

Autonomous fault detection and diagnostics, an enabler to control risks of military operations

Chris Rijdsdijk and Tiedo Tinga, Netherlands Defence Academy, c.rijdsdijk.01@mindef.nl, t.tinga@mindef.nl

Willem van der Sluis, Naval Maintenance and Sustainment Agency, w.vd.sluis.01@mindef.nl

Abstract

This paper aims to reveal how the Netherlands Armed Forces may mature in autonomous fault detection and diagnostics. Generally, fault detection and diagnostics is important as lacking asset health may amplify the risks of military operations. Autonomous fault detection and diagnostics could make an individual's knowledge explicit and it could be essential to process the expanding amount of data from (health) monitoring systems. The state-of-the-art in autonomous fault detection and diagnostics will be projected on a realistic case study. It will be shown that the autonomous fault detection and diagnostics in this typical case study benefits from model-based limits that were already implicitly used. Model-based limits are known to better cope with the varying and unprecedented nature of military operations than value-based limits. The underlying problem is that the risks assigned to faults tend to evolve more rapidly than the designed-in autonomous fault detection and diagnostics. The challenge to make autonomous fault detection and diagnostics robust against varying risk assessments remains unresolved.

1 Introduction

The Netherlands Armed Forces operate in high risk environments of large diversity. These risks may amplify when asset health is lacking. Therefore, fault detection and diagnostics (FDD) may highly affect risks of military operations. In principle, FDD could be autonomous but in practice, it often involves human effort. Reducing human involvement in FDD may be advantageous, not to cut labour costs but to make an individual's expertise transparent to the organisation. Moreover, individuals just appear to be incapable to process the expanding data sets from (health) monitoring systems. Autonomous FDD may become important as newer generations of weapon systems are expected to generate more data.

This paper will project the state-of-the-art in FDD from review papers on the FDD that has been implemented at the platform systems of the vessels at the Royal Netherlands Navy. This exercise should reveal the maturity of the implemented FDD and some directions to improve or develop.


Section 2 will classify FDD methods from review papers by their limits and by their model selection process. Section 3 will introduce autonomous asset health control and various degrees of non-autonomous asset health control. Section 4 will present a case study to illustrate the FDD of the platform management systems of the Royal Netherlands Navy. Section 5 will reflect on the results and finally Section 6 forwards some conclusions.

2 Fault detection and diagnostics

This Section will introduce FDD as an essential element of asset health control. An asset is an item, thing or entity that has potential or actual value to an organisation [1]. This paper confines to assets being artefacts that are valued for fulfilling requirements. If these requirements have been fulfilled, the asset is said to be healthy. So, any asset health assessment is encumbered with subjective requirements. Still, *shared* requirements, i.e. a set of commonly accepted requirements for a specific asset, should be assessable by common sense. Common sense about *shared* requirements enables individuals to agree

upon the actions to control the asset health. Agreement upon actions is essential in any collaboration [2]. The evolution of asset health will be represented by a process. Figure 1 depicts a process P that interacts with its environment by inputs U and outputs Y. Now, the asset health follows from common sense requirements on the inputs U, on the outputs Y or on the process that relates U and Y.

Figure 1. A process



FDD triggers the actions to control the asset health. Figure 2 classifies FDD methods primarily by the limit that detects the fault. A value-based limit has been built on some instantaneous measurement of an input U or an output Y. A value-based limit holds only when all its influential factors remain constant. A model-based limit has been built on an error of a model that relates the inputs U to the outputs Y. Here, an error expresses some distance between a model property and (i) an observation or (ii) a known physical quantity. A model-based limit also holds when its influential factors vary. Most research effort has therefore been directed to model based FDD methods [3-9]. Figure 2 further classifies the model based FDD methods by the model selection process. Knowledge-based model selection relies on knowledge of the physics that prescribes the variables, the parameters, and the structure of the model. Knowledge-based model selection is superior in making claims about unprecedented circumstances. History-based model selection is a resort when knowledge of the model properties is lacking. History-based model selection employs data to estimate unknown model properties or to weigh some arbitrary set of candidate models. Because history-based model selection relies on data from the past, claims about unprecedented circumstances become more risky.

Figure 2. Classification of FDD methods

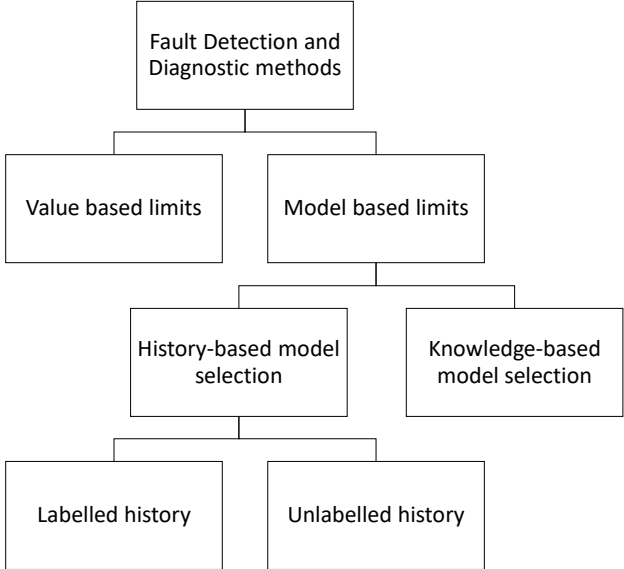


Figure 2 finally shows that history-based model selection may take place from labelled data or from unlabelled data. In the specific case of FDD, the label indicates the presence of the fault. Labelled data allows for supervised machine learning whereas unlabelled data only allows for unsupervised machine learning.

3 Asset health control

This Section will outline asset health control methods. Asset health control expands the objectives of regular FDD:

- Early fault detection. Fault detection is the determination of the faults present in a system and the time of detection.
- Correct fault isolation. Fault isolation is the determination of the kind, location, and time of detection of a fault.
- Correct fault identification. Fault identification is the determination of the size and time-variant behaviour of a fault.

with a recovery objective:


- Timely recovery from the fault.

Section 3.1 will introduce autonomous asset health control and Section 3.2 will introduce non-autonomous asset health control.

3.1 Autonomous control

Many assets respond to disturbances by autonomous control (Figure 3). A disturbance is an unknown (and uncontrolled) input acting on a system [10]. Such a disturbance may lead to a fault. In Figure 3, a fault follows from a deviation from a required input U , from a required output Y , or from a required model property that relates U to Y . Faults could occur abruptly, incipiently, or intermittently. Further, faults induce failures, i.e. a permanent interruption of the asset's ability to perform a required function under specified operating conditions.

Figure 3. Autonomous control.



Autonomous control requires (i) precise knowledge of the control model C and (ii) a means to automate the control. In practice, the control model C imperfectly represents the process P which means that it could only recover from a subset of the disturbances. The output U of the control model C may be continuous (intermediate value control) or dichotomous (on-off control).

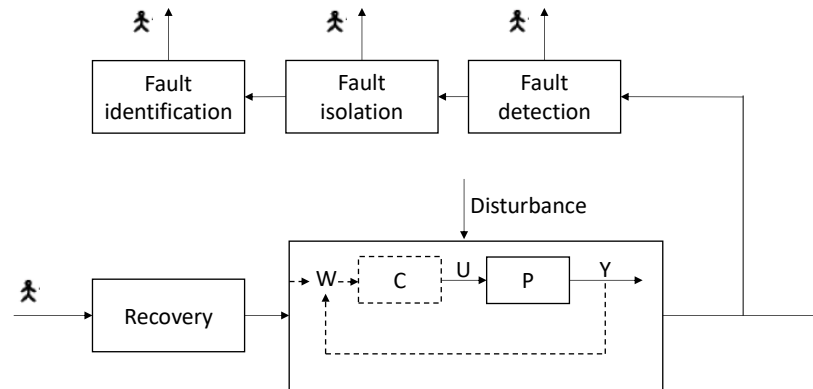
The complementary set of disturbances that is not covered by the control model should be controlled in another way. The autonomous control masks the fault [8],[9], i.e. the feedback loop may propagate a single abnormal process property to all others. Due to the feedback loop, the output Y is subject to influential control factors that vary. Therefore, model-based limits should then be used to control the asset health, as simple value-based limits do not work properly under varying conditions.

3.2 Non-autonomous control

This Subsection will introduce the asset health control of disturbances that involve human effort. Figure 4 shows that non-autonomous asset health control entails fault detection, fault isolation, fault identification and recovery as already mentioned in Section 3. The human involvement may only be limited to just the recovery, but it may also be complete.

The case studies in this paper exemplify a hybrid form of asset health control where fault detection is autonomous, but the follow up is human driven.

Figure 4. Non-autonomous control, where each “person symbol” represents human decision making.



A reduction of the human involvement can be achieved by untangling implicit human expertise and implementing this in the control system. This will make the non-autonomous asset health control more consistent. Ultimately, the non-autonomous control may become fully autonomous. Moreover, autonomous FDD is typically more efficient in analysing large amounts of data [3]. As the amount of data from (health) monitoring systems tends to grow, the importance of autonomous FDD may similarly grow.

4 Case study

This Section will introduce two realistic cases of a sea water cooling system. This sea water cooling system has been installed redundantly on a naval vessel. So, a failure of the duty system can be covered by standby backup systems.

Figure 5. Simplified layout of the sea water cooling system, including sensors measuring the pressure (p) and temperature (T).




Figure 5 depicts a simplified layout of the sea water cooling system. The designer implemented autonomous fault detection but the fault isolation, the fault identification and the recovery rely on human effort (Figure 4). The limits of the autonomous fault detection are all value-based:

- A low-pressure alarm at the discharge flange to indicate lacking hydraulic power.
- A low-temperature alarm at the inlet flange to indicate (the risk of) ice clogging.
- A high-temperature alarm at the inlet flange to indicate lacking cooling capacity.

As the designer predominantly applied value-based limits to identify faults in all platform systems of this vessel, the autonomous fault detection of the sea water cooling system is not exceptional. Section 2 mentioned that value-based limits may be appropriate as all its

influential factors remain constant. Here, the pressure at the discharge flange and the temperature at the inlet flange are not involved in any feedback control that *defines* their dependency on factors that vary. Moreover, the sea water cooling systems lacks pump speed control or flow control while the suction line and the discharge line are constantly connected to the sea. Therefore, value-based limits may be an appropriate fault detection at first glance.

After several years of deployment, the operator’s risk appreciation of the faults in the sea water cooling system evolved [11]. In fact, the operator implicitly implemented model-based limits as he learned to ignore fault detections under specific operating conditions. Moreover, the operator also learned to detect faults that have not explicitly been defined by the designer. The case study is just an attempt to disclose these implicit model-based limits. The model-based limits will appear to be imperfect, but they may be improved by extension (which often requires additional measurements), or by an evaluation of their error.

4.1 Sensor fault


A detected fault may appear to be in the sensor rather than in the asset. So, a fault detection may have several explanations. This case will exemplify a more refined fault isolation (Figure 2). More specifically, this case attempts to better isolate a sensor fault. This fault isolation relies on the knowledge-based model that posits that redundant temperature measurements T_i and T_j should be equal:

$$T_i - T_j = \begin{cases} = 0 & \text{then, no sensor fault} \\ \text{otherwise} & \text{then, sensor fault} \end{cases} \quad (1)$$

As the sea water cooling system (Figure 5) has been installed redundantly, the sea water temperature is measured by four different sensors. Then, Equation (1) may isolate a sensor fault from a faulty state of the sea water cooling system as defined by the designer’s autonomous fault detection system.

Figure 6 shows the evolution in $T_i - T_j$, sensor 1 is compared to sensors 2, 3, and 4. Firstly, all sensors are deemed healthy and halfway the plot sensor T_2 has a fault (as concluded from the large difference with sensor 1). However, even the healthy sensors still show a systematic nonzero difference $T_i - T_j$. A naïve implementation of equation (1) without threshold value would unjustifiably identify many sensor faults.

Figure 6. Evolution in $T_i - T_j$ while the sensors were deemed healthy.



Equation (1) may be improved by (i) an extension of the model or by (ii) an evaluation of the error. An *extension* would require knowledge of other explanations for the difference $T_i - T_j$ and a means to assess these explanations. An extended knowledge-based model

often requires additional measurements. An *evaluation* of the error would require access to a labelled history. From the labelled history, it follows that Figure 6 is initially free from sensor faults before the sensor T_2 turns to a faulty state. The classification in Figure 2 indicates that from this labelled history, a(n ensemble of) model(s) may be selected that could isolate the sensor fault. This history-based model is a resort as the knowledge (of the measurements) to use a(n extended) knowledge-based model is lacking.

Most likely, even the extended knowledge-based model would again appear to be incomplete and a history-based model on its errors may again become a resort. So, a hybrid approach that uses a history-based model *on the errors* in a knowledge-based model seems appropriate here.

4.2 Pump cavitation

This second case is about a cavitation fault that has not been covered by the autonomous fault detection system of the sea water cooling system, but that excessively occurs in practice. The initial risk assessment has been adjusted by operating experience. Coincidentally, the sensor suite (Figure 5) at some of the pump locations has been extended with a pressure measurement at the inlet flange of the pump which enables to construct an imperfect knowledge-based model that detects the cavitation fault.

A model-based limit of cavitation can be obtained by specifying a minimal pressure at the pump inlet. Such a limit can be defined by the required nett positive suction head ($NPSH_R$), a pump specific pressure. A cavitation fault follows from combining this with the available nett positive suction head ($NPSH_A$) as shown in equation (2).

$$\frac{NPSH_A}{NPSH_R} = \begin{cases} \geq 1 & \text{then, no cavitation} \\ < 1 & \text{then, cavitation} \end{cases} \quad (2)$$

Equation 3 defines $NPSH_A$ and Figure 7 specifies $NPSH_R$ for this specific pump.

$$NPSH_A = \frac{p_i}{\rho g} + \frac{c^2}{2g} - \frac{p_v}{\rho g} \quad (3)$$

Here, p_i is the static pressure at the pump inlet, c is the velocity of the sea water at the pump inlet, p_v is the vapour pressure of the sea water, ρ is the density of the sea water and g is the gravitational constant. The parameters g , ρ and p_v follow from deep knowledge and the temperature measurement, but the variables p_i and c remain unknown. Therefore, $NPSH_A$ cannot be compared with $NPSH_R$. However, at some sea water cooling systems (Figure 5) the static pressure at the pump inlet p_i has been measured, which also enables to approximate the velocity by equation (4).

$$c = \frac{q}{1/4\pi d_i^2} = \frac{f(H)}{1/4\pi d_i^2} \approx \frac{f((p_d - p_i)/\rho g)}{1/4\pi d_i^2} \quad (4)$$

Equation (4) states that the velocity equals the flow q divided by the area $1/4\pi d_i^2$ at the pump inlet. The flow q is a function of the pump head H by the pump characteristic in Figure 7 and the pump head may be approximated by the measured difference in the static pressure at the inlet and discharge flange $(p_d - p_i)/\rho g$. This approximation ignores the kinetic energy in Bernoulli's law due to the difference in the area at the inlet and discharge flange.

Figure 7. Pump characteristic and the specification of $NPSH_R$.





Figure 8 shows that the ratio $NPSH_A/NPSH_R > 1$ all the time which means that cavitation did not occur in this time bracket. Still, the ratio evolves non-stationary while showing “cavitation spikes” at some regime changes.

Figure 8. Evolution of the ratio $NPSH_A/NPSH_R$.



Evidently, the ratio $NPSH_A/NPSH_R$ suggests that the sea water cooling pumps are regularly and unnecessarily operating far away from their best efficiency point. As there is no operational need that drives the pump away from its best efficiency point (Eff) in Figure 7, the nonstationary evolution in the ratio $NPSH_A/NPSH_R$ may be seen as a fault. Possibly, some faults in Figure 8 are just explained by the fact that the knowledge-based model is imperfect. If additional explanations are known or measured, they may be addressed by an extended knowledge-based model.

Here, it is only known that cavitation is an issue across the fleet, but there is *no* labelled history that indicates whether all pump locations were evenly affected or (even better) that indicates *when* cavitation took place at a specific pump. Still, experts defined limits on the ratio $NPSH_A/NPSH_R$ beyond which cavitation *generally* takes place. The lower limit $NPSH_A/NPSH_R = 1,1$ in Figure 8 follows from [12] and the upper limit $NPSH_A/NPSH_R = 4$ follows from the maximum pump head in Figure 7. Then, Figure 8 shows that the pump has been operating quite close to the upper limit which makes it vulnerable to discharge cavitation. Discharge cavitation occurs when the pump hardly produces any flow. Then, the sea water will pass at a very high speed through the clearance between the impeller and the pump house. Then, discharge cavitation may occur due to the temporal low pressure behind the impeller blades.

So, unlike the previous case, it was impossible to similarly use a history-based model *on the errors* of a knowledge-based model. In this case, the history-based model *on the errors* of equation (2) has *not* been inferred from data about this specific pump, but from general guidelines that have been inferred from data about other pumps.

5 Discussion

The autonomous FDD of the platform systems at the Royal Netherlands Navy predominantly relies on value-based limits whereas model-based limits are known to better cope with varying and unprecedented operations. Although the operation of the sea water cooling system appeared to be quite stationary and free from autonomous control at first glance, the value-based limits were often implicitly replaced by better model-based limits. By disclosing these implicit model-based limits, the asset health control (Figure 4) could become less reliant on human effort.

Operating experience influences the risks assigned to faults (particularly the frequency component of a risk), whereas the designer's autonomous FDD did not change. If the autonomous FDD were to be updated more regularly, the autonomous FDD would have been better aligned with the concerns of the decision makers.

The case from Section 4.1 showed that the knowledge to extend a knowledge-based model may be lacking. A resort to a history-based model selection *on the errors* of the knowledge-based model appeared to be a pragmatic solution to cover known faults at regular operating conditions. This history-based model selection may follow from a labelled history of a specific asset (Section 4.1) or from guidelines that rely on a history of many assets (Section 4.2).

This specific case study was not hampered by the masking effects of autonomous control, but this does not hold for all platform systems. Then, the ones who intend to make FDD more autonomous should be informed about the applicable control models.

6 Conclusions

This paper observed that the autonomous FDD of the platform systems of the vessels of the Royal Netherlands Navy predominantly relied on value-based limits whereas there were good reasons to believe that model-based limits would in principle be more appropriate under changing and unprecedented operating conditions. This paper demonstrated the benefits of model-based limits in two simple cases that relied on some hybrid form of knowledge-based and history-based model selection.

This paper did not solve the fundamental problem that the risks assigned to faults tend to change more rapidly than the design of the autonomous FDD. This paper just untangled some implicit model-based limits that were already available, but lacking knowledge of (control) models or measurements may obstruct the use of improved model-based limits. If the autonomous FDD could become robust against varying risk assessments, this problem would alleviate.

References

1. International Organization for Standardization (2014) *ISO 55000; Asset management – Overview, principles and terminology*.
2. Rijdsdijk, C. (2016) *Maintenance is unjustifiable: an improved inference*, Doctoral thesis, <https://doi.org/10.3990/1.9789036541909>.
3. Gertler, J.J. (1988) Survey of Model-Based Failure Detection and Isolation in Complex Plants. *IEEE Control Systems Engineering*, vol. 8, issue 6, pp. 3-11.
4. Frank, P.M. (1996) Analytical and Qualitative Model-based Fault Diagnosis – A Survey and Some New Results. *European Journal of Control*, vol. 2, issue 1, pp. 6-28

5. Venkatasubramanian, V., Rengaswamy, R., Yin, K., & Kavuri, S. (2003) A review of process fault isolation and diagnosis Part I: Quantitative model-based methods. *Computers & Chemical Engineering*, vol. 27, issue 3, pp. 293-311.
6. Venkatasubramanian, V., Rengaswamy, R., & Kavuri, S. (2003) A review of process fault detection and diagnosis part II: Qualitative models and search strategies. *Computers & Chemical Engineering*, vol. 27, issue 3, pp. 313-326.
7. Venkatasubramanian, V., Rengaswamy, R., Kavuri, S., & Yin, K. (2003) A review of process fault detection and diagnosis part III: Process history based methods. *Computers & Chemical Engineering*, vol. 27, issue 3, pp. 327-346.
8. Isermann, R. (2005) Model-based fault-detection and diagnosis – status and applications. *Annual Reviews in Control*, vol 29, issue 1, pp. 71-85.
9. Park, Y.J., et al. (2020) A Review on Fault Detection and Process Diagnostics in Industrial Processes. *Processes*, vol. 8, issue 9: 1123, <https://doi.org/10.3390/pr8091123>.
10. Isermann, R., Ballé, P, (1997) Trends in the Application of Model-Based Fault Detection and Diagnosis of Technical Processes. *Control Engineering Practice*, Vol. 5, issue 5, pp. 709-719.
11. Rijdsdijk, C. et al. (2020) Using ship sensor data to achieve smart maintenance? *Proceedings of the International Naval Engineering Conference 2020, October 5-9 2020*, pp. 1-10.
12. Hydraulic Institute (2012) *ANSI HI 9.6.1-2012 Rotodynamic Pumps Guideline for NPSH Margin*.