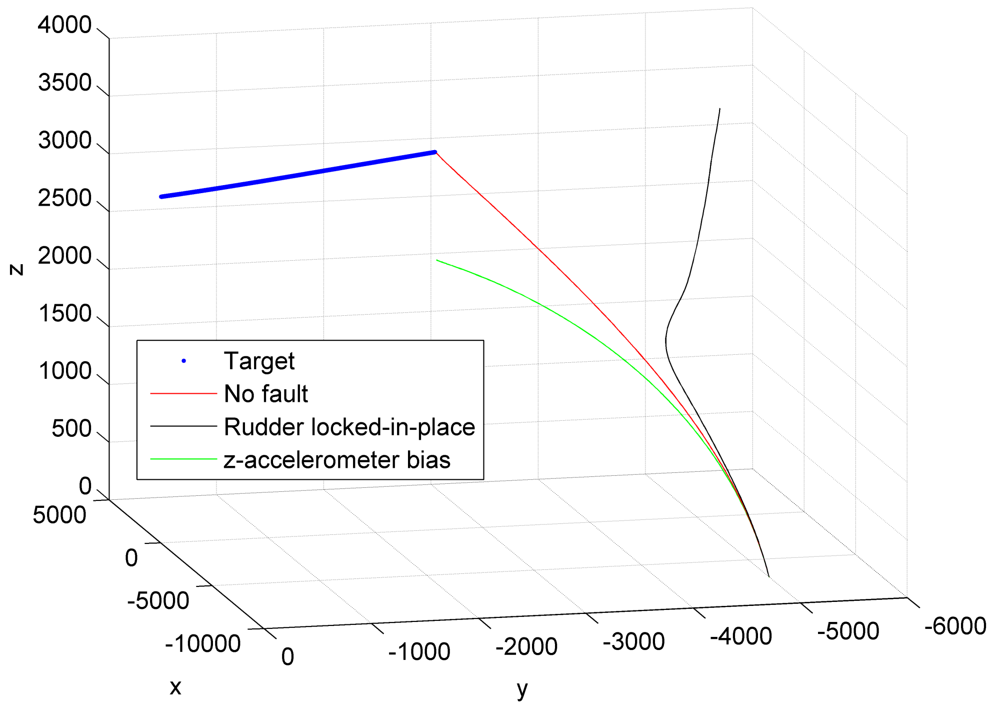

Figure 1. Target, healthy and faulty missile trajectories

3.4. *Methods to be investigated and quantitative performance indices*

As seen in Section 2, reliability on the knowledge about the system is a major criterion for method selection. The model described in Section 3.1 has been validated by previous works in flight mechanics. It represents the behavior of the vehicle satisfactory for given speed and altitude conditions. In actual flight conditions its inner parameters are inaccurate, especially aerodynamic coefficients whose in-flight variations are not well known. This suggests that modern pattern recognition approaches could be of great help to enhance diagnosis robustness to these sources of uncertainty. Nevertheless, these techniques cannot be used alone because no record is available before the mission, and prior knowledge on the system should not be ignored. An interesting approach would be to assist a model-based algorithm with a model-free one, based on in-flight measurements. A line of inquiry may be given by a semi-parametric kernel machine taking into account the influences of both model and measurements to regularize the estimation.

The highly nonlinear dynamics governing the missile is limiting the range of methods using an explicit knowledge-based dynamical model. Indeed, linearization or polynomial approximation would only add more uncertainty to an already inaccurate model. Therefore, we tend to favor particle filtering and nonlinear geometric decoupling as the analytical methods to be considered. Nonlinear observers and multiple model techniques are also interesting contenders, if applicable. The assumption of bounded errors is an attractive way to handle uncertainty, and the aforementioned methods could be adapted in this context.

It must be kept in mind that the missile is controlled by a closed-loop guidance algorithm. This control scheme may lessen the impact of failing components, by modifying the fault dynamics. This may be taken into account in the design of fault detection methods that will test consistency between computed control inputs (i.e., controller outputs) and measured system outputs. It is even possible to design control inputs in such a way as to facilitate FDI; this is known as *active diagnosis* [40].

Objective evaluation of all these methods is necessary in order to build an efficient aircraft FDI methodology. Method-indepedent performance indices must be defined to compare them. These indices should reflect the reliability of the decision rules, as it is a major concern in fault diagnosis. False-alarm and non-detection rates along with their relative amplitude provide such information. It is also essential that faults are detected as quickly as possible to allow the required safety measures to be taken. This is measured by the detection delay. See [37] for precise definitions of a number of performance indices, including those mentioned above.

4. Conclusions

A survey of the main fault diagnosis approaches has been conducted in this paper. A classification according to the type of knowledge available has been proposed, and the links between apparently dissimilar techniques coming from different communities have been emphasized. We defined a test-case for in-flight FDI, based on an interceptor missile with a predetermined set of sensors and actuators with no hardware redundancy. The nonlinear models of the aircraft, its components and the faults affecting them were described in detail. Finally, methods to be further investigated along with the performance criteria to compare them were set forth.

References

[1] G. Ducard and H.P. Geering. Efficient nonlinear actuator fault detection and isolation system for unmanned aerial vehicles. *Journal of Guidance Control and Dynamics*, 31(1):225, 2008.

[2] R.J. Patton. Fault detection and diagnosis in aerospace systems using analytical redundancy. *Computing & Control Engineering Journal*, 2(3):127–136, 1991.

[3] R. Isermann and P. Balle. Trends in the application of model-based fault detection and diagnosis of technical processes. *Control Engineering Practice*, 5(5):709–719, 1997.

[4] C. Angeli and A. Chatzinikolaou. On-line fault detection techniques for technical system: A survey. *International Journal of Computer Science & Applications*, 1(1):12–30, 2004.

[5] V. Venkatasubramanian, R. Rengaswamy, S.N. Kavuri, and K. Yin. A review of process fault detection and diagnosis Part III: Process history based methods. *Computers and Chemical Engineering*, 27(3):327–346, 2003.

[6] B. Dubuisson. *Diagnostic et reconnaissance des formes*. Hermes Paris, 1990.

[7] R.J. Patton, C.J. Lopez-Toribio, and F.J. Uppal. Artificial Intelligence Approaches to Fault Diagnosis for Dynamic Systems. *International Journal of Applied Mathematics and Computer Science*, 9(3):471–518, 1999.

[8] H.J. Shin, D.H. Eom, and S.S. Kim. One-class support vector machines: an application in machine fault detection and classification. *Comput. Ind. Eng.*, 48(2):395–408, 2005.

[9] R.J. Patton, J. Chen, and T.M. Siew. Fault diagnosis in nonlinear dynamic systems via neural networks. In *International Conference on Control, Coventry*, volume 2, pages 1346–1351, 1994.

[10] L. Györfi. *Principles of Nonparametric Learning*. Springer Verlag Wien New York, 2002.

[11] V.N. Vapnik. An overview of statistical learning theory. *IEEE Transactions on Neural Networks*, 10(5):988–999, 1999.

[12] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. Springer-Verlag New York, 2006.

[13] W. Li, H.H. Yue, S. Valle-Cervantes, and S.J. Qin. Recursive PCA for adaptive process monitoring. *Journal of Process Control*, 10(5):471–486, 2000.

[14] J.M. Lee, C.K. Yoo, S.W. Choi, P.A. Vanrolleghem, and I.B. Lee. Nonlinear process monitoring using kernel principal component analysis. *Chemical Engineering Science*, 59(1):223–234, 2004.

[15] L. Chittaro and R. Ranon. Hierarchical model-based diagnosis based on structural abstraction. *Artificial Intelligence*, 155(1):147–182, 2004.

[16] J. Montmain and S. Gentil. Dynamic causal model diagnostic reasoning for online technical process supervision. *Automatica*, 36(8):1137–1152, 2000.

[17] C. Combastel, S. Lesecq, S. Petropol, and S. Gentil. Model-based and wavelet approaches to induction motor on-line fault detection. *Control Engineering Practice*, 10(5):493–509, 2002.

[18] S.M. Castillo, E.R. Gelso, and J. Armengol. Constraint satisfaction techniques under uncertain conditions for fault diagnosis in nonlinear dynamic systems. In *16th IEEE Mediterranean Conference on Control and Automation,*, pages 1216–1221, 2008.

[19] M. Basseville and I.V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall Englewood Cliffs, NJ, 1993.

[20] M. Witczak. *Modelling and Estimation Strategies for Fault Diagnosis of Non-Linear Systems: From Analytical to Soft Computing Approaches*. Springer-Verlag, Berlin-Heidelberg, 2007.

[21] L. Cork and R. Walker. Sensor fault detection for UAVs using a nonlinear dynamic model and the IMM-UKF algorithm. *Information, Decision and Control, Aalborg*, pages 230–235, 2007.

[22] V. Verma, G. Gordon, R. Simmons, and S. Thrun. Real-time fault diagnosis . *IEEE Robotics & Automation Magazine*, 11(2):56–66, 2004.

[23] Y. Xiong and M. Saif. Sliding mode observer for nonlinear uncertain systems. *IEEE Transactions on Automatic Control*, 46(12):2012–2017, 2001.

[24] P.M. Frank and B. Köppen-Seliger. Fuzzy logic and neural network applications to fault diagnosis. *International Journal of Approximate Reasoning*, 16(1):67–88, 1997.

[25] E. Busvelle and J.P. Gauthier. High-gain and non high-gain observers for nonlinear systems. *Contemporary Trends in Nonlinear Geometric Control Theory*, pages 257–286, 2002.

[26] V. Andrieu and L. Praly. On the existence of a Kazantzis-Kravaris/Luenberger observer. *SIAM Journal on Control and Optimization*, 45(2):432–456, 2007.

[27] L. Jaulin, M. Kieffer, O. Didrit, and E. Walter. *Applied interval analysis*. Springer London, 2001.

[28] A. Ingimundarson, J.M. Bravo, V. Puig, T. Alamo, and P. Guerra. Robust fault detection using zonotope-based set-membership consistency test. *International Journal of Adaptive Control and Signal Processing*, 23(4), 2009.

[29] R. Orjuela, B. Marx, J. Ragot, and D. Maquin. State estimation for non-linear systems using a decoupled multiple model. *International Journal of Modelling, Identification and Control*, 4(1):59–67, 2008.

[30] M. Fliess, C. Join, and H. Sira-Ramirez. Non-linear estimation is easy. *International Journal of Modelling, Identification and Control*, 4(1):12–27, 2008.

[31] C. De Persis and A. Isidori. A geometric approach to nonlinear fault detection and isolation. *IEEE Transactions on Automatic Control*, 46(6):853–865, 2001.

[32] A. Shumsky. Redundancy relations for fault diagnosis in nonlinear uncertain systems. *International Journal of Applied Mathematics and Computer Science*, 17(4):477–489, 2007.

[33] R.J. Patton and J. Chen. Robust fault detection using eigenstructure assignment: a tutorial consideration and some new results. In *30th IEEE Conference on Decision and Control, Brighton*, pages 2242–2247, 1991.

[34] N. Meskin, T. Jiang, E. Sobhani, K. Khorasani, and CA Rabbath. Nonlinear geometric approach to fault detection and isolation in an aircraft nonlinear longitudinal model. In *American Control Conference, New York*, pages 5771–5776, 2007.

[35] S. Grenaille, D. Henry, and A. Zolghadri. A method for designing fault diagnosis filters for LPV polytopic systems. *Journal of Control Science and Engineering*, 2008(1), 2008.

[36] D. Düştegör, E. Frisk, V. Cocquempot, M. Krysander, and M. Staroswiecki. Structural analysis of fault isolability in the DAMADICS benchmark. *Control Engineering Practice*, 14(6):597–608, 2006.

[37] M. Bartyś, R.J. Patton, M. Syfert, S. de las Heras, and J. Quevedo. Introduction to the DAMADICS actuator FDI benchmark study. *Control Engineering Practice*, 14(6):577–596, 2006.

[38] L. Lavigne, A. Zolghadri, P. Goupil, and P. Simon. Oscillatory failure case detection for new generation Airbus aircraft: a model-based challenge. In *47th IEEE Conference on Decision and Control, Cancun*, pages 1249–1254, 2008.

[39] J.D. Boskovic, S.E. Bergstrom, R.K. Mehra, S.S. Co, and M.A. Woburn. Retrofit reconfigurable flight control in the presence of control effector damage. In *American Control Conference, Portland*, pages 2652–2657, 2005.

[40] F. Bateman, H. Noura, and M. Ouladsine. Active fault detection and isolation strategy for an unmanned aerial vehicle with redundant flight control surfaces. In *16th Mediterranean Conference on Control and Automation, Ajaccio*, pages 1246–1251, 2008.