




# EASA Concept Paper: First usable guidance for Level 1 machine learning applications

A deliverable of the EASA AI Roadmap



## Table of Contents

<b>A. Foreword .....</b>	<b>4</b>
<b>B. Introduction .....</b>	<b>5</b>
1. Statement of issue .....	5
2. AI trustworthiness framework overview .....	7
3. Terminology and scope of the document .....	8
4. Criticality of AI applications .....	9
5. Classification of AI applications .....	9
<b>C. AI trustworthiness guidelines .....</b>	<b>10</b>
1. Purpose and applicability .....	10
2. Trustworthiness analysis .....	12
2.1. Characterisation of the AI application .....	12
2.2. Safety assessment of ML applications .....	16
2.3. Information security considerations for ML applications .....	22
2.4. Ethics-based assessment .....	24
3. Learning assurance .....	31
3.1. Learning assurance process planning .....	33
3.2. Requirements and architecture management .....	33
3.3. Data management .....	34
3.4. Learning process management .....	40
3.5. Training and learning process validation .....	42
3.6. Learning process verification .....	43
3.7. Trained model implementation .....	44
3.8. Inference model verification .....	46
3.9. Independent data and learning verification .....	47
3.10. Requirements verification .....	48
3.11. Configuration management .....	48
3.12. Quality and process assurance .....	48
4. AI explainability .....	49
4.1. AI explainability scope .....	49
4.2. AI dimensions and level of explainability .....	51
4.3. AI Explainability motivations and objectives .....	52
4.4. Contributors to the AI-based systems .....	54
4.5. Human factors guidance for explainability .....	56
5. AI safety risk mitigation .....	57

5.1.	AI safety risk mitigation concept.....	57
5.2.	AI SRM top-level objectives .....	57
6.	Organisations .....	59
6.1.	High-level provisions and anticipated AMC .....	59
6.2.	Design organisation case.....	60
<b>D.</b>	<b>Proportionality of the guidance.....</b>	<b>62</b>
1.	Technical building blocks footprint overview .....	62
2.	Risk-based levelling of objectives .....	63
3.	Rationale for the introduced proportionality .....	66
<b>E.</b>	<b>Annex 1 — Anticipated impact on regulations and MOC for major domains .....</b>	<b>67</b>
1.	Product design and operations .....	67
1.1.	Anticipated impact of the introduction of AI/ML on the current regulations.....	67
1.2.	Anticipated impact of AI/ML guidance on the current AMC/MoC framework .....	68
2.	Training / FSTD .....	69
2.1.	Anticipated impact of the introduction of AI/ML on the current regulations.....	69
2.2.	Anticipated impact of AI/ML guidance on the current AMC/MoC framework .....	70
3.	Ground / ATM/ANS.....	70
3.1.	Current regulatory framework relevant to the introduction of AI/ML.....	70
3.2.	Anticipated impact of AI/ML guidance on the current AMC and GM .....	70
4.	Aircraft production and maintenance .....	71
4.1.	Anticipated impact of the introduction of AI/ML on the current regulations.....	71
4.2.	Anticipated impact of AI/ML guidance on the current MoC framework.....	72
5.	Environmental protection.....	72
5.1.	Current regulatory framework relevant to the introduction of AI/ML.....	72
5.2.	Anticipated impact of AI/ML guidance on the current MoC framework.....	72
6.	Aerodromes .....	73
6.1.	Current regulatory framework relevant to the introduction of AI/ML.....	73
6.2.	Anticipated impact of AI/ML guidance on the current AMC and GM .....	73
6.3.	Preliminary analysis .....	73
6.4.	Anticipated impact of AI/ML guidance on the current and future CSs for aerodrome design and safety-related aerodrome equipment.....	74
<b>F.</b>	<b>Annex 2 — Use cases for major domains .....</b>	<b>75</b>
1.	Introduction .....	75
2.	Use cases — aircraft design and operations.....	77
2.1.	Visual landing guidance system (derived from the CoDANN report use case).....	77
2.2.	Pilot assistance — radio frequency suggestion .....	83

3.	Use cases — ATM/ANS .....	84
3.1.	AI-based augmented 4D trajectory prediction — climb and descent rates .....	84
3.2.	Time-based separation .....	108
4.	Use cases — aircraft production and maintenance.....	114
4.1.	Controlling corrosion by usage-driven inspections.....	114
4.2.	Damage detection in images (X-Ray, ultrasonic, thermography) .....	118
5.	Training / FSTD .....	122
5.1.	Assessment of training performance.....	122
6.	Aerodromes .....	122
6.1.	Detection of foreign object debris (FOD) on the runway .....	122
6.2.	Avian radars .....	123
6.3.	Drone detection systems .....	123
7.	Environmental protection.....	123
7.1.	Engine thrust and flight emissions estimation.....	123
8.	Safety management .....	123
8.1.	Quality management of the European Central Repository (ECR).....	123
8.2.	Support to automatic safety report data capture .....	123
8.3.	Support to automatic risk classification.....	123
<b>G.</b>	<b>Annex 3 — Definitions and acronyms .....</b>	<b>124</b>
1.	Definitions.....	124
2.	Acronyms .....	130
<b>H.</b>	<b>Annex 4 — References .....</b>	<b>132</b>
<b>J.</b>	<b>Annex 5 — Full list of questions from the ALTAI adapted to aviation .....</b>	<b>134</b>
1.	Gear #1 — Human agency and oversight .....	134
2.	Gear #2 — Technical robustness and safety.....	135
3.	Gear #3 — Privacy and data governance.....	137
4.	Gear #4 — Transparency.....	138
5.	Gear #5 — Diversity, non-discrimination and fairness .....	139
6.	Gear #6 — Societal and environmental well-being .....	140
7.	Gear #7 — Accountability .....	142

## A. Foreword

In line with the first major milestone of the European Union Aviation Safety Agency (*EASA Artificial Intelligence (AI) Roadmap 1.0 Phase I* ('Exploration and first guidance development'), this concept paper presents a first set of objectives for *Level 1 Artificial Intelligence* ('assistance to human'), in order to anticipate future EASA guidance and requirements for *safety-related machine learning (ML)* applications.

It aims at guiding applicants when *introducing AI/ML technologies* into systems intended for use in safety-related or environment-related applications in all domains covered by the *EASA Basic Regulation* (Regulation (EU) 2018/1139).

It covers only an initial set of AI/ML techniques and will be enriched with more and more advanced techniques, as the EASA AI Roadmap is implemented.

This document provides a first set of usable objectives, however it does not constitute at this stage definitive or detailed guidance. It will serve as a basis for the *EASA AI Roadmap 1.0 Phase II* ('AI/ML framework consolidation') when formal regulatory development comes into force.

After setting the scene in an introductory Chapter (Chapter B), reminding the reader of the four *AI trustworthiness 'building blocks'*, Chapter C develops the guidelines themselves, dealing with:

- *trustworthiness analysis* (Section C.2);
- *learning assurance* (Section C.3);
- *explainability* (Section C.4); and
- *safety risk mitigation* (Section C.5).

Chapter D introduces *proportionality* which is intended to allow the customisation of the objectives to the specific AI applications.

Chapter E aims at identifying the possible *impacts* of the introduction of AI in the different *implementing rules (IRs)*, *certification specifications (CSs)*, *acceptable means of compliance (AMC)* and *guidance material (GM)* in the domains covered by the EASA Basic Regulation.

Chapter F provides the reader with a set of *use cases* from different aviation domains where the guidelines have been (partially) applied. These use cases serve as demonstrators to verify that the objectives defined in this guidance document are achieved.

Until IRs or AMC are available, this guidance can be used as an enabler or a Swiss Army knife facilitating the preparation of the approval or certification of products, parts and appliances introducing AI/ML technologies. In this respect, this guidance should benefit all aviation stakeholders, end users, applicants, certification or approval authorities.

## B. Introduction

This guidance document represents the first milestone in the implementation of the EASA AI Roadmap v1.0. It provides a first set of technical objectives and organisation provisions that EASA anticipates as necessary for the approval of **Level 1 AI applications** ('assistance to human') and, where practicable, a first set of anticipated means of compliance (MOC) and guidance material which could be used to comply with those objectives.

Note: The anticipated MOC will be completed based on the discussions triggered within certification projects, as well as based on the progress of industrial standards such as the one that is under work in the joint EUROCAE/SAE WG-114/G-34 or EUROCAE/RTCA WG-72/SC-216.

The goal of this document is therefore twofold:

- to allow applicants proposing to use AI/ML solutions in their projects to have an early visibility on the possible expectations of EASA in view of an approval. This material may be referred to by EASA through dedicated project means (e.g. a Certification Review Item (CRI) for certification projects);
- to establish a baseline for **Level 1 AI applications** that will be further refined for **Level 2 and Level 3 AI applications**.

*Disclaimer: To the best of EASA's knowledge, the information contained in these guidelines is accurate and reliable on the date of publication and reflects the state of the art in terms of approval/certification of AI/ML solutions. EASA does, however, not assume any liability whatsoever for the accuracy and completeness of these guidelines. Any information provided therein does not constitute in itself any warranty of fitness to obtain a final EASA approval. These guidelines will evolve over the next 3 years through publication of documents respectively for Level 2 and Level 3 AI applications, while being updated based on its application to Level 1 AI applications. They may evolve as well depending on the research and technology development in the dynamic field of AI research.*

### 1. Statement of issue

AI is a broad term and its definition has evolved as technology has developed. EASA therefore chose in the EASA AI Roadmap 1.0 a wide-spectrum definition that is 'any technology that appears to emulate the performance of a human'.

The EASA AI Roadmap has defined the following taxonomy for AI:

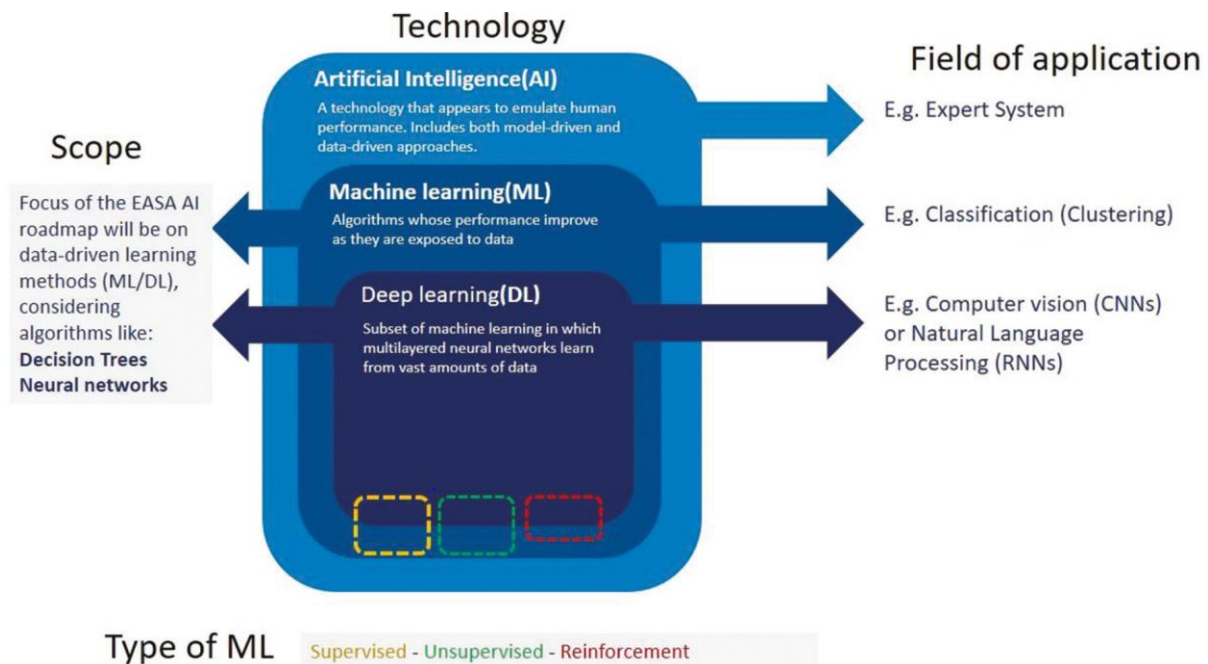


Figure 1 — AI taxonomy in EASA AI Roadmap

The current breakthrough is linked with ML, which is the ability of computer systems to improve their performance by exposure to data without the need to follow explicitly programmed instructions. **Deep learning (DL)** is a subset of ML that emerged with deeper neural networks (NNs), leading to large improvements in performance in the recent years. DL produced significant improvements for many problems in computer vision and natural language processing (NLP), enabling use cases which were not possible before. This is the reason why EASA AI Roadmap 1.0 and this Level 1 AI guidance are focusing on **data-driven AI** approaches.

Data-driven learning techniques are a major opportunity for the aviation industry but come also with a significant number of challenges with respect to the trustworthiness of ML and DL solutions. Here are some of the main challenges addressed through this first set of EASA guidelines:

- Adaptation of assurance frameworks to cover learning processes and address development errors in AI/ML components
- Creation of a framework for data management, to address the correctness (bias mitigation) and completeness/representativeness of data sets used for the ML items training and their verification
- Addressing model bias and variance trade-off in the various steps of ML processes
- Lack of predictability and explainability of ML/DL application behaviour, due to their statistical nature and to model complexity
- Lack of guarantee on robustness and on absence of ‘unintended function’ in ML/DL applications
- Mitigation of residual risk in ‘AI black box’. The expression ‘black box’ is a typical criticism oriented at AI/ML techniques, as the complexity and nature of AI/ML models bring a level of opaqueness that make them look like unverifiable black boxes (unlike rule-based software)
- Lack of trust by end users/operators.



## 2. AI trustworthiness framework overview

To address the challenges of data-driven learning approaches, EASA AI Roadmap 1.0 identifies four ‘**building blocks**’ that are considered essential in creating a framework for **AI trustworthiness** and for enabling readiness for use of AI/ML in aviation:

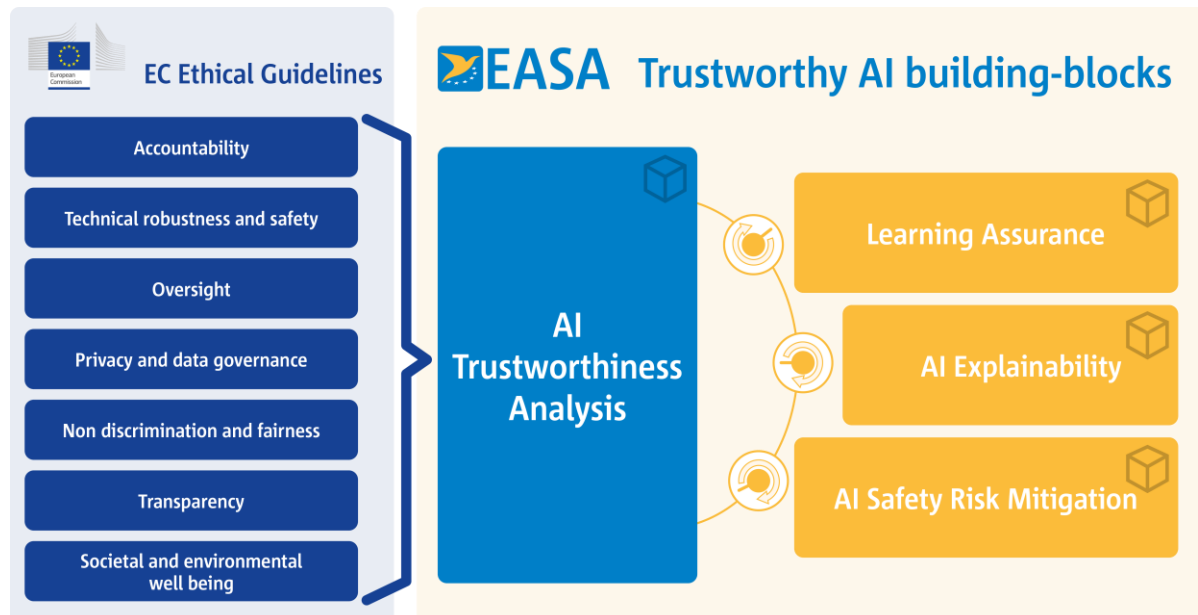


Figure 2 — EASA AI trustworthiness roadmap building blocks

The **trustworthiness analysis**, being one of those four building blocks (see light blue in the middle), creates an interface with the EU Ethical Guidelines created by the EU Commission (EU High-Level Expert Group on AI, 2019), and as such serves as a gate to the three other technical building blocks (see orange on the right). This block also encompasses the **safety assessment** and **security assessment** that are key elements of the trustworthiness analysis concept. All three **assessments (safety, security and ethics-based)** are important prerequisites in the development of any system developed with or embedding AI/ML and are not only preliminary steps but also integral processes towards approval of such innovative solutions.

**Learning assurance** is intended to cover the paradigm shift from programming to learning, as existing development assurance methods are not adapted to cover learning processes specific to AI/ML.

The **AI explainability** deals with the capability to provide the human with understandable and relevant information on how an AI/ML application is coming to its results.

The **AI safety risk mitigation** block considers that we may not always be able to open the ‘AI black box’ to the extent required and that the safety risk may need to be addressed to deal with the inherent uncertainty of AI.

All four building blocks have an importance in gaining confidence in the trustworthiness of an AI/ML application.

The trustworthiness analysis is always necessary and should be performed in its full spectrum for any application. For the other three building blocks, the potentiometers represented in Figure 2 indicate that the depth of guidance could be adapted depending on the application, as described in Chapter D.



### 3. Terminology and scope of the document

In EASA AI Roadmap 1.0, the initial focus has been put on **data-driven AI** approaches. Those can be further divided considering types of learning:

- **Supervised learning** — the process of learning a function that maps an input to an output based on example input-output training samples.
- **Unsupervised learning** — this strategy is used in cases where there is no labelled data set available to learn from. The neural network analyses the data set, and then a cost function tells the neural network how far off target it was. The neural network then adjusts to increase the accuracy of the algorithm.
- **Reinforcement learning** — in this algorithm, the neural network is reinforced for positive results, and punished for a negative result, forcing the neural network to learn over time.

The scope of this document includes at this initial stage **supervised learning** approaches and will be further expanded at a later stage to cover other types of learning.

Considering this scope, the following figure details the decomposition of an AI-based system and allows introducing the terminology that is used in the rest of the document when dealing with the system or portions of it.

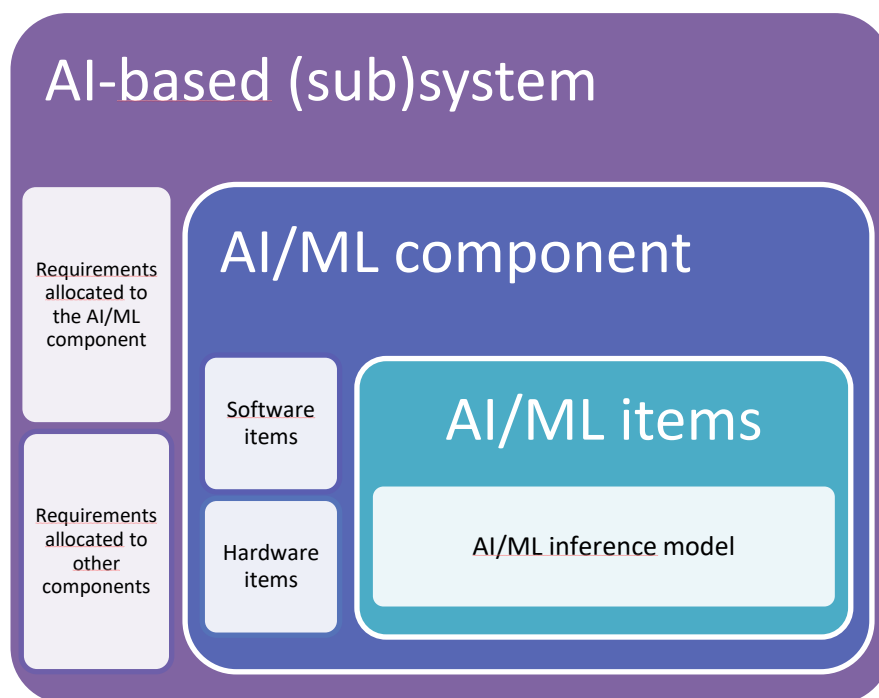


Figure 3 — Decomposition of the AI-based (sub)system

Where:

- an AI-based system is composed of several subsystems, with at least one of them being an AI-based subsystem;
- an AI-based subsystem embeds at least one AI/ML component;

- the AI/ML component is a collection of hardware and/or software items, and at least one AI/ML item;
- the AI/ML item is a specialised software and/or hardware item which contains at least one AI/ML inference model, as a result of the model training and implementation processes;
- the hardware and software items do not include any aspects related to the AI/ML inference model.

Note: The notation ‘(sub)system’ is used where their distinction is not necessary from a guidance perspective, and is equivalent to ‘system/subsystem’.

#### 4. Criticality of AI applications

Depending on the safety criticality of the application, and on the aviation domain, an assurance level is allocated to the AI-based (sub)system (e.g. development assurance level (DAL) for initial and continuing airworthiness or air operations, or software assurance level (SWAL) for air traffic management/air navigation services (ATM/ANS)).

A modulation of the objectives of this document based on the assurance level will be introduced in Chapter D ‘Proportionality of the guidance’ at a later stage, once the framework is completely developed.

#### 5. Classification of AI applications

The EASA AI Roadmap identifies three general levels of AI. This scheme has been proposed based on prognostics from industry regarding the types of use cases foreseen by AI-based systems. Indeed, these three scenarios can be related to the staged approach that most of the industrial stakeholders are planning for the deployment of AI applications, starting with assisting functions (Level 1 AI), then making a step towards more human-machine collaboration (Level 2 AI) and at last seeking for more autonomy of the machine (Level 3 AI).

Some refinement of the three scenarios is considered in the following figure:

Level 1 AI : assistance to human	Level 2 AI : human/machine collaboration	Level 3 AI : more autonomous machine
<ul style="list-style-type: none"> <li>• Level 1A: Human augmentation</li> <li>• Level 1B: Human cognitive assistance in decision and action selection</li> </ul>	<ul style="list-style-type: none"> <li>• Level 2: Human and AI-based system collaboration (Further subcategories not deemed necessary for Level 2)</li> </ul>	<ul style="list-style-type: none"> <li>• Level 3A: AI-based system performs decisions and actions, overridable by the human.</li> <li>• Level 3B: AI-based system performs non-overridable decisions and actions</li> </ul>

Figure 4 — Classification of AI applications

Detailed guidance on how to classify an AI-based system is provided in Section C.2.1.

Chapter D ‘Proportionality of the guidance’ already introduces the applicability of the objectives to each AI level (i.e. 1A and 1B), and will be completed at a later stage with considerations for Level 2 and 3 AI, as well with the modulation based on the assurance level.

## C. AI trustworthiness guidelines

### 1. Purpose and applicability

This chapter introduces a first set of objectives, in order to anticipate future EASA guidance and/or requirements to be complied with by safety-related ML applications. Where practicable, a first set of anticipated MOC has also been developed, in order to illustrate the nature and expectations behind the objectives.

The aim is to provide applicants with a first framework to orient choices in the development strategy for ML solutions. This first set of usable objectives does not however constitute either definitive or detailed means of compliance.

These guidelines apply to any system developed using ML techniques or incorporating ML algorithms (further referred to as AI-based system), and are intended for use in safety-related applications or for applications related to environmental protection covered by the Basic Regulation, in particular for the following domains:

- **Initial and continuing airworthiness**, applying to systems or equipment required for type certification or by operating rules, or whose improper functioning would reduce safety (systems or equipment contributing to failure conditions Catastrophic, Hazardous, Major or Minor);
- **Air operations**, applying to systems, equipment or functions intended to support, complement, or replace pilot tasks (examples may be information acquisition, information analysis, decision-making, action implementation and monitoring of outputs);
- **ATM/ANS<sup>1</sup>**, applying to equipment or procedures intended to support, complement or replace end-user tasks (examples may be information acquisition, information analysis, decision-making and action implementation) delivering ATS or non-ATS services;
- **Maintenance**, applying to systems supporting scheduling and performance of tasks intended to timely detect or prevent unsafe conditions (airworthiness limitation section (ALS) inspections, certification maintenance requirements (CMRs), safety category tasks) or tasks which could create unsafe conditions if improperly performed ('critical maintenance tasks');
- **Training**, applying to systems used for monitoring the training efficiency or for supporting the organisational management system, both in terms of compliance and safety;
- **Aerodromes**, applying to systems that automate key aspects of aerodrome operational services, such as the identification of foreign object debris, the monitoring of bird activities, and the detection of drones around/at the aerodrome;
- **Environmental protection**, applying to systems or equipment affecting the environmental characteristics of products. Note: While the use of AI/ML applications in such systems or equipment may not be safety-critical, the present guidance may still be relevant to establish the necessary level of confidence in the outputs of the applications.

---

<sup>1</sup> For the ATM/ANS domain, according to the currently applicable Regulation (EU) 2017/373, there is no separate approval for the ATM/ANS equipment, and all the activities related to the changes to the functional system (hardware, software, procedures and personnel) are managed under the change management procedures, as part of the air navigation service provider change management process. Competent authority approval is obtained for the introduced complete change. Furthermore, in this Regulation, only the air traffic service (ATS) providers are requested to perform a safety assessment as part of the change management process whereas the non-ATS providers (e.g. CNS) are requested to perform a safety support assessment, intended to assess and demonstrate that after the introduction of the change the associated services will behave as specified and will continue to behave as specified.

The introduction of AI/ML in these different aviation domains may thus imply (or ‘require’) as well adaptations in the respective organisational rules per domain (such as for design organisation approval (DOA) holders, maintenance organisation approval (MOA) holders, continuing airworthiness management organisations (CAMOs), air navigation service providers (ANSPs), approved training organisations (ATOs), operators, etc.). Each organisation would need to ensure compliance with EU regulations (e.g. for initial airworthiness, continued airworthiness, air operations, ATM/ANS, occurrence reporting, etc.) as applicable to each domain. Furthermore, each organisation would need to assess the impact on its internal processes in areas such as competence management, design methodologies, change management, supplier management, occurrence reporting, cybersecurity aspects or record-keeping.

The applicability of these guidelines is limited as follows:

- covering *Level 1 AI applications*, but not covering *Level 2 and 3 AI applications*;
- covering *supervised learning*, but not other types of learning such as *unsupervised or reinforcement learning*;
- covering *offline<sup>2</sup> learning* processes where the model is ‘frozen’ at the time of approval, but not *adaptive or online learning* processes.

---

<sup>2</sup> Offline versus online learning, i.e. learning relying on a ‘frozen’ training set, or learning relying on data acquired during operation [source (DEEL Certification Workgroup, 2021)].

## 2. Trustworthiness analysis

### 2.1. Characterisation of the AI application

#### 2.1.1. High-level function/task and the AI-based (sub)system

The first step should consist in identifying the operational environment in which the AI-based (sub)system will be used.

**Objective CO-01:** The applicant should identify the high-level function(s)/task(s) to be performed by the (sub)system either in interaction with the human or in autonomy.

At this stage, it is not yet decided what the interactions between the human (end user) and the machine will be.

In a second step, based on the high-level function(s)/task(s), the AI-based (sub)system should be determined.

**Objective CO-02:** The applicant should define the AI-based (sub)system taking into account domain-specific definitions of 'system'.

Note: The definition of system varies between domains. For example:

- for airborne systems, ARP4761 defines a system as 'combination of inter-related items arranged to perform a specific function(s)';
- for the ATM/ANS domain (ATS and non-ATS), Regulation (EU) 2017/373 defines a functional system as 'a combination of procedures, human resources and equipment, including hardware and software, organised to perform a function within the context of ATM/ANS and other ATM network functions'.

Once the AI-based (sub)system has been determined, two separate but correlated activities should be executed:

- Definition of the concept of operation (ConOps), with a focus on the identification of the end users and the authority allocation between the end user and the AI-based (sub)system; and
- A functional analysis of the AI-based (sub)system.

These activities will provide necessary inputs for the classification of the AI application, for safety, security, and ethical assessment, as well as for the different building blocks of the AI trustworthiness concept.

#### 2.1.2. Concept of operation for the AI application

To support compliance with the objectives of the AI trustworthiness guidelines, a detailed ConOps detailing precisely how the system will be operated is expected to be established.

**Objective CO-03:** The applicant should define and document the ConOps for all AI-based (sub)systems. A focus should be put on the definition of the operational design domain (ODD) and on the capture of specific operational limitations and assumptions.

**Anticipated MOC-CO-03:** The ConOps should be described at the level of the AI-based system, where the human is expected to achieve an objective, a function or a task.

Once this level of system is defined, the ConOps should consider:

- an end-user-centric operational description of the AI-based system;
- the list of potential end users (roles with their competencies, duties, and responsibilities);
- how the end users will interact with the AI-based system;
- how the end users will benefit from the AI-based system;
- descriptions of the operational scenarios in their ODD; and
- some already identified mitigations and limitations on the AI-based system.

Figure 5 — Interrelationship between ConOps and ODD shows the interrelationship between the operational scenarios for the concept of operation (ConOps) and the operating parameters for the operational design domain (ODD):




Figure 5 — Interrelationship between ConOps and ODD

Notes:

- ConOps limitations may be accounted for in activities related to the system safety assessment or safety support assessment, as described in Sections C.2.2.2 and C.2.2.3.
- Due to the data-driven nature of ML applications, the precise definition of the ConOps is an essential element to ensure that sufficient and representative data is collected for the data sets that are used for training, validation and testing purposes.

### 2.1.3. Functional analysis of the AI-based (sub)system

**Objective CO-04:** The applicant should perform a functional analysis of the (sub)system.

The functional analysis consists in identifying, proposing a break-down of the high-level function(s) into sub-function(s), allocating the sub-function(s) to the subsystem(s) and items in line with the architecture choices.

Notes:

- The functional analysis is an enabler to meet the objectives in Section C.3.2 ‘Requirements and architecture management’ of the learning assurance.
- The functional analysis is a means supporting the functional hazard analysis (FHA) as per Section C.2.2.2 ‘System safety assessment’.

#### 2.1.4. Classification of the AI application

This first usable guidance document focuses on **Level 1 AI applications**. It therefore provides classification guidelines for this level, including boundaries between the Levels 1A, 1B and 2, in order to avoid confusion of the applicants on the classification of their proposed AI-based systems.

To this purpose, EASA is taking advantage of the seminal ‘A model for Types of Human Interaction with Automation’ research paper (Parasuraman-et-al, 2000). According to the authors, the four-stage model of human information processing has its equivalent in system functions that can be automated. The authors propose that automation can be applied to four classes of functions:

- **Information acquisition** involves sensing and registration of input data; these operations are equivalent to the first human information processing stage, supporting human sensory processes.
- **Information analysis** involves cognitive functions such as working memory and inferential process.
- **Decision-making** involves selection from among decision alternatives.
- **Action implementation** refers to the actual execution of the action choice.

The research paper foresees several levels of automation (from Low to High) for each function. In early publications, the HARVIS research project (Javier Nuñez et al., 2019) made use of this scheme to develop a Level of Automation (LOAT), further splitting this scheme by distinguishing between an action performed to ‘automation support’ the human versus an action performed ‘automatically’ by the system.

To further refine this scheme, when considering the anticipated distinction between the **Level 2 AI** and **Level 3 AI** applications, a further decomposition is introduced for ‘automatic’ functions into ‘**overseen**’, ‘**overridable**’ or ‘**non-overridable**’ by the human.

- **Overseen**: capability for the human to closely monitor the functions allocated to the AI-based system (every decision-making and action implementation), with the ability to intervene in every decision-making and/or action implementation of the AI-based system.
- **Overridable**: capability for the human to supervise the operations of the AI-based system (some decision-making and some action implementation), with the ability to override the authority of the AI-based system (some decision-making and some action implementation) when it is necessary to ensure safety and security of the operations (e.g. upon alerting).
- **Non-overridable**: human has no capability to override the AI-based system’s operations.

Note: It is important to remind that the levels of **AI applications** introduced in **EASA AI Roadmap 1.0** is not meant to be an automation scheme but a typology of usage of AI-based systems. Therefore, the detailed levels introduced in the Raja Parasuraman paper (Low-High) or in the HARVIS deliverables (Levels A0 to D8) were not considered in this document.

The resulting classification scheme is as follows and provides a non-ambiguous scheme for the classification of AI-based (sub)systems. The applicant should consider the functions that are allocated to the AI-based (sub)system, excluding the ones that are performed by the human and the ones allocated to other (sub)systems (not based on AI or ML technology). An exception is identified which relates to AI-based (sub)systems which will operate at Level 3B. For such AI applications, there are no functions that can be performed by the human.



EASA AI Roadmap AI Level	High level function/task allocated to the (sub)systems
Level 1A Human augmentation	Automation support to information acquisition
	Automation support to information analysis
Level 1B Human assistance	Automation support to decision-making
Level 2 Human-AI collaboration	Overseen automatic decision-making
	Overseen automatic action implementation
Level 3A More autonomous AI	Overridable automatic decision-making
	Overridable automatic action implementation
Level 3B Autonomous AI	Non-overridable automatic decision-making
	Non-overridable automatic action implementation

*Table 1 — EASA AI typology and definitions*

**Objective CL-01:** The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.

## 2.2. Safety assessment of ML applications

### 2.2.1. AI safety assessment concept

#### 2.2.1.1. Statement of issue

The objective of a **safety assessment** is to demonstrate an acceptable level of safety as defined in the applicable regulations. A logical and acceptable inverse relationship must exist between the occurrence probability of a failure condition and the severity of its effect. Depending on the domain of applications in aviation, **safety assessment** methodologies may vary, but a common point is to consider that only hardware components are subject to random failure. The reliability of a given piece of software is not quantified per se. As an example, for airborne systems, it is usually considered that when known **development assurance** methodologies are used throughout the development, the risk of having an error resulting in a failure is minimised to an adequate level of confidence. Development errors are considered as a possible common source type and are mitigated by system architecture and analysed with other common mode errors and failures via dedicated techniques such as common mode analysis. Contribution of software components taken into account in the probabilistic risk assessment is then usually limited to the reliability of the software function input parameters and to the reliability of the platform executing the software code.

Due to their statistical nature and to model complexity, ML applications come with new limitations in terms of predictability. Taking this into consideration, this guidance is intended to assist applicants in demonstrating that systems embedding AI/ML components operate at least as safely as traditional system developed using existing **development assurance** processes and **safety assessment** methodologies<sup>3</sup>: the AI technology introduction should be done at no higher risk imposed to persons, personal properties (or critical infrastructure). Furthermore, the proposed guidance is also aimed at following as closely as possible existing aviation **safety assessment** processes to minimise the impact on existing safety processes.

#### 2.2.1.2. Safety assessment concept

To demonstrate that a proper safety level will be achieved and maintained throughout the product life, safety assessment is expected to be performed by the applicant:

- Initial safety assessment, during design phase by considering the contribution of an AI/ML component to system failure and by having particular architectural considerations when AI is introduced;
- Continuous safety assessment, with the implementation of a data-driven AI safety risk assessment based on operational data and occurrences. This 'continuous' analysis of in-service events may rely on processes already existing for domains considered in this guideline. The processes will need to be adapted to the AI introduction.

It is recognised that depending on the domains, the necessary activities to be performed and documented in view of EASA approval vary significantly. The table below summarises per domain the expected analysis to be performed in view of the approval by EASA of a system embedding an AI/ML component.

<sup>3</sup> In the ATM/ANS domain, for non-ATS providers, the safety assessment is replaced by a safety support assessment.

Aviation domains	'Initial' safety assessment	'Continuous' safety assessment
Initial and continuing Airworthiness	As per Section C.2.2.2 below	As per organisation section of this guideline (DOA – Continuous airworthiness)
Air operations	To be defined	As per organisation section of this guideline (Operators – SMS)
ATM/ANS	As per Section C.2.2.2 for ATS providers and Section C.2.2.3 for non-ATS providers – see note B	As per organisation section of this guideline (ANSPs – SMS)
Maintenance	None – see note A and D	As per organisation section of this guideline (CAMO or MOA – SMS)
Training	None – see note A and E	Managed from an organisation, operations and negative training, as per organisation section of this guideline (ATO – SMS)
Aerodromes	To be defined	To be defined
Environmental protection	None – see Note F	Currently not applicable

Note A: For some domains, only a 'continuous' safety assessment is expected to be presented to EASA at this stage. Applicants may however use the methodology described in Section C.2.2.2 to establish their design processes.

Note B: Regulation EU 2017/373 that addresses ATS and non-ATS providers has introduced the need of a 'safety support assessment' for non-ATS providers rather than a 'safety assessment'. The objective of the safety support assessment is to demonstrate that, after the implementation of the change to the functional system, the non-ATS will behave as specified and will continue to behave only as specified in the specified context. For these reasons, a dedicated Section C.2.2.3 has been created for non-ATS providers.

Note C: The terminology used in safety assessment/safety support assessment between the various domains varies. Footnotes have been used in the next paragraph to clarify the guideline intent with regard to domain specificities and domain-specific definitions reminded in Chapter G.

Note D: For the maintenance domain, whenever new equipment is used, it should be qualified and calibrated.

Note E: For the training domain, whenever an AI-based system is adopted, the entry into service period should foresee an overlapping time to enable validation of a safe and appropriate performance.

Note F: For the environmental protection domain, the initial safety assessment is to be interpreted as the demonstration of compliance with the applicable environmental protection requirements.

## 2.2.2. System safety assessment

In the following sections, the steps highlighted in **blue** are expected to be affected by AI introduction compared to a classical system safety assessment.

In the following sections, system safety assessment should be understood as safety assessment on the functional system when it applies to ATS providers in the ATM/ANS domain. Safety support assessment on the functional system applies to non-ATS providers and are addressed in Section C.2.2.3.

### 2.2.2.1. Impact assessment of AI introduction on system safety methodologies

The analyses below describe the typical system safety activities performed throughout the design phase.

- Perform functional hazard assessment in the context of the ConOps
- Perform safety assessment
  - Definition of safety objective<sup>4</sup>, proportionate with the hazard classification
  - Definition of a preliminary system architecture to meet safety objectives
  - Derive safety requirements including independence requirements to meet the safety objective and support the architecture
  - Allocate assurance level (e.g. DAL or SWAL)
  - Define AI/ML component performance/reliability metrics
  - Analyse and mitigate the effect of AI/ML component exposure to input data outside of the ODD
  - Perform AI/ML component failure mode effect analysis
- Verification — final safety assessment
- Consolidate the safety assessment to verify the safety objective

#### 2.2.2.2. Proposed objectives for the system safety methodologies

Based on the high-level impact assessment performed (see the previous paragraph), the following objectives are considered in the guideline:

**Objective SA-01:** The applicant should define metrics to evaluate the AI/ML component performance and reliability.

Depending on the application under consideration, a large variety of metrics may be selected to optimise the performance of AI/ML components. The selected metrics should also provide relevant information with regard to the actual AI/ML component reliability so as to support the system safety assessment in **Objective SA-02**.

---

<sup>4</sup> In the ATM/ANS domain, for ATS providers, this activity corresponds to the following ones:

- definition of safety criteria;
- definition of safety requirements proportionate to the risk assessment.

Performance evaluation is performed as part of learning assurance per **Objectives LM-09** (for the trained model) and **IMP-06** (for the inference model). The measured performance is fed back to the safety assessment process.

Once this step is complete, the applicant is expected to estimate the generalisation gap. This is done through **Objective LM-04** in the learning assurance chapter. The output of this objective may then be fed into the system safety assessment. There may be some iterations between **Objective LM-04** and **Objective SA-02** below in case the generalisation guarantee does not allow meeting the safety objective. In such a case, either stronger guarantee may be achieved by constraining further the learning process or changes to the system (e.g. system architecture consideration) may be considered.

**Objective SA-02:** The applicant should perform a system safety assessment for all AI-based (sub)systems.

**Anticipated MOC-SA-02-1:** In performing the safety assessment, the applicant should address the following aspects:

- Perform functional hazard assessment in the context of the ConOps;
- Define safety objectives<sup>5</sup> proportionate with the hazard classification;
- Analyse the system architecture to ensure that it meets the safety objectives. As part of this activity, the necessary architectural mitigation means should be identified<sup>6</sup>;
- Define and validate assumptions;
- Derive safety requirement including independence requirements;
- [Allocate AI/ML items with an appropriate assurance level \(e.g. DAL or SWAL\);](#)
- [Analyse and mitigate the effect of AI/ML component exposure to input data outside of the ODD;](#)
- [Establish AI/ML component failure modes<sup>7</sup>:](#)
  - [Establish a taxonomy of AI/ML component failures;](#)
  - [Evaluate possible failure modes and associated detection means \(see also MOC-SA-02-2 for considerations on the generalisation guarantees\);](#)
- Verify that the proposed implementation satisfies the safety requirements including the independence requirements.

**Anticipated MOC-SA-02-2:** Once the generalisation gap has been evaluated (per **Objective LM-04**), the applicant should assess the impact on the system safety assessment.

When quantitative assessment is required to demonstrate the safety requirements are met, AI/ML component failure rates may be evaluated from the ‘out-of-sample error’ ( $E_{out}$ ). One possible approach is to define the ‘in-sample error’ ( $E_{in}$ ) using a metric that reflects application-specific quantities commensurate with the safety hazard. Then, provided that  $E_{in}$  is defined in a meaningful and practical way,  $E_{out}$ , that reflects the safety performance in operations, can be estimated from the

<sup>5</sup> See Note 4.

<sup>6</sup> This particular step is not specific to systems embedding AI/ML components. However, particular attention is expected to be paid on this step when AI/ML components are used in critical systems as they are likely to drive the system architecture. Indeed, to meet safety objectives associated with the most critical functions, and based on the reliability of AI/ML components/items currently available, it seems likely that the use of redundant AI/ML components/items associated with various architecture mitigation means will be necessary. For example, means may be implemented to detect a failed AI/ML component and revert to a backup non-AI/ML component.

<sup>7</sup> Based on the state of the art in AI/ML, it is acknowledged that relating the notion of probability in AI/ML with safety analyses is challenging (e.g. as discussed in Section 4.2.4.1 ‘Uncertainty and risk’ in (DEEL Certification Workgroup, 2021)) and subject to further investigation.

$E_{in}$  and the generalisation gap. Such errors are however quantities on average and this should be taken into account.

Note: For non-AI/ML components, traditional safety assessment methodology should be considered.

The following standards and implementing rules with adaptation may be considered:

- For embedded systems:
  - ED79A/ARP4754A and ARP4761
- For ATS providers: the following implementing rule requirements are applicable:
  - ATS.OR.205 Safety assessment and assurance of changes to the functional system (and the associated AMC and GM);
  - ATS.OR.210 Safety criteria (and the associated AMC and GM).

### 2.2.3. Safety support assessment

This section is only applicable to non-ATS providers in the ATM/ANS domain. In the following paragraphs, the steps highlighted in blue are expected to be affected by AI introduction compared to a classical safety support assessment.

#### 2.2.3.1. Impact assessment of AI introduction in safety support assessment

The analyses below describe the typical safety support assessment activities performed during design phases. The steps highlighted in blue are expected to be affected by AI introduction compared to the usual process:

- Evaluate impact on the service specification, including service performance;
- Identify applicable service performance requirements;
- Define the AI/ML component performance/reliability metrics;
- Analyse and mitigate the effect of AI/ML component exposure to input data outside of the ODD;
- Perform AI/ML component failure mode effect analysis;
- Analyse degraded modes of the services;
- Elaborate safety support requirements;
- Allocate assurance level (e.g. SWAL);
- Verify safety support requirements.

#### 2.2.3.2. Proposed objectives for the safety support assessment

Based on the high-level impact assessment performed (see the previous paragraph), the following objectives are considered in the guideline:

**Objective SA-03:** The applicant should define metrics to evaluate the AI/ML component performance.

Depending on the application under consideration, a large variety of metrics may be selected to optimise the performance of AI/ML components. Performance evaluation is performed as part of learning assurance, per **Objectives LM-09** (for the trained model) and **IMP-06** (for the inference model). The measured performance is fed back to the safety assessment process.

Component performance in the above objective also encompasses elements including but not limited to reliability, continuity, and integrity.

Once this step is complete, the applicant is expected to estimate the generalisation gap. This is done through **Objective LM-04** in the learning assurance chapter. The output of this objective may then be

fed into the safety support assessment: There may be some iterations between **Objective LM-04** and **Objective SA-04** below in case the generalisation guarantee does not allow meeting these safety support requirements. In such a case, either stringer guarantees may be achieved by constraining further the learning process or changes to the system (e.g. system architecture changes) may be considered.

**Objective SA-04:** The applicant should perform a safety support assessment for any change in the functional (sub)system embedding a component developed using AI/ML techniques or incorporating AI/ML algorithms.

**Anticipated MOC-SA-04-1:** In performing the safety support assessment, the applicant should address the following aspects:

- Evaluate impact on the service specification, including service performance;
- Analyse degraded modes of the services. To this purpose,
  - establish AI/ML component failure modes<sup>8</sup>:
    - establish a taxonomy of AI/ML component failures;
    - evaluate possible failure modes and associated detection means (see also MOC-SA-04-2 for considerations on the generalisation guarantees);
  - analyse and mitigate the effect of AI/ML component exposure to input data outside of the ODD;
- Perform a safety support assessment. To this purpose,
  - define safety support requirements proportionate with degraded mode effects;
  - define a preliminary system architecture to meet the safety support requirements;
  - allocate AI/ML component(s) with an appropriate assurance level (e.g. SWAL);
  - verify that the proposed implementation satisfies the safety support requirements.

**Anticipated MOC-SA-04-2:** Once the generalisation gap has been evaluated (per **Objective LM-04**), the applicant should assess the impact on the system safety support assessment.

Note: For non-AI/ML components, traditional safety support assessment methodology should be considered. The following implementing rule requirement is applicable:

- ATM/ANS.OR.C.005 Safety support assessment and assurance of changes to the functional system (and the associated AMC and GM)

---

<sup>8</sup> See Note 7.



## 2.3. Information security considerations for ML applications

When dealing with ML applications, whatever the domain considered, the data-driven learning process triggers specific considerations from an information security perspective.

Focusing on the initial and continuing airworthiness domains, with Decision 2020/006/R, EASA has amended the Certification Specifications (CSs) for large aircraft and rotorcraft, as well as the relevant AMC and GM, introducing objectives aimed at assessing and controlling safety risks posed by information security threats. Such threats could be the consequences of intentional unauthorised electronic interaction (IUEI) with systems on the ground and on board of the aircraft.

For systems and equipment based on AI/ML applications, this work will serve as a basis to orient the specific guidelines for information security. To this extent, key aspects are:

- the identification of security risks and vulnerabilities through a product information security risk assessment (PISRA) or, more in general, an information security risk assessment;
- the implementation of the necessary mitigations to reduce the risks to an acceptable level (acceptability is defined in the relevant CS for the product); and finally
- the verification of effectiveness of the implemented mitigations. Effectiveness verification should entail a combination of analysis, security-oriented robustness testing and reviews.

For the initial and continued airworthiness of airborne systems embedding AI/ML applications, the guidance from AMC 20-42 'Airworthiness information security risk assessment' is applicable, although contextualised to take into account the peculiarities of the AI/ML techniques.

For other domains, as already stated in Section C.2.2.1.2 for the safety risk assessment, the necessary activities to be performed and documented in view of EASA approval may be different. However, the aforementioned key aspects remain applicable and before dedicated AMC are defined for those other domains, the principles of AMC 20-42 could be used to deal with AI/ML applications information security risk assessment and mitigation.

Since security aspects and properties of AI/ML applications are still object of study, there are no precise answers about the generalisation of the protection measures. Therefore, we have to consider that the initial level of protection of an AI/ML application may degrade more rapidly if compared to a standard aviation technology. In light of this, systems embedding an AI/ML component should be designed with the objective of being resilient and capable of failing safely if attacked by unforeseen and novel information security threats.

### 2.3.1. Proposed objectives for the information security risks management

Based on the high-level considerations made in the previous paragraph, the following objectives are considered in the guideline:

**Objective IS-01:** For each AI-based (sub)system and its data sets, the applicant should manage those information security risks with an impact on safety, identifying and accounting for specific threats introduced by AI/ML usage.

**Anticipated MOC IS-01:** In performing the system information security risk assessment and risk treatment, the applicant should address the following aspects:

- Consider those threat scenarios related to unauthorised modifications of the training data set commonly referred to as ‘training set poisoning’.
- Consider those threat scenarios related to inputs that are specifically crafted to look genuine to a human analysis, but can cause output errors, such as erroneous classification of images or patterns. These threat scenarios are commonly referred to as ‘adversarial attacks’.
- The list of scenarios that will be considered have to be communicated to EASA as soon as practical and need to be tailored for the specific application.

## 2.4. Ethics-based assessment

As already mentioned above, the EU Commission's AI High-Level Expert Group (HLEG) elaborated that, deriving from a fundamental-rights-based and domain-overarching list of **4 ethical imperatives** (i.e. respect to human autonomy, prevention of harm, fairness and explicability), the **trustworthiness of an AI-based system** is built upon **3 main pillars** or expectations, i.e. lawfulness, adherence to ethical principles, and technical robustness. The HLEG further 'operationalised' these expectations by means of a set of **7 gears** and sub-gears (i.e. human agency and oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination and fairness; societal and environmental well-being; accountability). To ease self-evaluation and provide orientation to applicants, the HLEG underpinned this set of gears by a so-called **Assessment List for Trustworthy AI (ALTAI)**<sup>9</sup>, containing several questions and explanation.

Overview of the different layers:

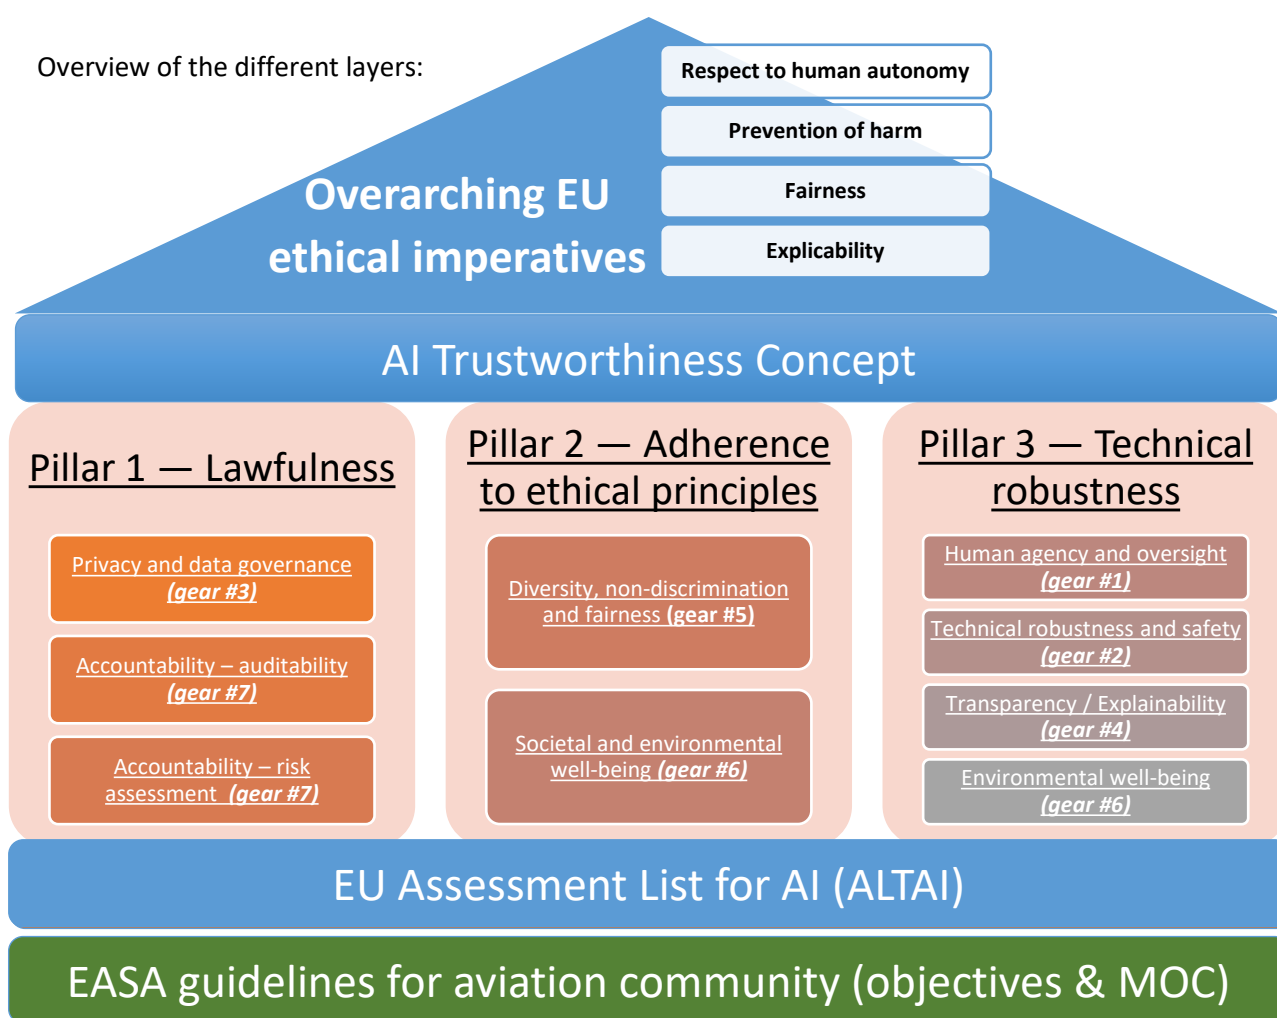


Figure 6 — EU Commission ethical assessment layers (visualisation of the gears distribution from the EASA perspective)

Building on this approach, the present EASA guidelines further clarify and tailor as necessary the above general gears and sub-gears of the HLEG to the needs of the aviation sector and its stakeholders. The

<sup>9</sup> [https://ec.europa.eu/newsroom/dae/document.cfm?doc\\_id=68342](https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=68342)

document therefore formulates a number of Objectives and Anticipated Means of Compliance (MOC) per gear. In addition, it provides a slightly adapted wording of the ALTAI itself.

**Objective ET-01:** The applicant should perform an ethics-based trustworthiness assessment for any AI-based (sub)system developed using ML techniques or incorporating ML algorithms.

**Objective ET-02:** In performing the ethics-based trustworthiness assessment, the applicant should address questions from the EU Commission Assessment List for Trustworthy AI (ALTAI), taking into account the clarifications brought in the following anticipated MOC.

#### 2.4.1. Gear #1 — Human agency and oversight

##### **Anticipated MOC ET-02-1: Clarifications for ‘human agency’ in aviation applications**

To respond to expectations related to human agency, the applicant should perform the necessary analyses and consider in particular the following questions of the ALTAI in the ethical assessment documentation.

- a. **Is the AI-based system designed to interact, guide or take decisions by end users that affect humans or society?**
  - Could the AI-based system generate confusion for end users and/or subjects on whether a decision, content, advice or outcome is the result of an algorithmic decision?
  - Are end users and/or other subjects adequately made aware that a decision, content, advice or outcome is the result of an algorithmic decision?
- b. **Could the AI-based system generate confusion for end users and/or subjects on whether they are interacting with a human or AI system?**
  - Are end users and/or subjects informed that they are interacting with an AI-based system?
- c. **Could the AI-based system affect human autonomy by generating over-reliance by end users?**
  - Did you put in place procedures to avoid that end users over-rely on the AI system?
- d. **Could the AI-based system affect human autonomy by interfering with the end user’s decision-making process in any other unintended and undesirable way?**
  - Did you put in place any procedure to avoid that the AI system inadvertently affects human autonomy?
- e. **Does the AI-based system simulate social interaction with or between end users and/or subjects?**

##### **Anticipated MOC ET-02-2: Clarifications for ‘human oversight’ in aviation applications**

To respond to expectations related to human oversight, the applicant should consider in particular the following questions of the ALTAI in the ethical assessment documentation.

**f. Please determine whether the AI-based system is overseen by a human-in-the-loop, human-on-the-loop, human-in-command, considering the definitions below.**

- From a design perspective, the **‘capability for the human to oversee the full design cycle of the AI-based system’** is a prerequisite in the aviation domain. It is therefore a strong requirement for any AI application, irrespective of the planned oversight mechanism. This also means that **‘self-learning’ approaches are not supported by these guidelines.**
- From an operational perspective, the EU Commission Guidelines on Trustworthy AI (EU High-Level Expert Group on AI, 2019), introduce definitions for the governance mechanisms human-in-command (HIC), human-in-the-loop (HITL) and human-on-the-loop (HOTL). As those definitions would require refinement for aviation and these mechanisms are not further used in the current version of this document, it was not deemed necessary to provide a different set of definitions at this stage.

Note: Applicants may find necessary to answer this question and characterise the functions/tasks of the AI-based system(s) with such oversight mechanisms. In such a case, the applicant should clarify the definitions used.

This guidance does not require a formal classification of the ‘oversight mechanism(s)’ applicable to the AI-based system, as this information is not used in any of the objectives of this document.

**g. Have the humans been given specific training on how to exercise human oversight?**

- From an AI-based system design perspective, this question is to be analysed and answered by the applicant, including specific assumptions made on the necessary training provided to the humans performing oversight.
- From an operational perspective, this question pertains to the organisations of the end users of the system. Therefore, coordination with the end users’ organisations is required to fully answer this question to ensure adequate level of awareness and competence when interacting with the AI-based system.

**h. Did you establish any detection and response mechanisms for undesirable adverse effects of the AI system for the operator?**

- This question is answered through compliance with the objectives of the safety and security assessments, the learning assurance, the explainability and the safety risk mitigation.

Question i. may not be applicable to all use cases; therefore, it has not been reproduced here (it can be found in Annex 5 — Full list of questions from the ALTAI adapted to aviation). It is answered through compliance with the requirements for the safety and security assessments and the safety risk mitigation objectives.

**j. Did you take any specific oversight and control measures to reflect the self-learning or autonomous nature of the AI system?**

- The **‘capability for the human to oversee the full design cycle of the AI-based system’** is a strong requirement for any AI application, irrespective of the planned oversight mechanism. Therefore, application for ‘self-learning’ AI/ML items will not be accepted by EASA at this stage.
- More autonomous systems are considered to be covered under Level 3 AI applications and will be addressed in a future revision of these guidelines.

### 2.4.2. Gear #2 — Technical robustness and safety

#### Anticipated MOC ET-02-3: Clarifications for all aspects of ‘Technical robustness and safety’ in aviation applications

All questions related to ‘Technical robustness and safety’ are considered by EASA to be addressed through compliance with the objectives of the safety and security assessments, the learning assurance, the explainability and the safety risk mitigation.

Note: Regarding adaptive (also known as continual or online) learning, it is not addressed in the current guidelines; therefore, such applications will not be accepted by EASA at this stage.

The following mapping is intended to clarify the precise links to the EASA guidelines (see Annex 5 — Full list of questions from the ALTAI adapted to aviation):

Gear	Question	Applicable section
Resilience to attack and security	a, b, c, d, e, f	C.2.3 Information security considerations for ML applications
General safety	g, h, i, j, k	C.2.2 Safety assessment
Accuracy	l	C.2.2 Safety assessment
Accuracy	m, n, o, p	C.3 Learning assurance
Reliability, fallback plans and reproducibility	q, r, t	C.2.2 Safety assessment
Reliability, fallback plans and reproducibility	s	C.5 Safety risk mitigation
Reliability, fallback plans and reproducibility	u	Not within the scope of the current guidelines (see Note above).

Table 2 — Mapping of technical robustness and safety questions to EASA AI Roadmap building blocks

### 2.4.3. Gear #3 — Privacy and data governance

**Objective ET-03:** The applicant should comply with national and EU data protection regulations (e.g. GDPR), i.e. involve their Data Protection Officer (DPO), consult with their National Data Protection Authority, etc.

#### Anticipated MOC ET-03-1: Clarifications for ‘Privacy and data governance’ in aviation applications

All ALTAI questions related to ‘Privacy and data governance’ in terms of personal data, are considered to be addressed through compliance with the EU and national data protection regulations.

The applicant should thus ensure and provide a confirmation that a ‘data protection’-compliant approach was taken, e.g. through a record or a data protection impact assessment (DPIA).

#### 2.4.4. Gear #4 — Transparency

##### Anticipated MOC ET-02-4: Clarifications for ‘Transparency’ in aviation applications

All questions related to ‘Transparency’ are considered to be addressed through compliance with the objectives of the safety and security assessments, the learning assurance, the explainability and the safety risk mitigation.

The following mapping is intended to clarify the precise links to the EASA guidelines:

Gear	Question	Applicable section
Traceability	a	C.3 Learning assurance
Traceability	b, e	C.5 Safety risk mitigation
Traceability	c, d, f	C.4.4.3 Data recording
Explainability	g, h	C.4 Explainability
Communication	i, j	C.4 Explainability

Table 3 – Mapping of transparency questions to EASA AI Roadmap building blocks

#### 2.4.5. Gear #5 — Diversity, non-discrimination and fairness

##### Anticipated MOC ET-02-5:

This gear may not be applicable to all use cases. Therefore, in a first analysis, applicants should check whether the AI-based system could have any impact on diversity, non-discrimination or fairness.

If no impact exists, the record of this analysis should be added to the ethical assessment documentation.

In case of an impact, please consider the questions from the ALTAI related to Gear #5. The original questions can be found in the Annex 5 — Full list of questions from the ALTAI adapted to aviation. The assessment of the answer to these questions does not fall under the remit of EASA and would be performed by a competent authority for ‘Diversity, non-discrimination and fairness’ matters, at a European Level or at a national Level as applicable.

#### 2.4.6. Gear #6 — Societal and environmental well-being

**Objective ET-04:** The applicant should assess the environmental impact of the AI-based (sub)system.

##### Anticipated MOC ET-04-1: Clarifications for ‘Environmental well-being’

To align with EU guidelines on societal and environmental well-being, the applicant should perform the necessary analyses and answer the following questions in the ethical assessment documentation. Adaptations from the ALTAI are in italics.

- a. **Did you identify potential negative impacts of the AI system on the environment or human health throughout its life cycle (development, deployment, use, end of life)?**
  - Does the AI system require additional energy and/or generates additional emissions?



- Does the AI system have adverse effects on the product's environmental compatibility, in particular on aircraft/engine noise and emissions arising from the evaporation or discharge of fluids?
- Does the AI system have adverse effects on the product's environmental performance in operation?
- Could the use of the AI system have rebound effects, e.g. lead to an increase in traffic, which in turn could become harmful for the environment or human health?

**b. Did you assess the impacts of the AI system on the environment or human health throughout its life cycle?**

**c. Did you define measures to reduce or mitigate these impacts?**

#### **Anticipated MOC ET-02-9: Work and skills, and impact on society at large or democracy**

While developing products assisted by AI systems, applicants should also consider the impacts on the human resources affected by these new developments.

This sub-gear may not be applicable to all use cases. Therefore, in a first analysis, applicants should check whether the AI-based system could have any impact on work and skills.

If no impact exists, the record of this analysis should be added to the ethical assessment documentation.

In case of an impact, please consider the questions from the ALTAI related to Gear #6 'Work and skills' and 'Impact on society at large or democracy'. The original questions can be found in the Annex 5 — Full list of questions from the ALTAI adapted to aviation. The assessment of the answer to these questions does not fall under the remit of EASA and would be performed by a competent authority for 'Work and skills, and impact on society at large or democracy' matters, at a European Level or at a national Level as applicable.

### **2.4.7. Gear #7 — Accountability**

#### **Anticipated MOC ET-02-11: Accountability / Auditability**

The AI system should be auditable by internal and external parties, including the approving authorities, as prescribed in the Basic Regulation.

**a. Did you establish mechanisms that facilitate the AI system's auditability (e.g. traceability of the development process, the sourcing of training data and the logging of the AI system's processes, outcomes, positive and negative impact)?**

- Do these mechanisms ensure the auditability of the AI system's components (including algorithms and datasets), independently of their source (internal or from suppliers)?

**b. Did you ensure that the AI system can be audited by independent third parties?**

### **Anticipated MOC ET-02-12: Accountability / Monitoring**

Questions c., d., e. and f. of the accountability gear may not be applicable to all use cases. Therefore, in a first analysis, applicants should check whether the AI-based system could have any impact on the monitoring of ethical concerns from an organisation's perspective.

If no impact exists, the record of this analysis should be added to the ethical assessment documentation.

In case of an impact, please consider the questions from the ALTAI related to Gear #7 'Accountability' 3rd, 4th and 5th and 8th bullet points. Those questions correspond to items c, d, e and f in Annex 5 — Full list of questions from the ALTAI adapted to aviation and can also be found in the ALTAI (EU High-Level Expert Group on AI, 2020).

### **Anticipated MOC ET-02-13: Accountability / Risk Management**

The AI system should foresee a risk management process, including interaction with third parties.

- g. Did you establish a process to discuss and continuously monitor and assess the AI system's adherence to this Assessment List for Trustworthy AI (ALTAI)?**
  - Does this process include identification and documentation of conflicts between the six aforementioned gears or between different ethical principles and explanation of the 'trade-off' decisions made?
  - Did you provide appropriate training to those involved in such a process and does this also cover the legal framework applicable to the AI system?
  
- h. Did you establish a process for third parties (e.g. suppliers, end users, subjects, distributors/vendors or workers) to report potential vulnerabilities, risks or bias in the AI system?**
  - Does this process foster revision of the risk management process?

### 3. Learning assurance

In the current regulatory framework, the associated risk-based approach for systems, equipment and parts is mainly driven by a requirements-based ‘development assurance’ methodology during the development of their constituents. Although the system-level assurance might still require a requirements-based approach, it is admitted that the design-level layers that rely on learning processes cannot be addressed with ‘development assurance’ methods.

Intuitively, the assurance process should be shifted on the correctness and completeness/representativeness of the data (training/validation/verification data sets) and on the learning and its verification. Most importantly, the main challenge lies in providing guarantees that the training performed on sampled data sets can generalise with an adequate performance on unseen operational data.

To this purpose, a new concept of ‘learning assurance’ is proposed to provide novel means of compliance. The objective is to gain confidence at an appropriate level that a ML application supports the intended functionality, thus opening the ‘AI black box’ as much as practicable.

#### Definition

**Learning assurance:** All of those planned and systematic actions used to substantiate, at an adequate level of confidence, that errors in a data-driven learning process have been identified and corrected such that the system satisfies the applicable requirements and provides sufficient generalisation guarantees.

To illustrate the anticipated learning assurance process steps, EASA proposes the following W-shaped process outline.

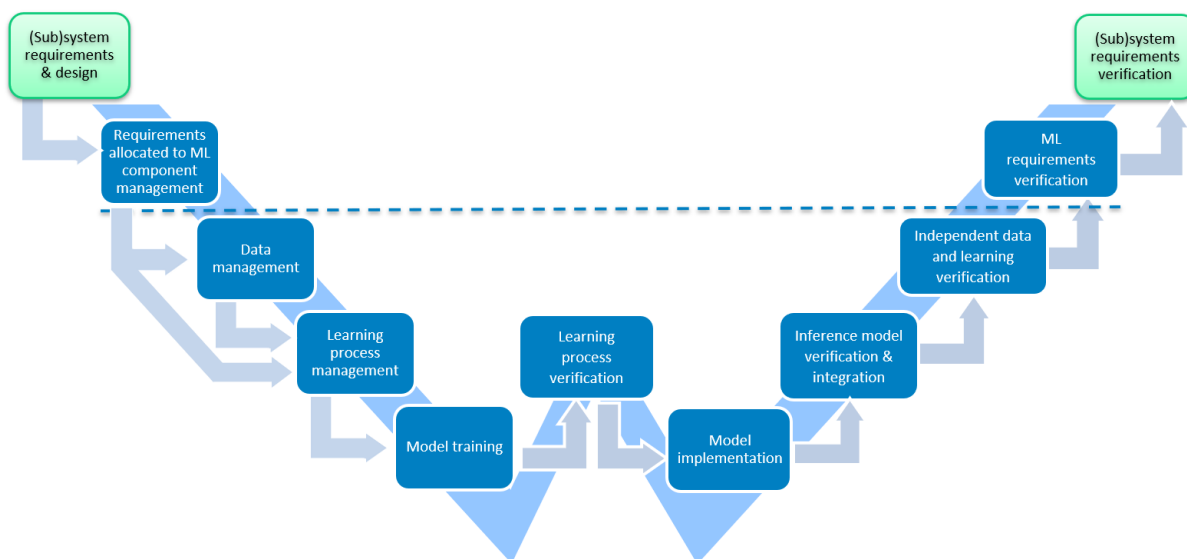


Figure 7 — Learning assurance W-shaped process

This cycle adapts the typical development assurance V-cycle to ML concepts and aims to be a framework for EASA’s MOC. It requires some further refinement, but already allows to structure the guidance through blocks composing the cycle.

The dotted line is here to make a distinction between the use of traditional development assurance processes (above) and the need for processes adapted to the data-driven learning approaches (below).

Note: The pure learning assurance processes start below the dotted line. It is however important to note that this dotted line is not meant to split specific assurance domains (e.g. system / software).

This W-shaped process is concurrent with the traditional V-cycle that is required for development assurance of non-AI/ML components.

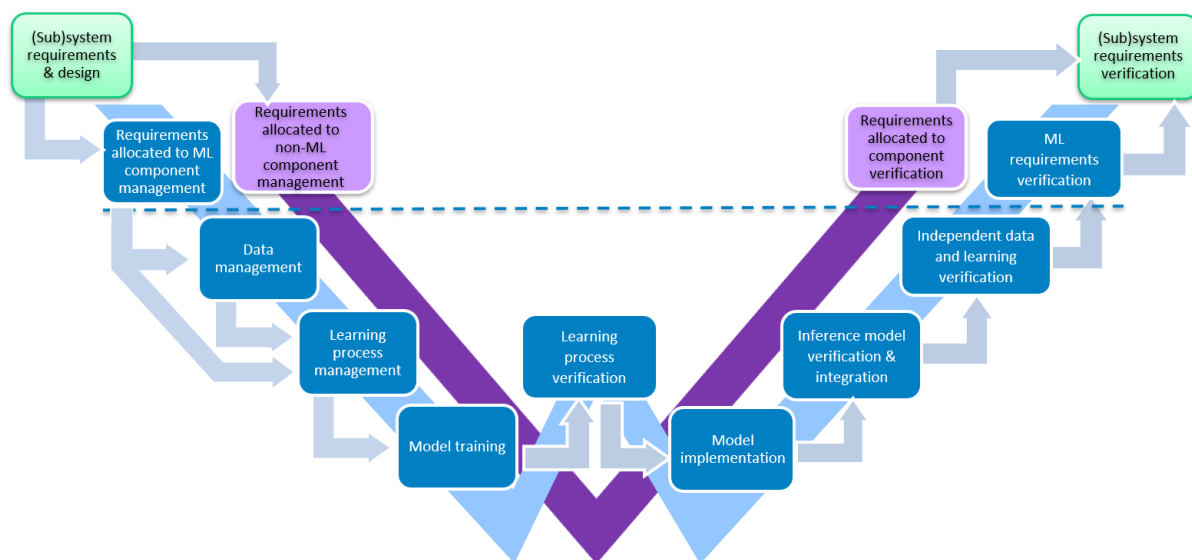


Figure 8 — Global view of learning assurance W-shaped process, non-AI/ML component V-cycle process and safety assessment process

This new learning assurance approach will have to account for the specific phases of learning processes, as well as to account for the highly iterative nature of certain phases of the process (orange and green arrows).

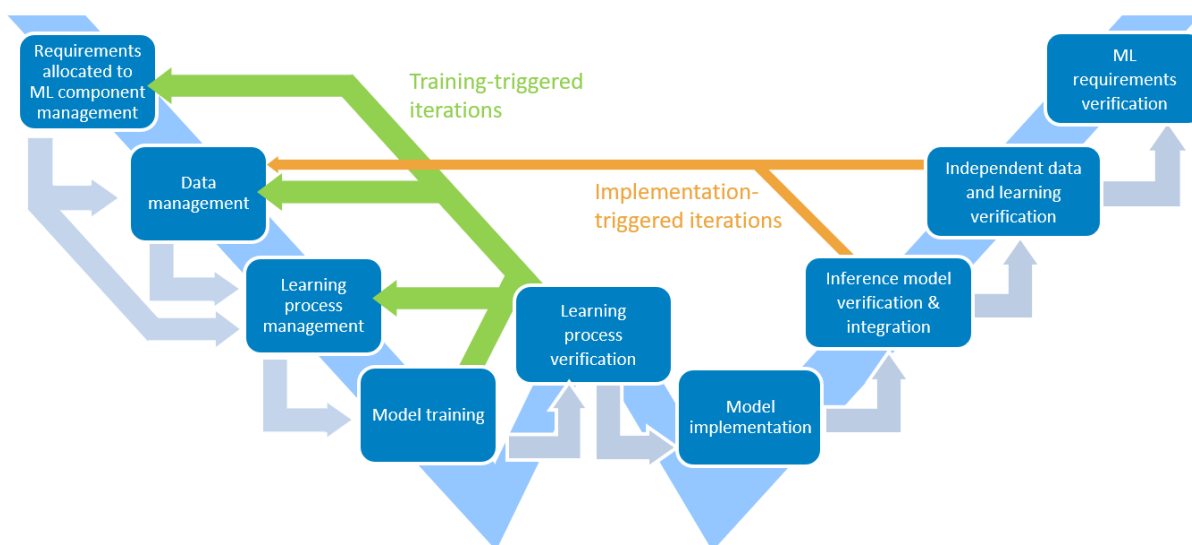


Figure 9 — Iterative nature of the learning assurance process

### 3.1. Learning assurance process planning

**Objective DA-01:** The applicant should describe the proposed learning assurance process, taking into account each of the steps described in Sections C.3.3 to C.3.12, as well as the interface and compatibility with development assurance processes.

Anticipated MOC: the set of plans should include a plan for learning assurance (PLAC), addressing all objectives from Section C.3 and detailing the proposed MOC.

### 3.2. Requirements and architecture management

The *requirements management* process covers the preparation of a complete set of requirements for the design of the *AI/ML component*. This step may be divided in several successive refinement steps and is preceded by a traditional flow-down of requirements (e.g. from aircraft to system for the *initial and continuing airworthiness* or *air operations* domains).

This step is further divided in:

- requirements capture;
- AI-based (sub)system architecture development<sup>10</sup>;
- requirements validation.

#### 3.2.1. Requirements capture

*Requirements capture* consists in the capture and unique identification of all requirements allocated to the AI/ML component, which are necessary to design and implement the ML component.

**Objective DA-02:** (Sub)systems requirements documents should be prepared and encompass the capture of the following minimum requirements:

- safety requirements allocated to the AI/ML component;
- information security requirements allocated to the AI/ML component;
- functional requirements allocated to the AI/ML component;
- operational requirements allocated to the AI/ML component, including ODD monitoring requirements;
- non-functional requirements allocated to the AI/ML component (e.g. performance, scalability, reliability, resilience, etc.).

#### 3.2.2. AI-based (sub)system architecture development

*AI-based (sub)system and components architecture development* is not a novel step compared to traditional systems development approaches; it is however an essential step in detailing the AI-based system, subsystem (if applicable) and AI/ML components architecture.

**Objective DA-03:** The applicant should describe the system and subsystem architecture, to serve as reference for related safety (support) assessment and learning assurance objectives.

<sup>10</sup> This step is different from the model architecture described in Section C.3.4.

### 3.2.3. Requirements validation

**Requirements validation** is considered to be covered by traditional system development methods. (e.g. ED-79A/ARP-4754A for product certification).

**Objective DA-04:** Each of the captured requirements should be validated.

## 3.3. Data management

The **data management** process is the first step of the data life cycle management.

It covers:

- data collection;
- data preparation (including labelling, pre-processing, data transformation and feature engineering);
- identification of the various data sets used in the training phase (typically the training, validation and test data sets);
- data sets validation and verification (including accuracy, completeness and representativeness, with respect to the ML requirements and ConOps);
- independence requirements between data sets;
- identification and elimination of the bias inherent to the data sets.

The data generated by the **data management** process is verified at each step of the process against the subset of data quality requirements (DQRs) pertaining to this step.

### 3.3.1. Data management requirements

The **data management** process will encompass its own requirements. In particular, some of these requirements will deal with DQRs.

**Objective DM-01:** The applicant should capture the DQRs for all data pertaining to the data management process, including but not limited to:

- the data needed to support the intended use;
- the ability to determine the origin of the data;
- the format, accuracy and resolution of the data;
- the traceability of the data from their origin to their final operation through the whole pipeline of operations;
- the mechanisms ensuring that the data will not be corrupted while stored or processed,
- the completeness and representativeness of the data sets; and
- the level of independence between the training, validation and test data sets.

**Anticipated MOC DM-01-1:** Starting from ED-76A Section 2.3.2 and accounting for specificities of data-driven learning processes, the DQRs should characterise the data by:

- the accuracy of the data;
- the resolution of the data;
- the confidence that the data has not been corrupted while stored, processed or transmitted (assurance level);
- the ability to determine the origin of the data (traceability);
- the level of confidence that the data is applicable to the period of intended use (timeliness);

- the data needed to support the intended use (completeness and representativeness), and
- the format of the data.

The MOC will need refinements based on progress in the standardisation (e.g. EUROCAE/SAE WG-114/G-34) and other best practices (e.g. reference: (DEEL Certification Workgroup, 2021)).

The **data management** process will also capture the requirements to be transferred to the implementation, regarding the pre-processing and feature engineering to be performed on the inference model.

**Objective DM-02:** The applicant should capture the requirements on data to be pre-processed and engineered for the inference model.

### 3.3.2. Data collection

The collection of data can be of different nature depending on the project (i.e. database, text, image, video, audio records); the applicant should always take into account that the data collected might fall under the category of *personal data* or affect *privacy*. In this case, there is a need to take into account Gear #3 of this Guidance since personal data requires a special protection.

The **data collection** should identify the different sources of data of relevance to the training.

**Objective DM-03:** To enable the data collection step, the applicant should identify explicitly and record the input space and the operating parameters that drive the selection of the training, validation and test data sets.

The sources of data are inherent to the AI/ML project. The sources can be internal or external to the applicant. External sources can be open-source or sourced via a contract to be established between the applicant and the data provider (e.g. weather data from a MET office, or databases shared between aeronautical organisations).

Depending on data sources, data sampling could be applied (simple random sampling, clustered sampling, stratified sampling, systematic sampling, multiphase sampling (reference: (DEEL Certification Workgroup, 2021))). The applicant should ensure completeness and representativeness on the sampling.

In order to address a lack of data completeness or representativeness, additional data may need to be collected or the existing data may be complemented with synthetic data (e.g. coming from models, digital twins, virtual sensors).

In the context of supervised learning techniques, the data set will need to be annotated or labelled.

**Objective DM-04:** Once data sources are collected, the applicant should make sure that the data set is correctly annotated or labelled.

All data items are annotated according to a specific set of annotation requirements, created, refined and reviewed by the applicant. Annotation can be a manual or automated process. Depending on the project, the annotation step can be effort-consuming (e.g. image annotations for detection purposes), and the applicant could decide to keep the annotation step insourced or outsourced, depending on its capabilities. In case of outsourcing of the activity, the applicant should decide on the DQRs to be achieved by the supplier.



### 3.3.3. Data preparation

The **data preparation** is paramount as it will be a key success factor for the ability of the AI/ML component to generalise. The **data preparation** is a multi-step process which involves a very significant part of the effort needed to implement an AI/ML component.

All operations on the data during **data preparation** should be performed in a way that ensures that the requirements on data are addressed properly.

**Objective DM-05:** The applicant should make sure that the operations on data properly address the captured requirements (including DQRs).

The main steps of the **data preparation** consist in:

- the pre-processing of the data, which is the act of cleaning and preparing the data for training;
- the feature engineering, aiming at defining the most effective input parameters from the data set to enable the training; and
- the data normalisation.

Note: Feature engineering does not apply to all ML techniques as depicted in Section C.3.3.3.2.

#### 3.3.3.1. Data pre-processing

The **data pre-processing** should consist in a set of basic operations on the data, preparing them for the **feature engineering** or the **learning process**.

**Objective DM-06:** The applicant should ensure a sequence of operations on the collected data in preparation of the training.

Depending on data sets, different aspects should be considered for cleaning and formatting the data:

- fixing up formats, typically for timestamp information, distances and temperatures;
- filling in missing values (e.g. some radar plot missing between different points on a trajectory); different strategies can apply in this case, either removing the corresponding row in the data set, or filling missing data (in general by inputting the mean value for the data in the data set);
- correcting erroneous values or standardising values (e.g. spelling mistakes, or language differences in textual data, cropping to remove irrelevant information from an image);
- identification and management of outliers (e.g. keeping or capping outliers, or sometimes removing them depending on their impact on the DQRs).

#### 3.3.3.2. Feature engineering

**Feature engineering** is a discipline consisting in transforming the pre-processed data so that it better represents the underlying structure of the data to be an input to the model training.

It is to be noted that **feature engineering** does not apply to all ML techniques. For example, many applications in computer vision use the feature extraction capabilities of a convolutional neural network, and do not apply any **feature engineering** step.

When **feature engineering** is applied, it should identify the relevant functional and operational parameters from the input space that are necessary to support the AI/ML model training.

**Objective DM-07:** When applicable, the applicant should transform the pre-processed data from the input space into features which are effective for the performance of the selected ML algorithm.

Considering the objective, and depending on the data in the input space, different technics could apply including:

- binning data (e.g. in computer vision, combining a cluster of pixels into one single pixel);
- breaking data into multiple parts (e.g. date in the year decomposed in week number and day of the week);
- consolidating data into features that better represent some patterns for the AI/ML model (e.g. positions and time transformed into speed, represent geospatial latitudes and longitudes in 3 dimensions then facilitate normalisation).

In relation with the objective, the applicant should limit the number of input variables, applying a dimensionality reduction step on the candidate features. This step aims at limiting the dimension of the feature space.

### 3.3.3.3. Data normalisation

**Objective DM-08:** The applicant should ensure that the data is effective for the stability of the model and the convergence of the learning process, possibly via normalisation.

**Anticipated MOC DM-08-1:** Depending on the data and the characteristics of the ODD, data normalisation could be achieved via different techniques:

- Min-Max Normalisation:  $X' = \frac{X - X_{min}}{X_{max} - X_{min}}$
- Mean Normalisation ( $X_{min}$  is replaced by the mean)
- Z Normalisation (Standardisation):  $X' = \frac{X - \mu}{\sigma}$ .

### 3.3.4. Data allocation

The **data allocation** corresponds to the last step of the **data preparation** process.

**Objective DM-09:** The applicant should distribute the data into three separate and independent data sets which will meet the specified DQRs:

- the training and validation data sets, used during the training and learning process validation;
- the test data set used during the learning process verification, and the inference model verification.

Particular attention should be given to the completeness and representativeness of each of the three data sets. Particular attention should be also paid to the independence of the data sets, in particular to that of the test data set.

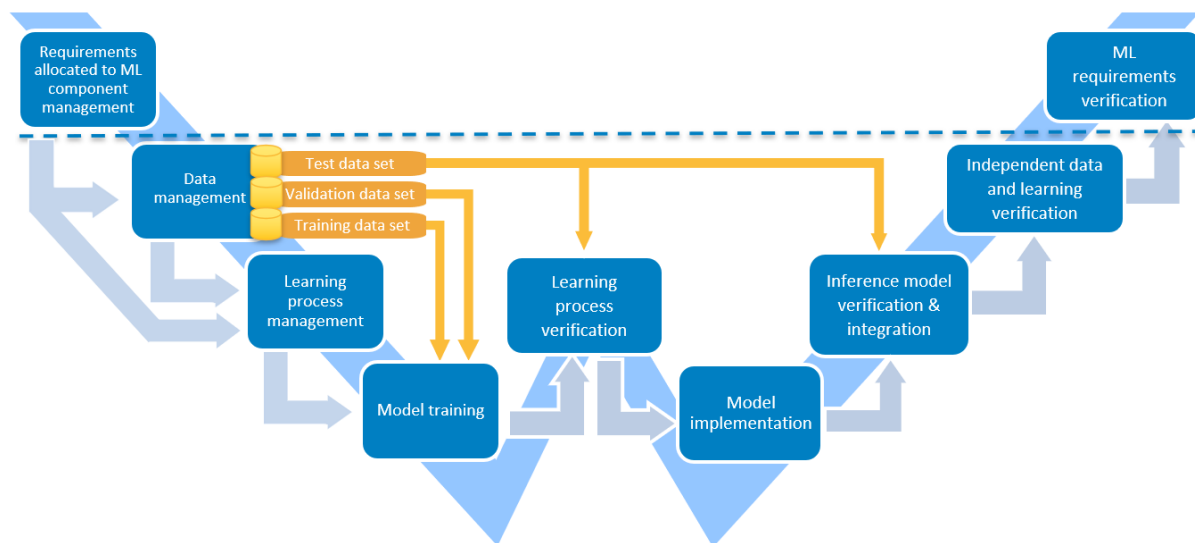


Figure 10 — Training, validation and test data sets usage in W-shaped cycle

### 3.3.5. Data validation

The **data validation** should be ensured all along the **data preparation** process, in order to provide the training with data aligned with the DQRs.

**Objective DM-10:** The applicant should ensure validation and verification of the data all along the data management process so that the DQRs are addressed.

The following represents a non-exhaustive list of anticipated MOC for a set of quality attributes which are expected for the data in the data set:

Completeness and representativeness of the data sets are prerequisites to ensure performance on unseen data and to derive generalisation guarantees for the trained model.

As surveyed in (Daedalean, 2020), two essential steps are considered: first, ensuring that the set of data covers the operational design domain sufficiently; second, that the data sets have been uniformly (according to the right distribution) and independently sampled from the input space.

#### Anticipated MOC DM-10-1: Data completeness

The data sets should be reviewed to evaluate their completeness with respect to the set of requirements and the ODD.

Various methods exist to assess the completeness of the data sets (training, validation or test). For example, the input space can be subdivided into a union of hyper-cubes whose dimensions are defined by the set of operating parameters, and the number of subdivisions for each dimension, by the granularity required for the associated operating parameter.

The completeness can be analysed through the number of points contained in the hypercubes.

The scalability of such an approach may be an issue and alternatives can be considered.

### **Anticipated MOC DM-10-2: Data representativeness**

There exist multiple methods to verify the representativeness of data sets according to a known or unknown distribution, stemming from the fields of statistics and ML.

To avoid the pitfalls of a posteriori justification or confirmation bias, it is important to first determine requirements to select and verify the chosen technique(s).

For parameters derived from operating parameters (e.g. altitude, time of day) or low-dimensional features from the data (e.g. image brightness), classical (parametric) methods apply to assess the goodness of fit of distributions.

However, considering only such parameters for high-dimensional spaces such as images might be too shallow, and techniques applying on images or other high-dimensional data might be necessary. For example, it is impossible to codify all possible sets of backgrounds on images.

There exist multiple methods adapted to high-dimensional data, sometimes by reducing to low-dimensional spaces. One of them is the distribution discriminator framework discussed in (Daedalean, 2020). A generic representativeness/completeness verification method is viewed as function  $D$  taking as input data sets, and returning a probability of them being in-distribution. Two opposite requirements must then hold:

- (1) The probability of  $D$  evaluated on in-distribution data sets is high.
- (2) The probability of  $D$  evaluated on out-of-distribution data sets is low.

The exact verification setting is to be determined depending on the required statistical significance and use case, but the framework remains method- and data-agnostic. Moreover, it is meant to allow easy verification as only in- or out-of-distribution (unannotated) data is required.

### **Anticipated MOC DM-10-3: Data accuracy, correctness**

In order to achieve correctness of the data, different types of errors and bias should be identified before bias in data sets is eliminated, and variance of data is controlled.

Errors and bias include:

- errors already present in the sourced data (e.g. data collected from databases or data lakes with residual errors or missing data);
- errors introduced by sensors (e.g. bias introduced by different cameras for the design and operational phases in the case of image recognition);
- errors introduced by collecting data from a single source;
- errors introduced by any sampling which could be applied during data collection from the data source;
- errors introduced by the human or tools when performing data cleaning or removal of presupposed outliers;
- annotation errors, especially when such an activity is performed manually by an annotation team.

#### Anticipated MOC DM-10-4: Data traceability

The applicant should establish an unambiguous traceability from the data sets to the source data, including intermediate data. Each operation should be shown to be reproducible.

Note: Traceability is of particular importance when data is cleaned during *data pre-processing* or is transformed as per the *feature engineering* activities.

#### Anticipated MOC DM-10-5: Data sets independence

The applicant should ensure that the training, validation and test data sets are verified against the independence requirement set in the DQRs.

Depending on the criticality of the AI application, more stringent requirements should be allocated to the independence of the test data set.

For highly critical applications, the applicant should ensure that the test data set is sampled independently from the training and validation data sets. The test data set should be ideally sampled from real data, complemented by synthetic data only where not avoidable (e.g. data at the limit or beyond flight envelope).

### 3.4. Learning process management

The *learning process management* considers the preparatory step of the formal *training phase*.

**Objective LM-01:** The applicant should describe the AI/ML components and model architecture (including computational graph and activation functions).

**Anticipated MOC LM-01-1:** The applicant should describe AI/ML components and model (computational graph) architecture in the planning documentation.

**Objective LM-02:** The applicant should capture the requirements pertaining to the learning management and training processes, including but not limited to:

- learning algorithm(s) selection;
- cost function selection describing the link to the performance and safety metrics;
- training environment (hardware and software) identification;
- model parameters initialisation strategy;
- hyper-parameters and parameters identification and setting;
- expected performance with training, validation and test data sets.

**Anticipated MOC LM-02-1:** The applicant should describe the selection and validation of key elements of the training process in the planning documentation.

**Objective LM-03:** The applicant should document the credit sought from the training environment and qualify the environment accordingly.

**Objective LM-04:** The applicant should provide quantifiable generalisation guarantees.

**Anticipated MOC LM-04-1:** The field of statistical learning theory (SLT) offers means to provide guarantees on the capability of generalisation of ML models. As introduced in the CoDANN report (Daedalean, 2020) Section 5.3.3, ensuring guarantees on the performance of a model on unseen data is one of the key goals of the field statistical learning theory. This is often related to obtaining ‘generalisation bounds’ or ‘measuring the generalisation gap’, that is the difference between the performance observed during development and the one that can be guaranteed during operations. The seminal work of Vapnik and Chervonenkis (On the Uniform Convergence of Relative Frequencies of Events to their Probabilities, 1971) established a relation of the generalisation capability of a learning algorithm with its hypothesis space complexity. Various forms of such bounds have been derived since then.

A good generalisation guarantee means that the ‘in-sample errors’ (i.e. the errors computed during the design phase) should be a good approximation of the ‘out-of-sample errors’ (i.e. the errors computed during the operations of the AI-based system). The generalisation gap of a model  $\hat{f}$  with respect to an error metric  $m$  and a data set  $D_{train}$  can be defined as:

$$G(\hat{f}, D_{train}) = |E_{out}(\hat{f}, m) - E_{in}(\hat{f}, D_{train}, m)|$$

where:

$\chi$  is the input space,  
 $E_{in}$  is the in – sample error (training error of the model) ,  
 $E_{out}$  is the out – of – sample error (expected operational error),  
 $D_{train}$  is the training dataset sampled from  $\chi$ ,  
 $\hat{f}$  is the model trained using  $D_{train}$ ,  
 $m$  is the error metric used to compute the errors.

A generalisation bound is by nature a probabilistic statement with the probability taken over possible data sets of a fixed size drawn from the input space  $\chi$ . Because of this, such bounds usually output a probability tolerance  $\delta$  for some given generalisation gap tolerance  $\varepsilon$ :

$$P_{D_{train} \sim \chi}(|E_{out}(\hat{f}, m) - E_{in}(\hat{f}, D_{train}, m)| < \varepsilon) > 1 - \delta$$

where:

$\varepsilon$  is the generalisation gap tolerance,  
 $\delta$  is the probability tolerance.

Such bounds can be dependent on the properties of the data or on the learning algorithm, with the amount of data typically tightening the generalisation gap and the model complexity loosening it. For example, VC dimension-based bounds give (in the above setting):

$$G(\hat{f}_V, D_{train}) < \sqrt{\frac{d_{vc} \cdot \log\left(\frac{2|D_{train}|}{d_{vc}}\right) + \log\left(\frac{1}{\delta}\right)}{|D_{train}|}}$$

where:

$d_{vc}$  is the VC – dimension of the model family

Other techniques like model compression can be used to reduce model complexity and also can help in obtaining stronger generalisation bounds (refer to (Stronger generalization bounds for deep nets via a compression approach., 2018)).

Based on the CoDANN report (Daedalean, 2020), it appears that, in the current state of knowledge, the values of the generalisation upper bounds obtained for large models (such as neural networks) are often too large without an unreasonable amount of training data. It is however not excluded that applicants could rely on such approaches with sharper bounds in a foreseeable future.

In the meantime, generalisation bounds not depending on model complexity can be obtained during the testing phase (refer to (Kenji Kawaguchi, 2018)). The drawback is that this requires the applicant to have a large test data set in addition to the training data set.

### 3.5. Training and learning process validation

The **training** consists primarily in applying the algorithm in the conditions defined in the previous step (typically an optimisation process for the weights of a defined architecture), using the training data set originating from the **data management** process step. Once trained, the model performance is evaluated, using the validation data set. Depending on the resulting performance, new training iteration is considered as necessary. The **training phase and its validation** can be repeated iteratively until the trained model reaches the expected performance.

**Objective LM-05:** The applicant should document the result of the model training.

**Anticipated MOC LM-05-1:** The records should include the training curves for the loss functions and for the error metrics.

The model performance with the validation data sets should also be recorded, linking this evaluation to the metrics defined under **Objective SA-01/03**.

**Objective LM-06:** The applicant should document any model optimisation that may affect the model behaviour (e.g. pruning, quantisation) and assess their impact on the model behaviour or performance.

**Anticipated MOC LM-06-1:** This step may need to be performed to anticipate the inference model implementation step (e.g. embedded hardware limitations). Any optimisation that can impact the behaviour of the model is to be addressed as part of the model training and validation step.

**Objective LM-07:** The applicant should estimate bias and variance of the selected model family and should provide evidence of the reproducibility of the training process.

**Anticipated MOC LM-07-1:** The bias-variance trade-off should be accounted for in the model family selection. The selection should aim for a model family whose complexity is high enough to minimise the bias, but not too high to avoid high variance, in order to ensure reproducibility.

The bias and variance linked to a given model family are not easy to evaluate. The applicant should identify methods to provide the best possible estimates of the bias and variance of the selected model family; for instance, using random resampling methods (e.g. 'Bootstrapping' or 'Jack-knife').



Regularisation is a typical method to avoid overfitting (high variance) with complex models like neural networks.

**Objective LM-08:** The applicant should ensure that the estimated bias and variance of the selected model meet the associated learning process management requirements.

**Anticipated MOC LM-08-1:** For the selected model, bias is measured as the mean of the ‘in sample error’ ( $E_{in}$ ), and variance is measured by the statistical variance of the ‘in sample error’ ( $E_{in}$ ).

The applicant should analyse the errors on the training dataset to identify and mitigate systematic errors.

### 3.6. Learning process verification

The **learning process verification** consists then in the evaluation of the trained model performance evaluation using the test data set. Any shortcoming in the model quality can lead to the need to iterate again on the data management process step or learning process management step, e.g. by correcting or augmenting the data set, or updating learning process settings. It is important to note that such an iteration may invalidate the test data set and lead to the need to create a new independent test data set.

**Objective LM-09:** The applicant should perform an evaluation of the performance of the trained model based on the test data set and document the result of the model verification.

**Anticipated MOC LM-09-1:** The final performance with the test data set should be measured and fed back to the safety assessment process, linking this evaluation to the metrics defined under the **Objective SA-01/03** and explaining any divergence in the metrics compared to the ones used to fulfil **Objective LM-04**.

**Objective LM-10:** The applicant should perform a requirements-based verification of the trained model behaviour and document the coverage of the ML component requirements by verification methods.

**Anticipated MOC LM-10-1:** Requirements-based testing and robustness testing methods are recommended to reach this objective. They consist in a systematic testing of both extrinsic (outputs) and intrinsic (architecture-dependent) behaviours of the model.

**Objective LM-11:** The applicant should provide an analysis on the robustness (or stability) of the algorithms and of the trained model.

**Anticipated MOC LM-11-1:** As outlined in (Daedalean, 2020) Section 6.4.1, there are (at least) two sources of instability which should be analysed:

- Perturbations in the design phase due to fluctuations in the training data set (e.g. replacement of data points, additive noise or labelling errors). This analysis supports the demonstration of the learning algorithm stability.
- Perturbations in the operational phase due to fluctuations in the data input and prediction output or in the model itself. This analysis supports the demonstration of the model stability.

The definition of abnormal range tests and adversarial tests is a necessary means to reach this objective. ‘Abnormal range’ is linked to defining tests at and out of the limits of the requirements-based tests range. ‘Adversarial testing’ consists in defining corner cases (not based on the requirements) that may affect the AI/ML component expected behaviour.

Also, the use of formal methods is anticipated to be a promising means of compliance with this objective, although in the current state of research, those methods appear to be limited to local evaluations.

**Objective LM-12:** The applicant should verify the anticipated generalisation bounds using the test data set.

**Anticipated MOC LM-12-1:** Evidence of the validity of the anticipated generalisation guarantees proposed to fulfil **Objective LM-04** should be recorded.

### 3.7. Trained model implementation

The implementation phase starts with the *requirements capture*.

**Objective IMP-01:** The applicant should capture the requirements pertaining to the implementation process, including but not limited to:

- data processing requirements originating from the data management process (C.3.3.1);
- conversion of the model to be compatible with the inference platform;
- optimisation of the model to adapt to the inference platform resources;
- deployment of the model into software and/or hardware items.

The *implementation* then consists in transforming the training model into an executable model that can run on a target hardware. This implementation follows different steps, at each of which transformations to the trained model might be performed:

- Model conversion
- Model optimisation
- Model deployment

**Objective IMP-02:** Any post-training model transformation (conversion, optimisation, deployment) should be identified and validated for its impact on the model behaviour and performance.

**Objective IMP-03:** For each transformation step, the environment (i.e. software tools and hardware) necessary to perform model transformation should be identified and any associated assumptions or limitations captured and validated.

**Anticipated MOC IMP-03-1:** The applicant should describe the environment for each transformation step.

### 3.7.1. Trained model conversion

One of the first activities after the learning process is the freezing of the model. The trained model is represented in formats specific to the framework on which it is trained. This conversion needs to be applied to the trained model in order to obtain a representation that is compatible with the inference platform. This step is the procedure of removing graph components that are not required during inference, as well as making changes that reduce the graph size and complexity without impacting the model behaviour and performance.

For example, since weights will not be updated anymore after training, gradients can be safely removed, the weight variables turned into constants and any other metadata that is relevant for training deleted. The result is a subset of the original training graph, where only the graph components that are required by the inference environment are kept, as captured in the set of requirements pertaining to implementation allocated to the AI/ML component.

Another conversion activity is the conversion of the model into an open format. The format in which frozen models are saved and restored is likely to be different between the learning and inference environment essentially due to the difference of framework.

**Anticipated MOC IMP-02-1:** Identification of the different conversion steps and confirmation that no impact on the model behaviour is foreseen.

### 3.7.2. Trained model optimisation

In the scope of the implementation, any transformation that does not affect the behaviour or performance of the model, is considered as optimisation, by opposition to the ones that can be performed during the training phase.

A list of possible implementation optimisations includes:

- Winograd algorithms for convolution: these algorithms are targeting high-performance inference. Their efficiency comes from the reduction of the number of multiplication operations due to linear and Fourier transforms.

**Anticipated MOC IMP-02-2:** Identification of the different optimisation steps performed during implementation and confirmation that no impact on the model behaviour is foreseen.

### 3.7.3. Trained model deployment inference environment

Once confirmed that the transformations of the model had no impact, the last step that could impact its behaviour or performance is its implementation in software and/or hardware items.

**Anticipated MOC IMP-02-3:**

- For software aspects, it is anticipated that the provisions of applicable software development assurance guidance (e.g. AMC 20-115D for product certification projects) would provide the necessary means to confirm that **Objective IMP-02** is fulfilled. This guidance may need to be complemented to address specific issues linked to the implementation of an AI/ML model into software, such as memory management issues.
- For hardware aspects, it is anticipated that the provisions of applicable hardware development assurance guidance (e.g. AMC 20-152A for product certification projects) would provide the

necessary means to confirm that **Objective IMP-02** is fulfilled. FPGAs, ASICs and COTS architectures are covered by the existing guidance; however, other ML architectures, such as graphics processing units (GPUs), have specificities that are not accounted for in the existing guidance (e.g. very complex interference mechanisms or non-deterministic pipelining).

- For multicore processors (MCP) aspects, it is anticipated that the provisions of applicable MCP development assurance guidance (e.g. AMC 20-193<sup>11</sup> for product certification projects) would provide the necessary means to confirm that **Objective IMP-02** is fulfilled.

### 3.8. Inference model verification

The **inference model verification** aims at verifying that the inference model behaves adequately compared to the trained model, in evaluating the model performance with the test data set, explaining any difference in the evaluation metric compared to the one used in the **training phase verification** (e.g. execution time metrics). This process step also should foresee verification that the model properties have been preserved (e.g. based on the implementation analysis or through the use of formal methods). Finally, it includes typical software verification steps (e.g. memory/stack usage, WCET, etc.) that could be strictly conventional (e.g. per DO-178C/ED-12C) but for which any specificity linked to the learning approach should be identified and managed.

The verification of the embedded inference model implies several steps of integration, as many as considered necessary to support adequate verification; an important one being the integration of the 'implemented' AI/ML component on the target hardware, together with the other AI-based subsystem components.

#### 3.8.1. Verification of inference model properties preservation

**Objective IMP-04:** The applicant should verify that any conversion, optimisation or transformation performed during the trained model implementation step have not altered the defined model properties.

**Anticipated MOC IMP-04-1:** First a set of model properties should be captured. Then the use of specific verification methods (e.g. formal methods) is expected to be necessary to comply with this objective.

#### 3.8.2. Hardware verification

**Objective IMP-05:** The differences between the hardware platform used for training and the one used for verification should be identified and assessed for impact on the inference model behaviour and performance.

**Anticipated MOC IMP-05-1:** The analysis of the differences, such as arithmetic precision, is an important means to reach this objective.

<sup>11</sup> AMC 20-193 is not yet published at the time of release of this document.

### 3.8.3. Inference model integration and verification

**Objective IMP-06:** The applicant should perform an evaluation of the performance of the inference model based on the test data set and document the result of the model verification.

**Anticipated MOC IMP-06-1:** The final performance with the test data set should be measured and fed back to the safety assessment process, linking this evaluation to the metrics defined under the **Objective SA-01/03** and explaining any divergence in the metrics compared to the ones used to fulfil **Objective LM-09**.

**Objective IMP-07:** The applicant should perform a requirements-based verification of the inference model behaviour and document the coverage of the ML component requirements by verification methods.

**Anticipated MOC IMP-07-1:** Requirements-based testing and robustness testing methods are necessary to reach this objective. They consist in a systematic testing of both extrinsic (outputs) and intrinsic (architecture-dependent) behaviours of the model.

The test environments should at least foresee the level of integration of the inference model on the target hardware platform (environment #1) and the integration of the ML component in the subsystem with representative interfaces to the other subsystems, in particular to the sensors (environment #2).

**Objective IMP-08:** The applicant should provide an analysis on the robustness (or stability) of the inference model.

**Anticipated MOC IMP-08-1:** The definition of abnormal range tests and adversarial tests is a necessary means to reach this objective. 'Abnormal range' is linked to defining tests at and out of the limits of the requirements-based tests range. 'Adversarial testing' consists in defining corner cases (not based on the requirements) that may affect the AI/ML component expected behaviour.

### 3.9. Independent data and learning verification

The **data verification** step is meant to close the data management life cycle, by verifying with independence that data sets were adequately managed, considering that the verification of the data sets can be achieved only once the inference model has been satisfactorily verified on the target hardware. It is important to mention however that this does not imply waiting for the end of the process to initiate this step, considering the highly iterative nature of learning processes.

**Objective DM-14:** The applicant should perform a data and learning verification step to confirm that the appropriate data sets have been used for the training, validation and verification of the model and that the expected guarantees (generalisation, robustness) on the model have been reached.

**Anticipated MOC IMP-08-1:** The associated activities include:

- independent verification that the data sets (training, validation, verification) comply with the data management requirements;
- independent verification of the correct identification of the input space, including a reassessment of the defined ODD;

- independent verification that the data sets (training, validation, verification) are complete and representative of the input space of the application;
- independent verification that the expected guarantees (generalisation, robustness) on the model have been reached.

Note: This independent verification step may be requested only for higher-criticality levels.

### 3.10. Requirements verification

The **requirements verification** is addressing the verification of the AI/ML component fully integrated in the overall system. It is considered to be covered by traditional assurance methodologies (e.g. ED-79A/ARP4754A).

**Objective DA-05:** Each of the captured (sub)system requirements should be verified.

### 3.11. Configuration management

The **configuration management** is an integral process to the development of an AI-based (sub)system.

**Objective CM-01:** The applicant should apply all configuration management principles to the AI-based (sub)system life-cycle data, including but not limited to:

- identification of configuration items;
- versioning;
- baselining;
- change control;
- problem reporting;
- archiving and retrieval.

### 3.12. Quality and process assurance

**Quality and process assurance** is an integral process that aims at ensuring that the life-cycle process objectives are met, and the activities have been completed as outlined in plans (as per **Objective DA-01**) or that deviations have been addressed.

**Objective QA-01:** The applicant should ensure that quality/process assurance principles are applied to the development of the AI-based (sub)system, with the required independence level.

## 4. AI explainability

The explainability building block is designed with the necessary objective of maintaining trust of all stakeholders involved during the whole life cycle of the AI-based system, thus also contributing to safe operations. To this purpose, the guidance introduced in this chapter covers the development, operations and post-operation phases of the life cycle.

Specifically, for the operations phase, the introduction of AI is expected to modify the paradigm of interaction between the operator and the AI-based systems. More specifically, it will affect the function allocation distribution<sup>12</sup> between the operator and the AI-based system, by progressively giving more authority to the AI-based systems. The corollary of increasing the levels of automation of the AI-based systems will be a reduction of operator awareness of the logics leading to the automated decisions or actions.

This decreasing awareness may potentially limit the efficiency of the interaction and lead to a reduction of trust from the operator. Trust is considered to be essential and critical to the general acceptability of the AI-based systems.

Therefore, in the context of operations, in order to maintain the trust and ensure an adequate efficiency of the interaction, there will be a need for the AI-based systems to artificially provide explanations with regard to their decisions and actions.

As a starting point for the guidelines, there is a need to define more precisely what AI explainability is. Starting from the EASA AI Roadmap text with some refinement.

### Definition

***AI explainability:*** Capability to provide the human with understandable and relevant information on how an AI/ML application is coming to its results.

Knowing that the field of AI research is very active on this topic, this definition will evolve as needed in the upcoming updates of this set of first guidelines.

### 4.1. AI explainability scope

The scope of AI explainability is covering both the development/post-operation of the AI-based systems and the operation of the AI-based systems.

The figure below illustrates the scope of application and the need for explainability.

<sup>12</sup> Function allocation distribution refers to strategies for distributing system functions and tasks across people and technology. (Function Allocation Considerations in the Era of Human Autonomy Teaming, December 2019).



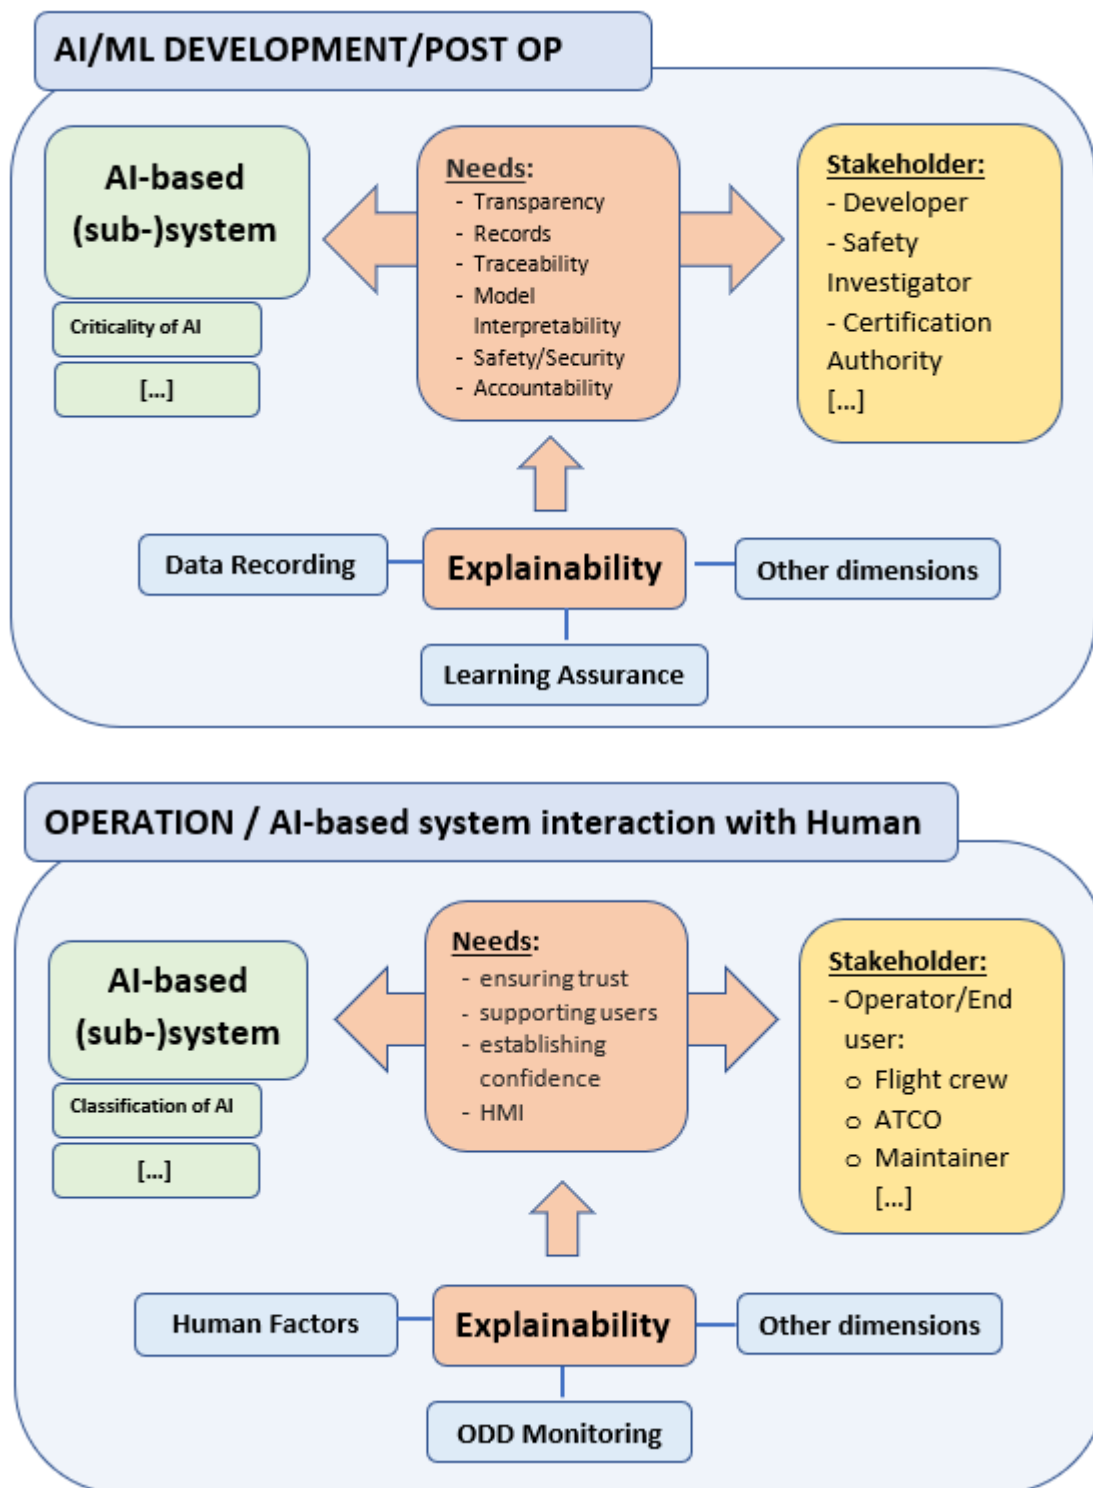


Figure 11 — AI explainability scope

The contributors to the AI-based system operation and to the AI-based system learning process building blocks are detailed in Section C.4.4.

## 4.2. AI dimensions and level of explainability

A lot of dimensions have to be taken into account in order to adequately design the acceptability of an AI-based (sub)system. A first review of the literature allows the identification of the following preliminary list of questions:

- Why should an AI-based (sub)system be explainable? Why does the operator require explanation?
- What explanation can/should be generated?
- When the information should be presented?
- How much details should be presented in an explanation?

The different dimensions are currently under study and the proposed list might evolve in accordance with the results.

### 4.2.1. Object of the explanation

One of the dimensions to be considered is the object of the explanation. The object could be either:

- the ML item itself (a priori/global explanation);
- an output of the ML item (post hoc/a posteriori/local explanation).

Naturally, one can expect that an important part of the explanation of the ML item itself (and the performance thereof) comes from the adoption of the learning assurance, namely following the W-shaped process. In other words, the application of correct data management practices, training and training verification, implementation, and inference model verification provides a mathematical proof of the correctness of the ML item. This proof is human-understandable given that the data and learning assurances are. However, this might not be fully satisfactory with respect to other dimensions, and might need to be complemented with other considerations.

While ML item-level explanations yield output-level ones, it is still desirable to have separate output-level explanations, since they can usually be simpler/more transparent, and then accessible to non-AI/ML experts like the end-user communities.

It can be anticipated that explanations at ML item level would be more dedicated to the stakeholders involved during development and post operations, whereas explanations on the output of the ML item would be more useful for the end users in the operations.

Explainability at ML item level or at ML output level are a focused area for many research groups, and it is expected that best practices and techniques will emerge soon. This emergence will enable new objectives to be captured as per these guidelines.

### 4.2.2. Levels of explainability

On top of the above, it will be important to consider the various possible levels of the AI-based systems. Indeed, the needs for explainability will be significantly dependent on the pattern of authority and function allocation distribution between the operator and the AI-based (sub)system.

The operation of the Level 1A AI-based systems will not be fundamentally different from the operation of existing systems. Therefore, there is no need to develop specific explainability mechanisms on top of the existing human factors requirements and/or guidance that are already in use (e.g. CS/AMC 25.1302 for cockpit design).

However, from level 1B and above, there will be a need to identify and characterise the need for explainability as well as its attributes. This will require the development of some disruptive material.

This following figure illustrates the relationship between AI-based (sub)system and the expected level of explainability needed **during operation**:

	<b>EASA AI Roadmap — AI Level (subsystem)</b>				
	<b>Level 1</b>		<b>Level 2</b>	<b>Level 3</b>	
	<b>Level 1A</b> <b>Human augmentation</b>	<b>Level 1B</b> <b>Human assistance</b>		<b>Level 3A</b> <b>More autonomous AI</b>	<b>Level 3B</b> <b>Complete autonomous AI</b>
Level of change in human-machine interaction — Expected impact of AI level on the operation.	<b>None</b> - There is no expected impact on the known operation.	<b>Medium</b> - Introduction of cognitive assistance	<b>Very High</b> - Introduction of human-machine teaming. AI/ML systems capable of taking decisions.	<b>High</b> - AI/ML system is taking over more tasks in autonomy. The operator is in full control of the final decision.	No more human-machine interaction
Expected level of explainability needed during operation	<b>None</b>	<b>Important</b>  The role of the explainability is there to help facilitate improved end user decisions. At this level, decision still requires human judgement or some agreement on the solution method.	<b>Critical or beneficial</b>  It will highly depend on the degree of the AI/ML system's autonomy and on the level of decision taken by the operator.	<b>Critical</b>  In order for the operator to override the AI/ML systems' decision, explainability is considered to be critical for the good operation of the system.	(Under development)
Need for specific guidelines for design for explainability	<b>None</b> Existing guidelines and requirements for interface design should be used. (e.g. CS/AMC 25.1302)	<b>Specific guidance needed</b>  (Under development)	<b>Specific guidance needed</b>  (Under development)	<b>Specific guidance needed</b>  (Under development)	(Under development)

### 4.3. AI Explainability motivations and objectives

Considering the different roles and responsibilities of humans that we attempt to cover with the future guidelines, the motivations to get explanations can be various. An initial list of needs can be found in DEEL's white paper (DEEL Certification Workgroup, 2021).

This list has been reworked to better match the objectives selected for the explainability:

Development/post-operation	Operation
<ul style="list-style-type: none"> <li>ensuring trust</li> <li>establishing causal relationships between the input and the output of the model</li> <li>catching the boundaries of the model and help in its fixing</li> <li>highlighting undesirable bias (data sets and model bias)</li> <li>allowing the relevant receivers identify errors in the model</li> <li>enable recording of relevant data to support continuous analysis of the AI-based system behaviour.</li> </ul>	<ul style="list-style-type: none"> <li>ensuring trust</li> <li>supporting its use by providing information about how the decisions have been taken</li> <li>establishing how confident a model is with its own decisions</li> </ul>

With the current state of knowledge, EASA considers AI explainability as a separate set of top-level objectives that are instrumental in supporting the human to safely perform tasks and to foster trust when interacting with the AI-based (sub)system, whatever their role.

Note: These objectives will be improved and complemented based on future work on this building block.

**Objective EXP-01:** The applicant should identify the list of humans that are intended to interact with the AI-based (sub)system, at any stage of its life cycle, together with their roles, their responsibilities and their expected expertise (including assumptions made on the level of training, qualification and skills).

**Anticipated MOC EXP-01-1:** The list of humans, roles, responsibilities and expected expertise should be included in the plan for aspects of learning assurance.

**Objective EXP-02:** For each role, the applicant should identify which task(s) the humans are intended to perform in interaction with the AI-based (sub)system, as well as the task allocation pattern.

**Anticipated MOC EXP-02-1:** The level at which the tasks are identified should be considered at the level of the interaction between the human and the AI-based system, not at the level of each single task performed by the AI-based subsystem or component. The list of task(s) relevant to the AI-based (sub)system explainability should be documented.

**Objective EXP-03:** For each output of the (sub)system relevant to the task(s), the applicant should identify the need for an explanation depending on several criteria (including nature of the task, authority of the human and the level of AI-based system) and specify the set of necessary explanations to be provided to the human.

**Objective EXP-04:** For each output relevant to the task(s), the applicant should ensure the validity of the specified explanation, based on actual measurements (e.g. monitoring) or on a quantification of the level of uncertainty.

## 4.4. Contributors to the AI-based systems

### 4.4.1. Learning assurance

Learning assurance is a prerequisite to ensure confidence in the performance and intended function of the ML-based system. Without this confidence, no AI explainability is practicable, therefore learning assurance is considered to be one of the fundamentals for the explainability guidelines.

In complement of the guidelines already provided in Section C.3, some additional MOC linked to learning assurance are anticipated to be of importance to fulfil the overall objective of explainability.

**Anticipated MOC EX-03-1:** Considerations on trained model interpretability.

**Anticipated MOC EX-03-2:** Identification of the trained features relevant to the explanation objective.

### 4.4.2. ODD monitoring in AI-based system operation

Learning assurance guarantees are given in the frame of the defined ODD (e.g. requirements on data completeness and representativeness in Section C.3.3.1) and one important objective is to monitor whether or not the operational conditions remain within acceptable boundaries. The feedback of this monitoring is considered to be another fundament of the explainability guidelines. The following MOC is anticipated.

**Anticipated MOC EX-03-3:** The AI-based system inputs should be monitored to be within the operational boundaries in which the AI/ML component performance is guaranteed, and deviations should be indicated to the human.

**Anticipated MOC EX-03-4:** The training and instructions available for the human should include procedures to act on the possible outputs of the ODD monitoring.

**Anticipated MOC EX-03-5:** Information concerning unsafe system operating conditions should be provided to the human operator/end user to enable them to take appropriate corrective action in a timely manner.

### 4.4.3. AI data recording capability

**Objective EX-05:** The applicant should provide the means to record operational data that is necessary to explain the behaviour of the AI-based (sub)system post operations.

With regard to the recording of data for the purpose of post-operation assessment, at least two distinct types of use should be addressed:

- **Data recording for the purpose of monitoring the safety of AI-based system operations (as part of safety management and/or continued operation approval)**
  - This monitoring consists in recording and processing data from day-to-day operation to detect and evaluate deviations from the expected behaviour of the AI-based system, as well as issues affecting interactions with human users or other systems.
  - This monitoring is usually performed by (or on behalf of) the organisation using the AI-based system.

- The purpose of this monitoring is to support the continuous or frequent assessment of the safety of the operations in which the AI-based system is used and to assess whether mitigation actions are effective.
  - This monitoring is meant to be part of the safety management system (SMS) of the organisation using the AI-based system.
  - 'This monitoring may also serve the purpose of continued operation approval, by providing the designers of the AI-based system with data to monitor the in-service performance of the system.
- **Data recording for the purpose of accident or incident investigation**
- This recording is meant for analysing an accident or incident for which the operation of the AI-based system could have been a contributing factor.
  - There are many kinds of accident or incident investigations (internal investigation, judicial investigation, assurance investigation, etc.) but in this document, only the official safety investigation (such as defined in ICAO Annex 13 and Regulation (EU) 996/2010) is considered. An official safety investigation aims at preventing future incidents and accidents, not at establishing responsibilities of individuals.
  - The recorded data is used to accurately reconstruct the sequence of events that resulted in the accident or serious incident.

Notes:

- It is not forbidden to address these two types of use with a single data recording solution.
- The recording of data does not need to be a capability of the AI-based system. It is often preferable that the relevant data is output for recording to a dedicated recording system.

**Anticipated MOC EX-05-1 (applicable to both types of use):**

The recording capability should automatically start recording before or when the AI-based system is operating, and it should keep recording until the AI-based system is no longer operating. The recording capability should automatically stop recording when or after the AI-based system is no longer operating.

#### 4.4.3.1. Data recording for the purpose of monitoring the safety of AI-based system operations

**Anticipated MOC EX-05-2:**

The recorded data should contain sufficient information to detect deviations from the expected behaviour of the AI-based system, whether it operated alone or interacting with an end user. In addition, this information should be sufficient:

- (a) to accurately determine the nature of each individual deviation, its time and the amplitude/severity of that individual deviation (when applicable);
- (b) for the organisation using the AI-based system, and when applicable, the end users of the AI-based system, to determine the particular conditions or events likely to have triggered the individual deviation;
- (c) for monitoring trends regarding deviations over longer periods of time.

**Anticipated MOC EX-05-3:**

The recorded data should be made available to those entitled to access and use it in a way so that they can perform an effective monitoring of the safety of AI-based system operations. This includes:

- (a) timely and complete access to the data needed for that purpose;
- (b) access to the tools and documentation necessary to convert the recorded data in a format that is understandable and appropriate for human analysis;
- (c) possibility to gather data over longer periods of time for trend analyses and statistical studies. In any case, the data should be retained for a minimum of 30 days.

**4.4.3.2. Data recording for the purpose of accident or incident investigation**

**Anticipated MOC EX-05-4:**

The recorded data should contain sufficient information to accurately reconstruct the operation of the AI-based system before an accident or incident. In particular, this information should be sufficient to:

- (a) accurately reconstruct the chronological sequence of inputs to and outputs from the AI-based system;
- (b) determine in this chronological sequence, which inputs and outputs most likely contributed to the accident or incident;
- (c) identify any unexpected behaviour of the AI-based system that is relevant for explaining the accident or incident; and
- (d) determine the causes of each unexpected behaviour of the AI-based system, either directly or by running simulations that accurately reproduce the unexpected behaviour.

**Anticipated MOC EX-05-5:**

The data should be recorded in a way so that it can be retrieved and used after an accident or an incident. This includes:

- (a) crashworthiness of the memory media if they could be exposed to severe environmental conditions resulting from an accident;
- (b) means to facilitate the retrieval of the data after an accident, if applicable (e.g. means to locate the accident scene and the memory media, tools to retrieve data from damaged memory media);
- (c) provisions of tools and documentation necessary to convert the recorded data in a format that is understandable and appropriate for human analysis; and
- (d) The data relevant for incident or accident investigation should be retained for a minimum of 60 days.

**4.5. Human factors guidance for explainability**

For Level 1A, existing guidelines and requirements for interface design should be used.

For Level 1B and above, the guidance is under development.



## 5. AI safety risk mitigation

### 5.1. AI safety risk mitigation concept

AI safety risk mitigation (SRM) is based on the anticipation that the ‘AI black box’ may not always be opened to a sufficient extent. Indeed, for some applications, it could be unpractical to fully cover all the objectives defined in the explainability and learning assurance building blocks of this guideline. This partial coverage of some objectives would result in a residual safety risk that may be accommodated by implementing some mitigations called hereafter SRM.

Furthermore, it is also recognised that the use of AI in the aviation domain is quite novel and until field service experience is gained, appropriate safety precautions should be implemented to reduce the risk to occupants, third parties and critical infrastructure.

This could be achieved by several means, among others:

- real-time monitoring of the output of the AI/ML component and passivation of the AI-based system with recovery through a traditional backup system (e.g. safety net);
- in a wider horizon by considering the notion of ‘licensing’ to an AI, as anticipated in (Javier Nuñez et al., 2019) and developed further in (ECATA Group, 2019).

Note that SRM is solely meant to address a partial coverage of the applicable explainability and learning assurance objectives. SRM is not aimed at compensating partial coverage of objectives belonging to the trustworthiness building blocks (e.g. safety assessment, cybersecurity, ethical objectives).

### 5.2. AI SRM top-level objectives

**Objective SRM-01:** Once activities associated with all other building blocks are defined, the applicant should determine whether the coverage of the objectives associated with the explainability and learning assurance building blocks is sufficient or if an additional dedicated layer of protection, called hereafter safety risk mitigation (SRM), would be necessary to mitigate the residual safety risks to an acceptable level.

**Anticipated MOC SRM-01:** In establishing whether AI SRM is necessary and to which extent, the following considerations should be accounted for:

- coverage of the explainability building block;
- coverage of the learning assurance building block;
- relevant in-service experience, if any;
- AI-level: the higher the level, the more likely SRM will be needed;
- criticality of the AI/ML component: The more the ML/AI component is involved in critical functions, the more likely SRM will be needed.

In particular, the qualitative nature of some building block mitigations/analysis should be reviewed to establish the need for an SRM.

The SRM strategy should be commensurate with the residual risk/unknown.

**Objective SRM-02:** The applicant should establish SRM means as identified in Objective SRM-01.

**Anticipated MOC SRM-02-1:** The following means may be used to gain confidence that the safety risk is properly mitigated:

- monitoring of the output of the AI/ML component and passivation of the AI-based system with recovery through a traditional backup system (e.g. safety net);
- when relevant, the possibility may be given to the operator to switch off the AI/ML-based function to avoid being distracted by erroneous outputs.

Note: The SRM functions should be evaluated as part of the safety assessment<sup>13</sup>, and, if necessary, appropriate safety requirements should be defined and verified.

---

<sup>13</sup> In the ATM/ANS domain, for non-ATS providers, the safety assessment is replaced by a safety support assessment.

## 6. Organisations

Prior to obtaining approval of AI applications in the field of civil aviation, organisations that are required to be approved as per the Basic Regulation (Regulation (EU) 2018/1139 might need to introduce adaptations in order to ensure the adequate capability to meet the objectives defined within the AI trustworthiness building blocks (see Figure 2 — EASA AI trustworthiness roadmap building blocks), and to maintain the compliance of the organisation with the corresponding implementing rules.

The introduction of the necessary changes to the organisation would need to follow the process established by the applicable regulations. For example, in the domain of initial airworthiness, the holder of a DOA would need to apply to EASA for a significant change to its design assurance system prior to the application for the certification project.

This chapter introduces some high-level provisions and anticipated AMC with the aim of providing guidance to organisations on the expected adaptations. It provides as well, as an example case, more detailed guidance on the affected processes for holders of a DOA.

### 6.1. High-level provisions and anticipated AMC

**Provision ORG-01:** The organisation should review its processes and adapt them to the introduction of AI technology.

**Provision ORG-02:** Implement a data-driven ‘AI continuous safety assessment system’ based on operational data and in-service events.

#### **Anticipated AMC ORG-02:**

The AI continuous safety assessment system should:

- ensure data gathering on safety-relevant areas for AI-based systems;
- perform analyses to support the identification of in-service risks, based on:
  - the organisation scope;
  - a set of safety-related metrics;
  - available relevant data.

The system should be able to refine the identification of risks based on the results of previous interactions with the AI-based systems and incorporating the human evaluation inputs.

When defining the metrics, the data set and gathering methodology should ensure:

- the acquisition of safety-relevant data related to accidents and incidents including near-miss events; and
- the monitoring of in-service data to detect potential issues or sub-optimal performance trends that might contribute to safety margin erosion; and
- the definition of target values, thresholds and evaluation periods; and
- the possibility to analyse data to determine the possible root cause and trigger corrective actions.

The following implementing rule requirements, associated AMC and GM may be considered with appropriate adaptations:

For ATS providers:

- ATS.OR.200(2) and (3) ‘Safety management system’
- GM1 ATS.OR.200(3)(i) and GM1 ATS.OR.200(3)(iii)
- AMC1 ATS.OR.200(3)(iii)

## 6.2. Design organisation case

This section aims to provide an example, for the case of DOA holders by identifying those processes that might need to be assessed and adapted. Some aspects are not yet covered by published EU regulations, like information security or SMS, but it is mentioned for completeness.

The following figure illustrates the potentially affected DOA processes and the key activities in relation to the implementation of AI/ML technologies:

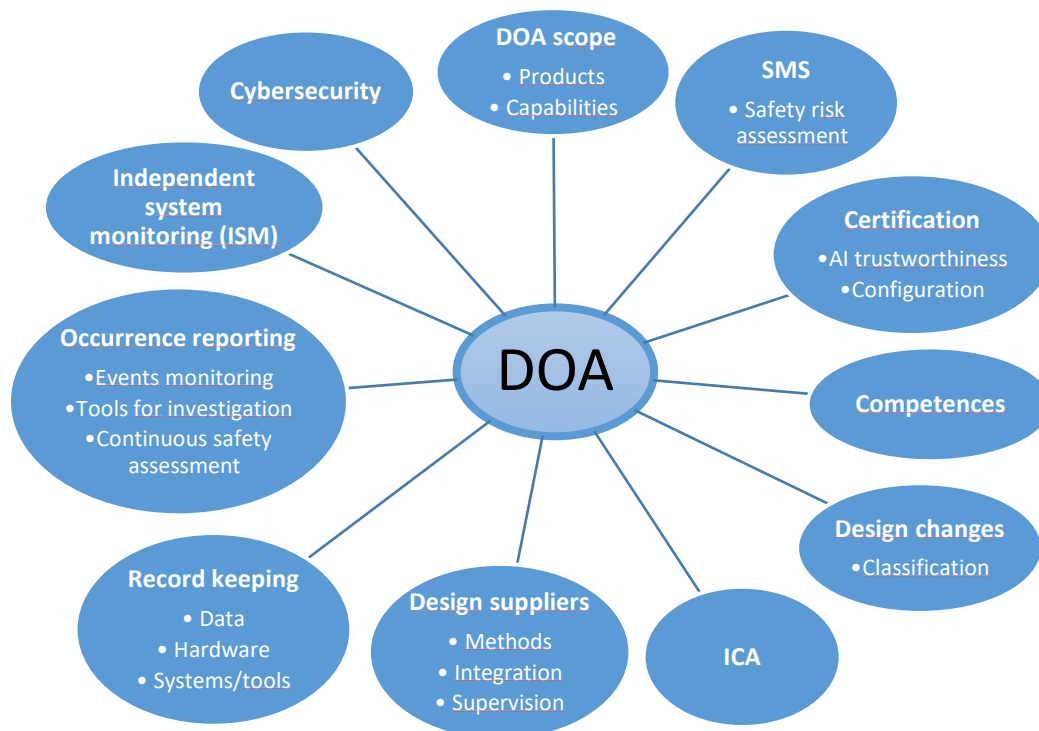


Figure 12 — DOA processes potentially affected by the introduction of AI/ML

Although almost all DOA processes are affected, the nature of the adaptation would be different depending on the interrelation of the process and the specificities of the AI technology.

The certification process would need to be deeply adapted to introduce new methodologies that will ensure compliance with the AI trustworthiness objectives as introduced in the previous chapters of this guidance. Similarly, new methodologies might be required for the record-keeping of AI-related data, for the independent system monitoring (ISM) process with regard to both compliance with and

adequacy of procedures, and for the continuous safety assessment of events when the root cause might be driven by the AI system.

With regard to design changes, new classification criteria may be required when approved type design related to AI is intended to be changed.

Other processes such as competences would need to be implemented considering the new AI technologies and the related certification process.

Finally, the DOA scope would need to reflect the capabilities of the organisation in relation to product certification and to privileges for the approval of related changes.



## D. Proportionality of the guidance

### 1. Technical building blocks footprint overview

	AI Level 1A	AI Level 1B
Learning Assurance	85 % (32/38)	100 % (38/38)
AI Explainability	20 % (1/5)	100 % (5/5)
AI Safety Risk Mitigation	100 % (2/2)	100 % (2/2)

#### Notes:

- The 100 % footprint for Level 1B is based on the fact that the objectives of this document have been defined for Level 1 AI applications only. Future work on Level 2 and Level 3 will increase the number of objectives and virtually reduce the footprints for both Level 1A and Level 1B applications.
- Proportionality would also depend on the criticality of the AI/ML application. Such considerations are not implemented at the level of the objectives of this document and are left open for further consideration in the lower-level objectives that will be defined in industrial standards such as the one that is under work in the joint EUROCAE/SAE WG-114/G-34.

## 2. Risk-based levelling of objectives

Building block	Objectives	1A	1B
Trustworthiness analysis	<b>CO-01:</b> The applicant should identify the high-level function(s)/task(s) to be performed by the (sub)system either in interaction with the human or in autonomy .	✓	✓
	<b>CO-02:</b> The applicant should define the AI-based (sub)system taking into account domain-specific definitions of 'system'.	✓	✓
	<b>CO-03:</b> The applicant should define and document the ConOps for all AI-based (sub)systems. A focus should be put on the definition of the operational design domain (ODD) and on the capture of specific operational limitations and assumptions.	✓	✓
	<b>CO-04:</b> The applicant should perform a functional analysis of the (sub)system.	✓	✓
	<b>CL-01:</b> The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.	✓	✓
	<b>SA-01:</b> The applicant should define metrics to evaluate the AI/ML component performance and reliability. *This objective applies only to product certification and ATS providers.	✓*	✓*
	<b>SA-02:</b> The applicant should perform a system safety assessment for all AI-based (sub)systems. *This objective applies only to product certification and ATS providers.	✓*	✓*
	<b>SA-03:</b> The applicant should define metrics to evaluate the AI/ML component performance. **This objective applies only to non-ATS providers.	✓**	✓**
	<b>SA-04:</b> The applicant should perform a safety support assessment for any change in the functional (sub)system embedding a component developed using AI/ML techniques or incorporating AI/ML algorithms. **This objective applies only to non-ATS providers.	✓**	✓**
	<b>IS-01:</b> For each AI-based (sub)system and its data sets, the applicant should manage those information security risks with an impact on safety, identifying and accounting for specific threats introduced by AI/ML usage.	✓	✓
	<b>ET-01:</b> The applicant should perform an ethics-based trustworthiness assessment for any AI-based (sub)system developed using ML techniques or incorporating ML algorithms .	✓	✓
	<b>ET-02:</b> In performing the ethics-based trustworthiness assessment, the applicant should address questions from the EU Commission Assessment List for Trustworthy AI (ALTAI), taking into account the clarifications brought in the following anticipated MOC.	✓	✓
	<b>ET-03:</b> The applicant should comply with national and EU data protection regulations (e.g. GDPR), i.e. involve their Data Protection Officer (DPO), consult with their National Data Protection Authority, etc .	✓	✓
	<b>ET-04:</b> The applicant should assess the environmental impact of the AI-based (sub)system.	✓	✓

Learning assurance	<b>DA-01:</b> The applicant should describe the proposed learning assurance process, taking into account each of the steps described in Sections C.3.3 to C.3.12, as well as the interface and compatibility with development assurance processes.	✓	✓
	<b>DA-02:</b> (Sub)systems requirements documents should be prepared and encompass the capture of the following minimum requirements : ...	✓	✓
	<b>DA-03:</b> The applicant should describe the system and subsystem architecture, to serve as reference for related safety (support) assessment and learning assurance objectives .	✓	✓
	<b>DA-04:</b> Each of the captured requirements should be validated.	✓	✓
	<b>DA-05:</b> Each of the captured (sub)system requirements should be verified.	✓	✓
	<b>DM-01:</b> The applicant should capture the DQRs for all data pertaining to the data management process, including but not limited to: ...	✓	✓
	<b>DM-02:</b> The applicant should capture the requirements on data to be pre-processed and engineered for the inference model.	✓	✓



Learning assurance	<b>DM-03:</b> To enable the data collection step, the applicant should identify explicitly and record the input space and the operating parameters that drive the selection of the training, validation and test data sets.	✓	✓
	<b>DM-04:</b> Once data sources are collected, the applicant should make sure that the data set is correctly annotated or labelled .	✓	✓
	<b>DM-05:</b> The applicant should make sure that the operations on data properly address the captured requirements (including DQRs).	✓	✓
	<b>DM-06:</b> The applicant should ensure a sequence of operations on the collected data in preparation of the training.	✓	✓
	<b>DM-07:</b> When applicable, the applicant should transform the pre-processed data from the input space into features which are effective for the performance of the selected ML algorithm.	✓	✓
	<b>DM-08:</b> The applicant should ensure that the data is effective for the stability of the model and the convergence of the learning process, possibly via normalisation.	✓	✓
	<b>DM-09:</b> The applicant should distribute the data into three separate and independent data sets which will meet the specified DQRs:...	✓	✓
	<b>DM-10:</b> The applicant should ensure validation and verification of the data all along the data management process so that the DQRs are addressed .	✓	✓
	<b>LM-01:</b> The applicant should describe the AI/ML components and model architecture (including computational graph and activation functions).	✓	✓
	<b>LM-02:</b> The applicant should capture the requirements pertaining to the learning management and training processes, including but not limited to: ... .	✓	✓
	<b>LM-03:</b> The applicant should document the credit sought from the training environment and qualify the environment accordingly.		✓
	<b>LM-04:</b> The applicant should provide quantifiable generalisation guarantees.	✓	✓
	<b>LM-05:</b> The applicant should document the result of the model training.	✓	✓
	<b>LM-06:</b> The applicant should document any model optimisation that may affect the model behaviour (e.g. pruning, quantisation) and assess their impact on the model behaviour or performance.		✓
	<b>LM-07:</b> The applicant should estimate bias and variance of the selected model family and should provide evidence of the reproducibility of the training process .		✓
	<b>LM-08:</b> The applicant should ensure that the estimated bias and variance of the selected model meet the associated learning process management requirements .	✓	✓
	<b>LM-09:</b> The applicant should perform an evaluation of the performance of the trained model based on the test data set and document the result of the model verification.	✓	✓
	<b>LM-10:</b> The applicant should perform a requirements-based verification of the trained model behaviour and document the coverage of the ML component requirements by verification methods.	✓	✓
	<b>LM-11:</b> The applicant should provide an analysis on the robustness (or stability) of the algorithms and of the trained model.		✓
	<b>LM-12:</b> The applicant should verify the anticipated generalisation bounds using the test data set.	✓	✓
	<b>IMP-01:</b> The applicant should capture the requirements pertaining to the implementation process, including but not limited to: ...	✓	✓
	<b>IMP-02:</b> Any post-training model transformation (conversion, optimisation, deployment) should be identified and validated for its impact on the model behaviour and performance.	✓	✓
	<b>IMP-03:</b> For each transformation step, the environment (i.e. software tools and hardware) that are necessary to perform model transformation should be identified and any associated assumptions or limitations should be captured and validated.	✓	✓
	<b>IMP-04:</b> The applicant should verify that any conversion, optimisation or transformation performed during the trained model implementation step have not altered the defined model properties.	✓	✓
	<b>IMP-05:</b> The differences between the hardware platform used for training and the one used for verification should be identified and assessed for impact on the inference model behaviour and performance.	✓	✓
	<b>IMP-06:</b> The applicant should perform an evaluation of the performance of the inference model based on the test data set and document the result of the model verification.	✓	✓

<b>Learning assurance</b>	<b>IMP-07:</b> The applicant should perform a requirements-based verification of the inference model behaviour and document the coverage of the ML component requirements by verification methods .	✓	✓
	<b>IMP-08:</b> The applicant should provide an analysis on the robustness (or stability) of the inference model.		✓
	<b>DM-14:</b> The applicant should perform a data and learning verification step to confirm that the appropriate data sets have been used for the training, validation and verification of the model and that the expected guarantees (generalisation, robustness) on the model have been reached.	✓	✓
	<b>CM-01:</b> The applicant should apply all configuration management principles to the AI-based (sub)system life-cycle data, including but not limited to: ...	✓	✓
	<b>QA-01:</b> The applicant should ensure that quality/process assurance principles are applied to the development of the AI-based (sub)system, with the required independence level.	✓	✓

<b>Explainability</b>	<b>EXP-01:</b> The applicant should identify the list of humans that are intended to interact with the AI-based system, at any stage of its life cycle, together with their roles, their responsibilities and their expected expertise (including assumptions made on the level of training, qualification and skills).		✓
	<b>EXP-02:</b> For each role, the applicant should identify which task(s) the humans are intended to perform in interaction with the AI-based system, as well as the task allocation pattern.		✓
	<b>EXP-03:</b> For each output of the system relevant to the task(s), the applicant should identify the need for an explanation depending on several criteria (including nature of the task, authority of the human and the level of AI-based system) and specify the set of necessary explanations to be provided to the human.		✓
	<b>EXP-04:</b> For each output relevant to the task(s), the applicant should ensure the validity of the specified explanation, based on actual measurements (e.g. monitoring) or on a quantification of the level of uncertainty.		✓
	<b>EXP-05:</b> The applicant should provide the means to record operational data that is necessary to explain the behaviour of the AI-based system post operations.	✓	✓

<b>Safety risk mitigation</b>	<b>SRM-01:</b> Once activities associated with all other building blocks are defined, the applicant should determine whether the coverage of the objectives associated with the explainability and learning assurance building blocks is sufficient or if an additional dedicated layer of protection, called hereafter safety risk mitigation (SRM), would be necessary to mitigate the residual safety risks to an acceptable level.	✓	✓
	<b>SRM-02:</b> The applicant should establish SRM means as identified in Objective SRM-01.	✓	✓

### 3. Rationale for the introduced proportionality

**Trustworthiness analysis:** AI application classification, ConOps definition, as well as the safety, security and ethics-based assessments constitute fundamental components of the trustworthiness analysis for the AI-based system. Therefore no proportionality can be introduced from an AI Level perspective. Modulation is however foreseen based on the criticality of the application; in particular, for the ethics-based assessment anticipated MOC.

**Learning assurance:** The proportionality introduced for Level 1A compared to Level 1B is based on an alleviation of objectives related to training phase, considering that for lower-complexity applications, the focus could be put on the inference model verification activities. This assumption will be revisited as the guidance evolves.

**AI explainability:** The rationale for alleviating the four first objectives is based on the assumption that existing guidelines and requirements for interface design should be sufficient to address Level 1A applications. This assumption will be revisited as AI explainability guidance will evolve for Level 1B and higher; therefore, this proportionality will evolve accordingly.

**AI safety risk mitigation:** This building block could be used by applicants for any level in case applicable learning assurance or AI explainability objectives could not be achieved. Therefore, no proportionality can be considered; however, fulfilment of all applicable learning assurance and AI explainability objectives practically alleviates the need to comply with the two SRM objectives.

## E. Annex 1 — Anticipated impact on regulations and MOC for major domains

The EASA Basic Regulation, beyond its main objective to establish and maintain a high uniform level of civil aviation safety in the Union, further aims to promote innovation, particularly by laying down requirements and procedures that are performance-based.

Considering the potential application of AI/ML solution in all the domains under the remit of the Agency, EASA intends to define a common policy that can be applied to the whole of the EU civil aviation regulatory framework, rather than issue domain-specific guidance.

This Annex provides an analysis of the anticipated impact on aviation regulations and on the means of compliance to the current regulations for the various impacted domains.

### 1. Product design and operations

#### 1.1. Anticipated impact of the introduction of AI/ML on the current regulations

In the product design and certification domain, the current implementing rules (Part 21) and CSs already offer an open framework for the introduction of AI/ML solutions.

In particular, requirements such as CS25.1301, 1302, 1309/SC-VTOL.2500, 2505, 2510 are considered to still be valid for evaluating the safety of AI-based systems, provided additional means of compliance and standards are developed to answer the gap identified in the building blocks of the AI Roadmap.

For AI Level 1 applications, no impact on the EU regulatory framework in relation to certification is necessary. For higher AI Levels (2 and 3), this assumption will need to be revisited when working on further updates to this document.

In the Air Operations domain, the current regulatory framework (Regulation (EU) No 965/2012 (Air OPS Regulation) in its general parts related to organisation requirements (Part-ORO) contains provisions based on safety management principles that allow operators to identify risks, adopt mitigating measures and assess the effectiveness of these measures in order to manage changes in their organisation and their operations (ORO.GEN.200). This framework permits the introduction of AI/ML solutions; however, certain existing AMC and GM will need to be revised and new AMC and GM will need to be developed in relation to AI/ML applications.

More specific provisions in the Air OPS Regulation, related to specific type of operations and specific categories of aircraft, may also need to be revised depending on the specific AI Level 1 application.

AI Levels 2 and 3 might require a deeper assessment on their regulatory impact on Air Operations particularly on the requirements for air crew. This assumption will need to be revisited when working on further updates to this document.

## 1.2. Anticipated impact of AI/ML guidance on the current AMC/MoC framework

### 1.2.1 Summary

The objectives identified in this document are anticipated to provide a sufficient framework in view of approving Level 1 AI applications (both for 1A and 1B). For higher AI Levels (2 and 3), this assumption will need to be revisited when working on further updates to this document.

The anticipated MOC will surely need to be completed based on the discussions triggered within certification projects, as well as based on industrial standards such as the one that is under work in the joint EUROCAE/SAE WG-114/G-34.

For the first applications, it will be necessary to establish the certification framework addressing the installation and certification of AI-based systems for a given project. That could be achieved by the preparation of Certification Review Items (CRI) using the guidelines from this document.

### 1.2.2 Detailed analysis

#### *Trustworthiness analysis*

From a safety and security assessment perspective, the current guidance (e.g. AMC25.1309 or MOC VTOL.2510) is fully applicable, as reflected in Sections C.2.2 and C.2.3. The Ethical guidelines provided in Section C.2.4 are mostly novel and constitute one of the impacts of considering AI/ML solutions compared to traditional product certification approaches.

#### *Learning assurance*

When dealing with development assurance, the current means of compliance for system, software and hardware development assurance are not sufficient to address the specificities of learning processes (i.e. data management + learning assurance), and need to be complemented through the guidelines for learning assurance (Section C.3) when dealing with the development of the AI/ML-based subsystem. For other (sub)systems not developed with or not embedding AI/ML solutions, the current applicable system, software and hardware development assurance guidance still applies.

#### *Explainability*

The need for explainability is specific to the use of AI/ML solutions and as such Section C.4 is a new MOC. It builds however on some existing guidance; in particular, the applicable human factors guidance already used in certification could provide a sufficient layer of MOC for Level 1A AI/ML applications.

#### *Safety risk mitigation*

The need for residual risk assessment is dependent on the capacity of the applicant to meet the applicable objectives of the learning assurance and AI explainability building blocks. Even if the risk mitigation foresees the use of traditional MOC (e.g. safety nets), the development of novel methods of mitigation will need to be investigated.

### *Part 21 AMC/GM for Design*

The technical particularities of AI technology might require a need to adapt or introduce new AMC & GM related to the following Part 21 points:

- 21.A.3A ‘Failures, malfunctions and defects’ with regard to potentially new methodologies needed for the analysis of data required to identify deficiencies in the design of AI components;
- 21.A.31 ‘Type design’ with regard to guidance in the identification of the AI-related data that constitutes the type design;
- 21.A.33 ‘Inspections and tests’ and 21.A.615 ‘Inspection by the Agency’ with regard to guidance to ensure adequate Agency review of data and information related to the demonstration of compliance;
- 21.A.55, 21.A.105 and 21.A.613 ‘Record-keeping’ with regard to guidance in the identification of the AI-related design information that needs to be retained and accessible;
- 21.A.91 ‘Classification of changes to a type-certificate’ with regard to guidance in the major/minor classification of changes to AI-related approved type design.

## 2. Training / FSTD

### 2.1. Anticipated impact of the introduction of AI/ML on the current regulations

The regulatory requirements for aircrew training are to be found in different Annexes to Regulation (EU) 1178/2011 (Aircrew Regulation).

In more detail, regulatory requirements are set in:

- Annex I (Part-FCL) in relation to licensing and training;
- Annex II (Part-ORA) in relation to organisational approvals.

Additional elements of flight crew training pertaining to the crew employed by operators are contained in the Air OPS Regulation.

Those regulations are mainly based on former Joint Aviation Authorities (JAA) regulatory requirements that were drafted almost 2 decades ago. All the structure of licensing and organisation approval is therefore referring to traditional methodologies in which the technological contribution is limited to the use of computer-based training (CBT) solutions for the delivery of theoretical elements and to aircraft and flight simulation training devices (FSTDs) to deliver practical flight training elements. Additionally, some reference to distance learning provisions are present allowing a certain flexibility for remote training.

The field of support of AI/ML solutions in the training domain may range from organisational aspects to monitoring functions up to more practical solutions in training delivery and performance assessment. The main impact will be on:

- the definition section to include the AI/ML components;
- the description of training programme delivery methodologies to address new technologies for administering the training courses;
- the crediting criteria for the use of AI/ML solutions; and
- organisation requirements in which the data management, analysis and correlation may play a role.

In any case, it is advisable that the initial use of AI/ML solutions in Aircrew training should be targeted to ground elements and simulator tasks.

## 2.2. Anticipated impact of AI/ML guidance on the current AMC/MoC framework

In support of the previous considerations, the AMC for the above-mentioned implementing rules shall be reviewed and updated to foresee the new technological solutions and to address the specificities of AI/ML solutions.

This review could run in parallel to the update of the regulatory framework which is already ongoing to incorporate new technologies and to accommodate emerging needs stemming from:

- new training needs for new categories of aircraft (e.g. VTOL or RPAS);
- new training devices (e.g. virtual or augmented reality).

The Aircrew Regulation is not intended to certify products and does not address the design process, therefore all the elements of the AI/ML model:

- trustworthiness analysis;
- learning assurance;
- explainability;
- safety risk mitigation

would need an effort to be created or tailored to the purpose.

## 3. Ground / ATM/ANS

### 3.1. Current regulatory framework relevant to the introduction of AI/ML

In addition to the Basic Regulation, Regulation (EU) 2017/373, applying to providers of ATM/ANS and other air traffic management network functions, lays down common requirements for:

- (a) the provision of ATM/ANS for general air traffic, in particular for the legal or natural persons providing those services and functions;
- (b) the competent authorities and the qualified entities acting on their behalf, which perform certification, oversight and enforcement tasks in respect of the services referred to in point (a);
- (c) the rules and procedures for the design of airspace structures.

Regulation (EU) 2017/373 is supplemented with Regulation (EC) No 552/2004<sup>14</sup> for interoperability, and Regulation (EU) No 376/2014 for occurrence reporting.

These Regulations open the path to the use of Level 1 AI. For higher AI Levels (2 and 3), this assumption will need to be revisited when working on further updates to this document.

### 3.2. Anticipated impact of AI/ML guidance on the current AMC and GM

#### 3.2.1. Summary

The current Easy Access Rules for ATM/ANS (Regulation (EU) 2017/373) embeds a comprehensive set of AMC and GM guiding the ATM/ANS service providers in complying with the requirements of the Regulation.

---

<sup>14</sup> Note: Regulation (EC) No 552/2004 was repealed by the Basic Regulation, but some provisions remain in force until 12 September 2023. To replace those provisions, a rulemaking task has been initiated.



The objectives identified in this document are anticipated to provide a sufficient framework in view of approving Level 1 AI applications (both for 1A and 1B). For higher AI Levels (2 and 3), this assumption will also need to be revisited when working on further updates to this document.

The current AMC will surely need to be completed based on the guidance material delivered, as well as based on industrial standards such as the one that is under work in the joint EUROCAE/SAE WG-114/G-34.

### 3.2.2. Detailed analysis

The following is an initial list of the AMC which could need adaptations:

ANNEX III — Part-TM/ANS.OR — AMC6 ATM/ANS.OR.C.005(a)(2) Safety support assessment and assurance of changes to the functional system

ANNEX III — Part-TM/ANS.OR — AMC1 ATM/ANS.OR.C.005(b)(1) Safety support assessment and assurance of changes to the functional system

ANNEX IV — Part-ATS — AMC4 ATS.OR.205(a)(2) Safety assessment and assurance of changes to the functional system

ANNEX XIII — Part-PERS — AMC1 ATSEP.OR.210(a) Qualification training

Of course, the associated GM could be impacted as well.

## 4. Aircraft production and maintenance

### 4.1. Anticipated impact of the introduction of AI/ML on the current regulations

Regulation (EU) No 1321/2014, covering continuing airworthiness and approval of related organisations, is not very specific about technical details and generally contains higher-level requirements. It already addresses the use of software or the use of test equipment and tools (e.g. ‘use of a software tool for the management of continuing airworthiness data’, ‘software that is part of the critical maintenance task’). Software making use of AI and/or ML could be covered under those requirements, including such software within test equipment.

However, the wording, being generic in many areas, still assumes a conventional way of planning and performing maintenance, meaning a *task-based approach*. Maintenance is divided into manageable portions of work (called ‘tasks’) which means human interference with the product at a defined point in time as a closed action which is signed off by humans when finished, with the product being released to service by explicit human action and signature.

Level 1 systems, with the human in command and in the specific case of maintenance closing out any activity by human signature and explicit release to service by human action, do not contradict this philosophy.

For Level 2 systems, this may require more attention, as humans still need to not only oversee, but also to explicitly close off the work performed by the systems with their signature. This may be possible within the frame of the current regulation but may limit the actions which can be carried out by systems.

Level 3 systems are not in line with the current regulation and would definitely require major changes, as the philosophy of explicit demonstration of airworthiness and release to service by humans would basically change to a withdrawal from service by systems finding lack of airworthiness.

It should also be noted that maintenance is a much more international business with more than a hundred states of registry being responsible compared to type certification with only about a dozen of states of design of large aeroplanes being responsible. This includes states with completely different regulations and hence will probably require a lot of international cooperation to harmonise the applicable regulations, guidance and standards.

#### 4.2. Anticipated impact of AI/ML guidance on the current MoC framework

In the maintenance domain, there is no MoC framework comparable to the one used in certification.

Additionally, a significant part of the approval is done by the competent authorities (NAAs), and the regulation makes specific reference to ‘officially recognised standards’ (industry standards, national standards) so the complete overall framework of applicable guidance is not that clearly defined, rendering thus the impact of AI/ML not that easy to be evaluated. Industry standards (e.g. SAE) may be used to show compliance with certain requirements.

‘Officially recognised standards’ as mentioned in the AMC material ‘means those standards established or published by an official body, being either a natural or legal person, and which are widely recognised by the air transport sector as constituting good practice.’ This allows the use of future standards on AI/ML developed by recognised official bodies (like SAE, ISO) for demonstrating compliance with certain requirements to the approving authority.

### 5. Environmental protection

#### 5.1. Current regulatory framework relevant to the introduction of AI/ML

The essential environmental protection requirements for products are laid out in the Basic Regulation Articles 9 and 55 for manned and unmanned aircraft respectively, and in its Annex III. These requirements are further detailed in Part 21 (in particular point 21.B.85) as well as in CS-34 ‘Aircraft engine emissions and fuel venting’, CS-36 ‘Aircraft noise’ and CS-CO2 ‘Aeroplane CO2 Emissions’. For the majority of manned aircraft, the AMC and GM linked to these requirements are defined in the appendices to ICAO Annex 16 and in Doc 9501 ‘Environmental Technical Manual’.

#### 5.2. Anticipated impact of AI/ML guidance on the current MoC framework

The AI/ML guidance for Level 1 systems is anticipated to have no impact on the current MoC framework for environmental protection. The impact of Level 2 or 3 AI/ML guidance will be assessed at a later stage. The safety-related guidelines in Chapter C of this document are anticipated to help provide adequate confidence in the functioning of AI/ML applications when demonstrating compliance with environmental protection requirements.

## 6. Aerodromes

### 6.1. Current regulatory framework relevant to the introduction of AI/ML

In addition to the Basic Regulation, Regulation (EU) No 139/2014<sup>15</sup> lays down requirements and administrative procedures related to:

- (a) aerodrome design and safety-related aerodrome equipment;
- (b) aerodrome operations, including apron management services and the provision of groundhandling services;
- (c) aerodrome operators and organisations involved in the provision of apron management and groundhandling services<sup>16</sup>;
- (d) competent authorities involved in the oversight of the above organisations, certification of aerodromes and certification/acceptance of declarations of safety-related aerodrome equipment<sup>17</sup>.

This regulation, in its consolidated form, does not represent a hinderance to the use of Level 1 AI use cases. For higher AI Levels (2 and 3), this statement might be revisited when the need would be brought to the attention of EASA by industry and overseen organisations, as well as manufacturers of safety-relevant aerodrome equipment.

### 6.2. Anticipated impact of AI/ML guidance on the current AMC and GM

The AMC and GM related to Regulation (EU) No 139/2014 support the implementation of the implementing rule requirements by the organisations concerned.

Most of the AMC and GM do not refer to specific technologies, so they do not impede the approval of Level 1 AI applications. For higher AI Levels (2 and 3), this statement might need to be revisited when the need by industry and overseen organisations, as well as manufacturers of safety-relevant equipment, would be brought to the attention of EASA.

### 6.3. Preliminary analysis

The following IRs and the related AMC and GM are relevant to the AI use cases further below:

- ADR.OPS.B.015 Monitoring and inspection of movement area and related facilities
- ADR.OPS.B.020 Wildlife strike hazard reduction
- ADR.OPS.B.075 Safeguarding of aerodromes

---

<sup>15</sup> As subsequently amended by Commission Regulation (EU) 2018/401 regarding the classification of instrument runways, Commission Implementing Regulation (EU) 2020/469 as regards requirements for air traffic management/air navigation services, Commission Delegated Regulation (EU) 2020/1234 as regards the conditions and procedures for the declaration by organisations responsible for the provision of apron management services, and Commission Delegated Regulation (EU) 2020/2148 as regards runway safety and aeronautical data.

<sup>16</sup> For groundhandling services and providers of such services, there are at this stage no detailed implementing rules. These are expected by 2023 at the latest.

<sup>17</sup> The oversight framework for safety-related aerodrome equipment will be developed in due course but is at the time of writing not yet in place, neither are the European certification specifications for such equipment.

#### 6.4. Anticipated impact of AI/ML guidance on the current and future CSs for aerodrome design and safety-related aerodrome equipment

The current CSs and to Regulation (EU) No 139/2014 provide a comprehensive set of requirements for the design of aerodrome infrastructure and for some aerodrome equipment (as far as it exists stemming from the transposition of Annex 14). Once the future framework for safety-related aerodrome equipment exists, EASA will issue European certification specifications for such equipment. This process will allow for the further introduction of AI/ML solutions at aerodromes, if they fulfil the demands placed on them with respect to safety.



## F. Annex 2 — Use cases for major domains

### 1. Introduction

With the objective of ensuring that its guidelines will remain practical for the applicants, EASA has engaged with the aviation industry and stakeholders, in order to support the elaboration of the guidelines with actual use cases from the major aviation domains.

It is not the intention that each use case is complete and fulfils the full set of objectives described in this guidance document, but rather to evaluate that the objectives and proposed anticipated MOC are practical. This may result in a number of use cases not implementing all AI trustworthiness building blocks.

Before entering into the use cases, Table 4 — Classification applied to use cases below provides the audience with a description of how each use case has been classified as per Table 1 — EASA AI typology and definitions.

		Domain					
		Aircraft design and operations		Ground / ATM/ANS		Aircraft production and maintenance	
EASA AI Roadmap AI Level (subsystem)	Function allocated to the (sub)systems (adapted HARVIS LOAT terminology)	Visual landing guidance system	Pilot assistance – radio frequency suggestion	AI-based augmented 4D trajectory prediction	Time-based separation + Optimum runway delivery	Controlling corrosion by usage-driven inspections	Damage detection in images
<b>Level 1A Human augmentation</b>	Support to information acquisition	camera + pre-processing component	ATC radio communication	Data acquisition (FPL and updates, radar + weather)	Data acquisition (weather + radar)	Maintenance, environment, operator / manufacturer databases	infrared camera
	Support to information analysis	Runway object classification + bounding box + tracking/filtering algorithm	Voice recognition	4D trajectory calculation – Curtain + Climb and descent rate	Information preparation (pairs, applicable separation)	Predicted corrosion level + Time to inspect for corrosion	Damage classification
<b>Level 1B Human assistance</b>	Support to decision/action selection	x	Radio frequency suggestion for pilot validation	x		x	Support decision to repair for inspector validation
<b>Level 2 Human-AI collaboration</b>	Overseen automatic decision/action selection	x	x	x	Trajectory prediction + uncertainty calculation	x	x
	Overseen automatic action implementation	x	x	x	x	x	x
<b>Level 3A More autonomous AI</b>	Overridable automatic decision/action selection	x	x	x	x	x	x
	Overridable automatic action implementation	x	x	x	x	x	x
<b>Level 3B More autonomous AI</b>	Non-overridable automatic decision/action selection	x	x	x	x	x	x
	Non-overridable automatic action implementation	x	x	x	x	x	x

Table 4 — Classification applied to use cases

Where:

represents the AI-based system or subsystem; and  
The ML component is in blue

## 2. Use cases — aircraft design and operations

### 2.1. Visual landing guidance system (derived from the CoDANN report use case)

#### 2.1.1. Trustworthiness analysis — description of the system and ConOps

##### 2.1.1.1. Description of the system

**Objective CO-02:** The applicant should define the AI-based (sub)system taking into account domain-specific definitions of ‘system’.

The application facilitates landing operations by identifying the runway through an image recognition system and providing advisory information to the pilot.




Figure 13 – System overview (source (Daedalean, 2020))

The system is composed by a combination of a camera unit, a pre-processing component (both constituting a first subsystem contributing to the ‘perception’ function), a neural network and a tracking/filtering component (both constituting a second subsystem contributing to the ‘analysis’ function). This second subsystem takes as input the output of the perception subsystem, namely an indication whether a runway is present or not, corner coordinates (these are relevant only if the runway likelihood is high enough) and a link to an avionics display (third subsystem) to support the landing operations. It then uses those to provide the actual visual landing guidance.

Definition of ‘system’ from ED-79A/ARP-4754A is taken as reference for this airborne application (i.e. a combination of inter-related items arranged to perform a specific function(s)).

**Objective CO-04:** The applicant should perform a functional analysis of the (sub)system.

Considering the CoDANN report Section 9.2.2, a possible functional decomposition of the system is the following:

- Function 1: To support the pilot in landing on a runway/vertiport
  - Function 1.1: To detect the runway/vertiport position
    - Function 1.1.1: To sense the aircraft’s environment and provide the flight computer with an image of the environment
    - Function 1.1.2: To pre-process the image
    - Function 1.1.3: To detect the runway/vertiport in a given image
    - Function 1.1.4: To track the target position
  - Function 1.2: To compute the flight director order to the runway/vertiport
- Function 2: To monitor the system
  - Function 2.1: To monitor sensors
  - Function 2.2: To monitor internal data buses



- Function 2.3: To monitor the neural network behaviour
- Function 2.4: To monitor the flight director
- Function 3: To interface with the aircraft systems
  - Function 3.1: To receive the GPS data
  - Function 3.2: To receive the digital terrain elevation data
  - Function 3.3: To receive the phase of flight
  - Function 3.4: To receive electrical power
  - Function 3.5: To provide visual guidance to the pilot
  - Function 3.6: To provide monitoring data to the display

The functional allocation to the subsystems and components can be done as follows:

Sub-system	Component	Allocated functions
#1	Optical sensor	F.1.1.1, F.2.1, F.3.4
#2	Main processing unit	F.1.1.2, F.1.1.4, F.1.2, F.2.2, F.2.4, F.3.1, F.3.2, F.3.3, F.3.4, F.3.5, F.3.6
#2	NN processing	F.1.1.3, F.2.3
#3	Avionics display	F.3.5, F.3.6

Table 5 — Functional allocation to the subsystems and components

#### 2.1.1.2. Concept of operations

**Objective CO-03:** The applicant should define and document the ConOps for all AI-based (sub)systems. A focus should be put on the definition of the operational design domain (ODD) and on the capture of specific operational limitations and assumptions.

A detailed description of possible ConOps related to this use case AI-based system can be found in Section 4.1 of the CoDANN IPC Report (Daedalean, 2020). The use case under consideration here corresponds to the Operational Concepts 1a or 2a as described in Table 4.1 of the CoDANN Report, that is to say is limited to displaying the output of the AI/ML-based subsystem on a glass cockpit display, with no flight computer guidance involved.

#### 2.1.1.3. Description of the system(s) involved (inputs, outputs, functions)

The system is composed of three subsystems, #1 implementing the perception function composed of the high-resolution camera, #2 implementing the pre-processing, image analysis and post-processing function composed of two flight computers and including the convolutional neural network (CNN) and the tracking/filtering algorithm, and #3 composed of the portions of the avionics display system supporting the visual landing guidance system's operations.

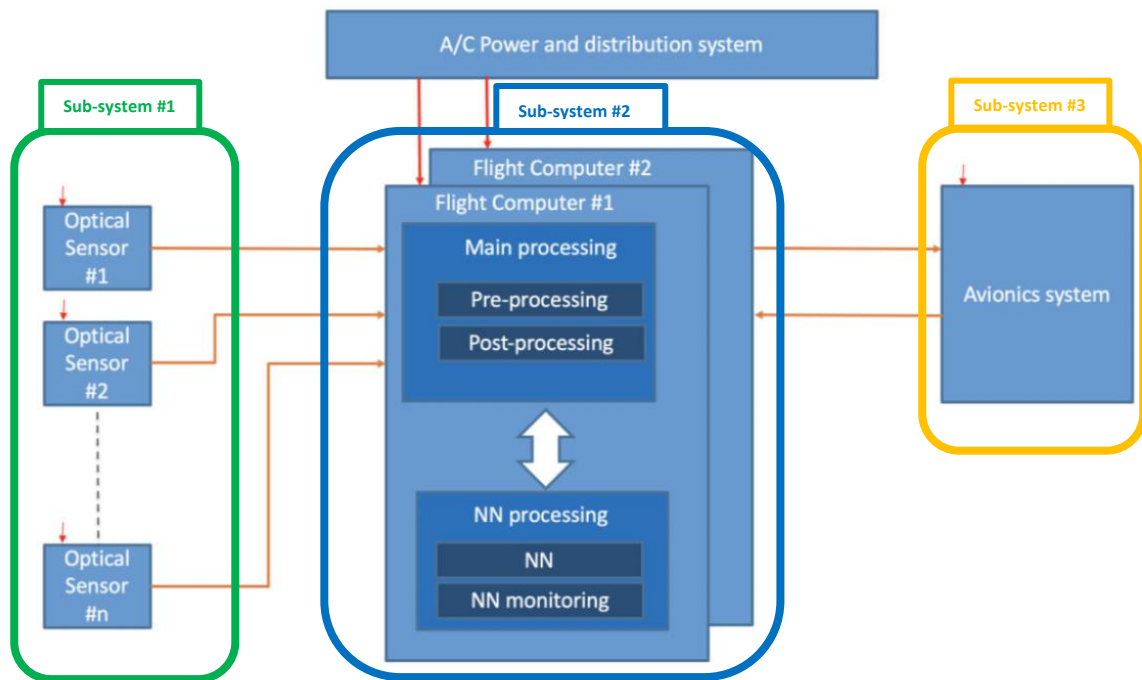


Figure 14 — System breakdown in subsystems and components (source (Daedalean, 2020))

#### 2.1.1.4. Expected benefits and justification for Level 1

The application is intended to provide additional information to the pilot in the form of a runway image displayed in the cockpit from the moment the runway is detected (cruise phase/holding pattern) until the decision to land is confirmed or a go-around is performed.

**Objective CL-01:** The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.

The **AI Level 1A 'Human augmentation'** classification is justified by only providing additional/advisory information (**support to information analysis**) to the pilot without any suggestion for action or decision making.

#### 2.1.2. Trustworthiness analysis — safety assessment

**Objective SA-01:** The applicant should define metrics to evaluate the AI/ML component performance and reliability.

Based on the discussions from the CoDANN report (Daedalean, 2020) Chapter 8, two types of metrics are considered for this use case:

- For the evaluation of the binary classification of the runway object, the precision and recall measures can be used to first select the best model and then to evaluate the operational performance.
- For the evaluation of the bounding box, the use of the Jaccard distance can be a useful metric for model selection.

**Objective SA-02:** The applicant should perform a system safety assessment for all AI-based (sub)systems.

A preliminary FHA can be found in the CoDANN report (Daedalean, 2020) Section 9.2.4. For the purpose of this use case discussion, the system can contribute to failure conditions up to Hazardous (as defined in applicable CSs). More severe failure conditions should be considered in case of linking the system to an autopilot, but this would trigger another classification for this AI-based system, probably up to a Level 2.

### 2.1.3. Learning assurance

#### 2.1.3.1. Data preparation

**Objective DM-03:** To enable the data collection step, the applicant should identify explicitly and record the input space and the operating parameters that drive the selection of the training, validation and test data sets.

The input space for this use case is the space of 512 x 512 RGB images that can be captured from a specific camera mounted on the nose of an aircraft, flying over a given region of the world under specific conditions, as defined in the ConOps and in the requirements.

A possible list of relevant operating parameters for the collection of data sets includes the following:

Parameter	Meaning	Domain
#1	Altitude	0, MAX_ALT
#2	Angle of approach when landing	0, 90
#2	Time of day	6, 21
#4	Binary variable denoting rain	0,1
#n	...	...

*Table 6 — Use-case-relevant operating parameters*

**Objective DM-04:** Once data sources are collected, the applicant should make sure that the data set is correctly annotated or labelled.

In the context of this use case, the annotation task consists of marking each of the four runway corners in every image. The review of the annotation is performed through a manual review.

#### 2.1.3.2. Data validation

**Objective DM-10:** The applicant should ensure validation and verification of the data all along the data management process so that the DQRs are addressed.

— Data completeness and representativeness

The set of operating parameters are first reviewed with respect to the set of requirements and with the ODD, to make a first evaluation of their intrinsic completeness relatively to the use case application.

The approach is completed by the definition of a distribution discriminator  $D$  using the ODIN method from the paper (Enhancing the reliability of out-of-distribution image detection in neural networks, 2018).

Refer to the CODANN Report (Daedalean, 2020), Section 6.2.8, for more information.

— Data accuracy

To demonstrate that the model was provided with correct data samples during the design phase, several sources of errors need to be shown to be minimal and independent, or else to be mitigated.

First, the systematic errors in the data, also called data bias are identified using statistical testing and mitigated.

In addition, specific attention is put on single-source errors which could introduce bias in the resulting data sets. This type of error has been avoided by using the same source for data collection in operations as well.

Furthermore, labelling errors have been addressed by involving multiple independent actors in the labelling activity and its validation.

— Data traceability

The data sets undergo a conversion from the raw images format to 8bit RGB, removal of irrelevant information as necessary and may be modified to enhance colour, brightness and contrast. These transformations are fully reproducible and a trace of the changes to the origin of each data pair is recorded. This applies also to the annotations.

— Data sets independence

The training/validation and test data sets are created by independent groups. The test data set is not accessible during the design phase.

### 2.1.3.3. Learning process management

**Objective LM-01:** The applicant should describe the AI/ML components and model architecture (including computational graph and activation functions).

**Objective LM-02:** The applicant should capture the requirements pertaining to the learning management and training processes.

The data indicated in Objectives LM-01 and LM-02 are documented, including substantiation for the selection of the model architecture, algorithm selection as well as for the learning parameters selection.

**Objective LM-03:** The applicant should document the credit taken from the training environment and qualify the environment accordingly.

The open-source software library TensorFlow is chosen and the training is run on a on-premises computer equipped with a GPU (NVIDIA K80), running in a Linux-based operating system, with device-specific libraries such as CUDA and cuDNN.

**Objective LM-04:** The applicant should provide quantifiable generalisation guarantees.

The approach to estimate the generalisation guarantees of the use case CNN is based on a combination of generalisation bounds based on model capacity and on the ‘evaluation-based’ approach summarised in Section 5.3.7 of the CoDANN report (Daedalean, 2020), the latter requiring a large test data set.

A failure probability of less than  $10^{-4}$  per image provided by generalisation guarantees would still accumulate exponentially over time, and this would therefore be insufficient to meet the overall safety objectives (e.g.  $10^{-7}$  per flight hour for a Hazardous failure condition). To prevent this, other architectural mitigations (such as filtering or tracking) are needed, and their analysis is tightly integrated with the one of the neural network component.

**Objective LM-05:** The applicant should document the result of the model training.

The resulting training curves and performance on the training and validation sets are recorded in the learning accomplishment summary (LAS).

**Objective LM-06:** The applicant should document any model optimisation that may affect the model behaviour (e.g. pruning, quantisation) and assess their impact on the model behaviour or performance.

No optimisation is performed at the level of the learning process.

**Objective LM-07:** The applicant should estimate bias and variance of the selected model family and should provide evidence of the reproducibility of the training process.

The Bootstrapping and Jack-knife methods have been used to estimate bias and variance and support the model family selection.

To this purpose, the learning process is repeated several times with variations in the training data set to show that:

- the models have similar performance scores on training and validation data sets;
- the selected model is not adversely impacted by a small change in the training data set.

**Objective LM-08:** The applicant should ensure that the estimated bias and variance of the selected model meet the associated learning process management requirements.

- The bias and variance of the selected model have been identified. No systematic error has been detected.

**Objective LM-09:** The applicant should perform an evaluation of the performance of the trained model based on the test data set and document the result of the model verification.

The resulting performance of the model on the test data set is recorded in an accomplishment summary.

#### 2.1.4. Safety risk mitigation

**Objective SRM-01:** Once activities associated with all other building blocks are defined, the applicant should determine whether the coverage of the objectives associated with the explainability and learning assurance building blocks is sufficient or if an additional dedicated layer of protection, called hereafter safety risk mitigation (SRM), would be necessary to mitigate the residual safety risks to an acceptable level.

In this use case, it is considered that all objectives related to the trustworthiness analysis, learning assurance and explainability building blocks can be fully covered.

**Objective SRM-02:** The applicant should establish SRM means as identified in Objective SRM-01.

No SRM mitigations are identified in SRM-01.

### 2.2. Pilot assistance — radio frequency suggestion

#### 2.2.1. Trustworthiness analysis — description of the system and ConOps

##### 2.2.1.1. Description of the system

**Objective CO-02:** Define the AI-based system taking into account domain-specific definitions of 'system'.

An example of AI Level 1B application for pilot assistance may be voice recognition and suggestion of radio frequencies.

The application recognises radio frequencies from ATC voice communications and suggests to the pilot a frequency that has to be checked and validated by the pilot before tuning the radio accordingly (e.g. tuning the standby VHF frequencies).

##### 2.2.1.2. Expected benefits and justification for Level 1

The application is expected to reduce workload or help the pilot to confirm the correct understanding of a radio frequency in conditions of poor audio quality.

**Objective CL-01:** The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.

The Level 1B classification is justified by providing support to the pilot in terms of gathering the information and suggesting it to the pilot for validation before any action is taken, i.e. support to decision-making. The frequency may be either displayed to the pilot who then will tune it manually or may be pushed automatically into the avionics after acceptance of the pilot. The two cases will require a different level of assessment.

#### 2.2.2. Trustworthiness analysis — safety and security assessment

**Objective SA-02:** The applicant should perform a system safety assessment for all AI-based (sub)systems.

A risk of complacency and over-reliance on the applications exists.

**Objective IS-01:** For each AI-based (sub)system and its data sets, the applicant should manage those information security risks with an impact on safety, identifying and accounting for specific threats introduced by AI/ML usage.

If the application is integrated with the avionics with the possibility to exchange data, the check and validation function, as well as data integrity and security aspects, will have to be further assessed.

### 3. Use cases — ATM/ANS

#### 3.1. AI-based augmented 4D trajectory prediction — climb and descent rates

The objective of the use case is to improve the accuracy of a predicted 4D trajectory by better estimating the climb and descent rates with the use of deep learning techniques. To this purpose, a deep neural network (DNN) is introduced to replace the software item in charge of the estimation of the climb and descent rates.

##### 3.1.1. Description of ConOps and systems involved in the use case

**Objective CO-03:** The applicant should define and document the ConOps for all AI-based (sub)systems. A focus should be put on the definition of the operational design domain (ODD) and on the capture of specific operational limitations and assumptions.

##### 3.1.1.1. Introduction

All information in this section has been derived from both the ATFCM Users Manual (EUROCONTROL, 2020) and the IFPS Users Manual (EUROCONTROL, 2020).

A 4D trajectory of a flight during pre-tactical phase, tactical phase, or when the flight is airborne is a fundamental element for correct network impact assessment and potential measures to be taken on congested airspace.

The 4D trajectory is (re)calculated in the context of many different services delivered by the Network Manager. Many different roles are interested in the 4D trajectory. Many different triggering events can generate the computation of a 4D trajectory.

Note: 4D trajectory and flight profile are to be considered as synonyms in this document.

##### 3.1.1.2. Roles

Four different categories of end users with the following roles are involved in the operations of the 4D trajectory:

- (Aircraft operator (AO)) flight dispatcher;
- Air traffic controller (ATCO), with the area or en-route (ATC in this document) and the aerodrome or tower (TWR in this document);
- Flow management position (FMP); and
- Network Manager (NM) tactical team: The NM tactical team is under the leadership of the Deputy Operations Manager in charge of managing the air traffic flow and capacity management) (ATFCM) daily plan during the day of operation. The tactical team is formed by the tactical Senior Network Operations Coordinator, the Network Operations Controllers, the Network Operations Officer and the Aircraft Operator Liaison Officer on duty.



#### 3.1.1.3. 4D trajectory before flight departure

- Initial 4D trajectory based on flight plan (flight plan or filed flight plan (FPL))

A first version of the 4D trajectory is computed on the reception of a valid FPL by the AO.

The 4D trajectory is distributed to all ATCOs and TWR responsible for the ATC where the flight takes place.

- Reception of a change message (CHG)

When an individual FPL has been filed but it is decided, before departure, to use an alternative routing between the same aerodromes of departure and destination, the AO may decide to send a CHG for any modification.

Reception of a CHG triggers the re-calculation of the 4D trajectory and distribution to all ATCOs and TWRs responsible for the ATC where the flight takes place.

- Reception of a delay(ed) message (DLA)

On receipt of a DLA by the AO, the IFPS shall re-calculate the 4D trajectory of that flight based on the revised estimated off-block time (EOBT).

- ATFCM solutions to capacity shortfalls

Where overloads are detected and the collaborative decision-making (CDM) process is initiated, different ATFCM solutions should be considered between the NM and the respective FMP(s).

This consists in:

- (a) optimisation of the utilisation of available capacity, and/or utilisation of other available capacities (rerouting flows or flights, flight Level (FL) management) or advancing traffic; and/or
- (b) regulation of the demand.

Most of the time, such ATFCM solutions will generate computation of 4D trajectories for the flights impacted.

#### 3.1.1.4. 4D trajectory all along the life cycle of the flight

- Updating Central Airspace and Capacity Database (CACD) Data in Predict / Enhanced Tactical Flow Management System (ETFMS)

Updates to a subset of the environmental data (i.e. taxi time, runway in use for departures and arrivals, time to insert in the sequence (TIS), time to remove from the sequence (TRS), etc.) will trigger the re-computation of the flight profile of the aircraft concerned.

Taxi time updates and actual SID used by aircraft originating from A-CDM (from EOBT-3h up to target take-off time (TTOT)) are communicated to the ETFMS via departure planning information (DPI) messages for each individual aircraft.

The above parameters may be updated for each different (active) runway and the flight profiles are re-computed using this information.

- Airport CDM

Most advanced airports have engaged with NM in a CDM process aiming at improving the quality of information based on which decisions are made, then leading to enhanced operational efficiency and facilitating optimum use of available capacity.

Some of the DPI messages received by the ETFMS will have as a consequence the re-computation of the 4D trajectory for this specific flight (e.g. taxi time updates and actual SID used by aircraft originating from A-CDM (from EOBT-3h up to TTOT)).

- ETFMS flight data message (EFD) / publish/subscribe flight data (PSFD)

The EFD is basically an extract of flight data that is available in the ETFMS of which the flight profile is the most important part.

The EFD is sent by ETFMS to ANSPs of flight data processing areas (FDPAs) that request such information.

In the last years, EFDs have been complemented with PSFDs accessible via the NM B2B services.

#### 3.1.1.5. 4D trajectory after departure

- Flight data information

On departure, the AO should send a departure message (DEP). Some AOs are sending aircraft (operator) position report (APR) messages to ETFMS. This data will then be used by the ETFMS to update the 4D trajectory in the current flight model (current tactical flight model (CTFM)) of the flight and also all other times (estimated times over (ETOs)) in the flight profile are updated accordingly.

Upon the flight's entry into the NM area, the flight's profile is then updated by first system activation (FSA) and correlated position report (CPR) messages where applicable.

For trans-Atlantic flights, flight notification message (FNM) from Gander and message from Shanwick (MFS) are messages that are received which provide an estimate for the oceanic exit point. MFS and FNM are processed first by integrated initial flight plan processing system (IFPS), that sends then the information to ETFMS. IFPS also sends it to AOs.

These estimates are used by the ETFMS to update the corresponding flight profiles.

- Correlated position reports (CPRs)

A flight may deviate from its last computed profile triggering a profile recalculation.

#### 3.1.1.6. other usage of 4D trajectory

- Network simulations

The NM is responsible for the management of strategic ATFCM plans. Such plans rely on many simulations running in parallel and involve FMPs and AOs. Some simulation can imply the 4D trajectory calculations for flows under scrutiny.

- Post OPS analysis and reporting

The NM regularly reports on its activities and deliveries.

Among these post-operations activities, some reports elaborate on alternative 4D trajectories of the flown ones for further analysis in terms of flight efficiency (improved use of airspace, fuel consumption, etc.).

### 3.1.1.7. Measures

Considering a normal day of operations with:

- 30 000 flights;
- 5 000 000 CPR messages received;
- multiplicity of scenarios being launched in the context of ATFCM operations;
- new requests coming from A-CDM airports,

a rough estimation gives **300 000 000** of 4D trajectories computed every day.

### 3.1.2. Expected benefits and justification for Level 1

**Objective CL-01:** The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.

The **AI Level 1A 'Human augmentation'** classification is justified by only augmentation of the precision of the climb and descent phases, which participate to the computation of the 4D trajectory distributed to the roles involved with the flight profile. All decisions based on the predicted 4D trajectory are performed by a human or a machine with many indirections to the flight profile. It is then considered that this augmentation (**support to information analysis**) does not suggest any action or decision-making.

### 3.1.3. Trustworthiness analysis

#### 3.1.3.1. Safety support assessment

**Objective SA-04:** The applicant should perform a safety support assessment for any change in the functional (sub)system embedding a component developed using AI/ML techniques or incorporating AI/ML algorithms.

The following describes the process that has been supporting the safety support assessment of the use case. The execution of the process takes into account the existence of a baseline safety support case (BSSC) for the NM services currently in the operations.

For reasons of conciseness, only the main outcomes of the process are presented in this document. For more information, please refer to Section 4.1 of the full report available by EUROCONTROL.

- Safety support assessment process

The safety support assessment of the change has been carried out in compliance with the requirements included in Regulation (EU) 2017/373 and its associated AMC and GM for service providers other than ATS providers.

The first step is the understanding and scoping of the change. It includes determination of the changed/new components of the NM functional system (FS), impacted (directly and indirectly) components of the NM FS, interfaces and interactions, and its operational context.

The second step of the safety support assessment used the failure mode and effect analysis (FMEA) technique to identify functional system failures. These failures can cause the services to behave in a non-specified manner, resulting in a different to the specified service output (e.g. lost, incorrect, delayed). Failure modes are linked (traceable) to the degraded mode(s) that can be caused by the failure. Where appropriate, internal (safety support requirements) and external mitigations (assumptions) have been derived to reduce or prevent undesired failure effects.

The third step of the safety support assessment, the degraded mode causal analysis, has been performed by means of facilitated structured brainstorming. It enabled the identification of the potential contribution of the changed and impacted elements of the NM FS to the occurrence of the degraded modes, as well as the establishment of safety support requirements to control the occurrence of the degraded modes and hence the service behaviour.

The fourth step will be the provision of the needed arguments and justification to demonstrate compliance with the safety support requirements.

#### — Safety support requirements

The table below contains the inventory of the safety support requirements, i.e. the necessary means and measures derived by the safety support assessment to ensure that NM operational services will behave as specified following the implementation of AI for the estimation of aircraft climb and descend rates. The table below provides traceability to the mitigated service degraded modes and to the service performance.

No transition safety support requirements have been derived as the implementation of AI for the aircraft climb and descent rate estimation does not require a transition period.

ID	Safety support requirement	Mitigated degraded mode	Impacted service performance
R-01	Curtain shall implement alternative way of prediction calculation (e.g. based on fallback BADA table).	DGM06 DGM10 DGM11 DGM15 DGM17 DGM19	integrity  availability
R-02	The AI component shall return error code in case it is able to detect an incorrect prediction.	DGM10	integrity
R-03	Curtain shall implement means to detect incorrect prediction provided by the AI component.	DGM10	integrity
R-04	Curtain shall perform validation check of the AI prediction using a set of established criteria.	DGM10 DGM15 DGM19	integrity
R-05	Rules for use of alternative prediction computation by curtain shall be implemented.	DGM-10	integrity
R-06	Learning assurance shall be applied to the AI module to optimise the model generalisation.	DGM10	integrity
R-07	Carry out adequate tests of the AI module.	DGM10	integrity
R-08	Carry out focused TensorFlow tests.		
R-09	Measure the time to obtain a prediction and trigger alarm in case a defined threshold has been reached.	DGM06 DGM11 DGM17	availability
R-10	Design and execute dedicated test to refine the prediction validity threshold.	DGM10 DGM15 DGM19	integrity
R-11	Carry out load tests (at development and verification level).	DGM06 DGM11 DGM17	availability
R-12	Ensure resources (e.g. memory, disk space, CPU load) monitoring in operations.		
R-13	Comply with the SWAL4 requirement for IFPS/ETFMS.	DGM10 DGM15 DGM19	integrity

*Table 7 — Safety support requirements*

— Behaviour in the absence of failures

To ensure the completeness of the change argument, there is a need to analyse the behaviour of changed and impacted components of the NM FS in the absence of failures in order to prove that the NM services continue to behave as specified in the respective service specifications.

As a result of this analysis, the following safety support requirements have been placed on the changed and impacted by the change FS elements:

- **R-14.** The AI/ML component shall use industry-recognised technology (e.g. deep neural network) for training the prediction model. The use of TensorFlow shall be considered.
- **R-15.** The AI/ML component shall ensure correct generalisation capabilities which shall be verified by means of pre-operational evaluation with real flight plan data and, if necessary, improved.
- **R-16.** The AI/ML component shall expose an interface which shall be consumed by Curtain.
- **R-17.** The AI/ML component shall be able to process up to 100 requests per second. Curtain shall send a prediction request to the AI/ML component upon identification of the need to build a new or update an existing 4D trajectory.
- **R-18.** Curtain shall process the climb and descent rate predictions delivered by the AI/ML component.

— Assumptions

The table below contains the list of assumptions made during the safety support assessment that may apply and impact on the effectiveness and/or availability of the mitigation means and measures. It traces the assumptions and conditions to the associated degraded modes where they have been raised. The table also provides justification why the assumptions are correct and valid.

ID	Assumption/ Condition	Degraded Modes	Justification
A-01	Exhaustion of system resources will not only affect the AI module, but Curtain and other system processes, too.	DGM06 DGM11 DGM17	The AI module, Curtain and other critical system processes use the same computing resources (disk, memory and CPU).
A-02	By design, consecutive incorrect rate prediction for different flights cannot occur.	DGM10 DGM19	Successive incorrect rate predictions due to AI design issues will be identified during the software development and integration testing phase, and the AI predictive model will be enhanced consequently.
A-03	Failure of Curtain to compute an alternative prediction cannot occur for all flights.	DGM10 DGM19	This is a legacy function that has been proven in operation since years.

*Table 8 — Use case assumptions*

— Safety support requirements satisfaction

This section will provide the needed assurance that the safety support requirements listed above are implemented as required in order to ensure that NM services (flight planning, ATFCM and centralised code assignment and management system (CCAMS)) will continue to behave only as specified in the respective service specifications.

### 3.1.3.2. Information security considerations

**Objective IS-01:** For each AI-based (sub)system and its data sets, the applicant should manage those information security risks with an impact on safety, identifying and accounting for specific threats introduced by AI/ML usage.

The following describes the process that has been supporting the security assessment conducted on the use case.

For reasons of conciseness, only the main outcomes of the process are presented in this document. For more information, please refer to Section 4.2 of the full report available by EUROCONTROL.

#### — Approach to security assessment

The high-level security assessment is based on the following works:

- Microsoft:
  - [AI/ML Pivots to the Security Development Lifecycle Bug Bar<sup>18</sup>](#)
  - [Threat Modeling AI/ML Systems and Dependencies<sup>19</sup>](#)
  - [Failure Modes in Machine Learning<sup>20</sup>](#)
- [A Survey on Security Threats and Defensive Techniques of Machine Learning: A Data Driven View \(Liu, 2018\)](#)
- MITRE [Adversarial ML Threat Matrix<sup>21</sup>](#).

The objective is to establish different potential attack paths and identify possible shortcomings.

As illustrated in Figure 15 — Illustration of defensive techniques of machine learning, we are considering the following security threats to the ML life cycle:

- Poisoning attacks: Those aim at corrupting the training data so as to contaminate the machine model generated in the training phase, aiming at altering predictions on new data.
- Evasion, impersonate & inversion attacks: Those aim at recovering the secret features used in the model through careful queries or other means.

<sup>18</sup> Source: <https://docs.microsoft.com/en-us/security/engineering/bug-bar-aiml>

<sup>19</sup> Source: <https://docs.microsoft.com/en-us/security/engineering/threat-modeling-aiml>

<sup>20</sup> Source: <https://docs.microsoft.com/en-us/security/engineering/failure-modes-in-machine-learning>

<sup>21</sup> Source: <https://github.com/mitre/advmthreatmatrix/blob/master/pages/adversarial-ml-threat-matrix.md>. Latest commit: Oct 23, 2020

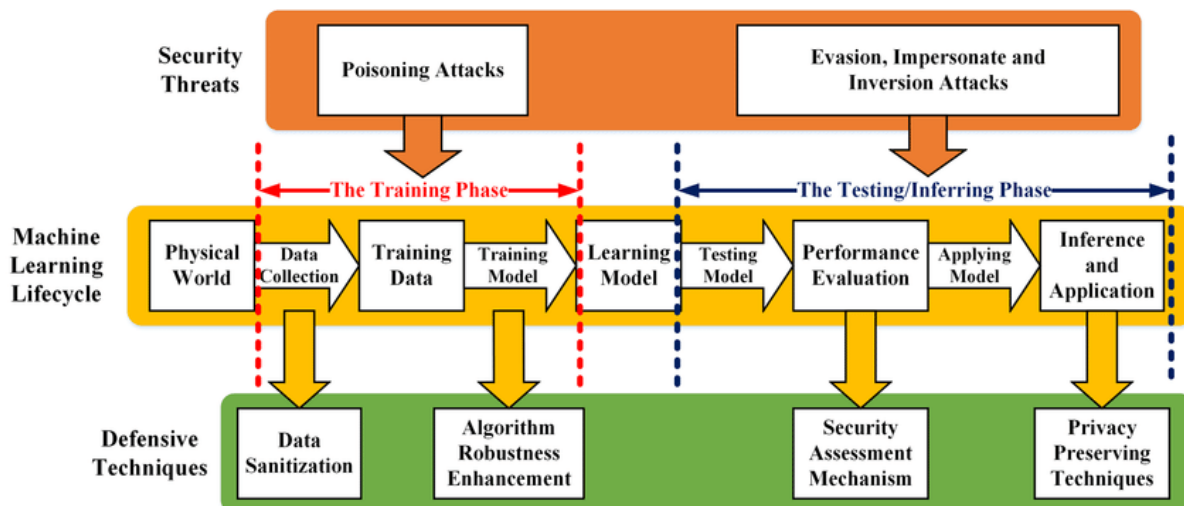


Figure 15 — Illustration of defensive techniques of machine learning

The 'Threat Modeling AI/ML Systems and Dependencies' questionnaires developed by Microsoft were used to capture the various aspects of the project and facilitate the security assessment. The 'Adversarial ML Threat Matrix' developed by MITRE was further used to focus the exercise on ML-specific techniques.

#### — System model for security assessment

Figure 16 — Modelisation of the 'climb and descent' ML system is a simplified modelisation of the interaction between the different elements of the system. It represents the principle data exchanges taking place in the system.



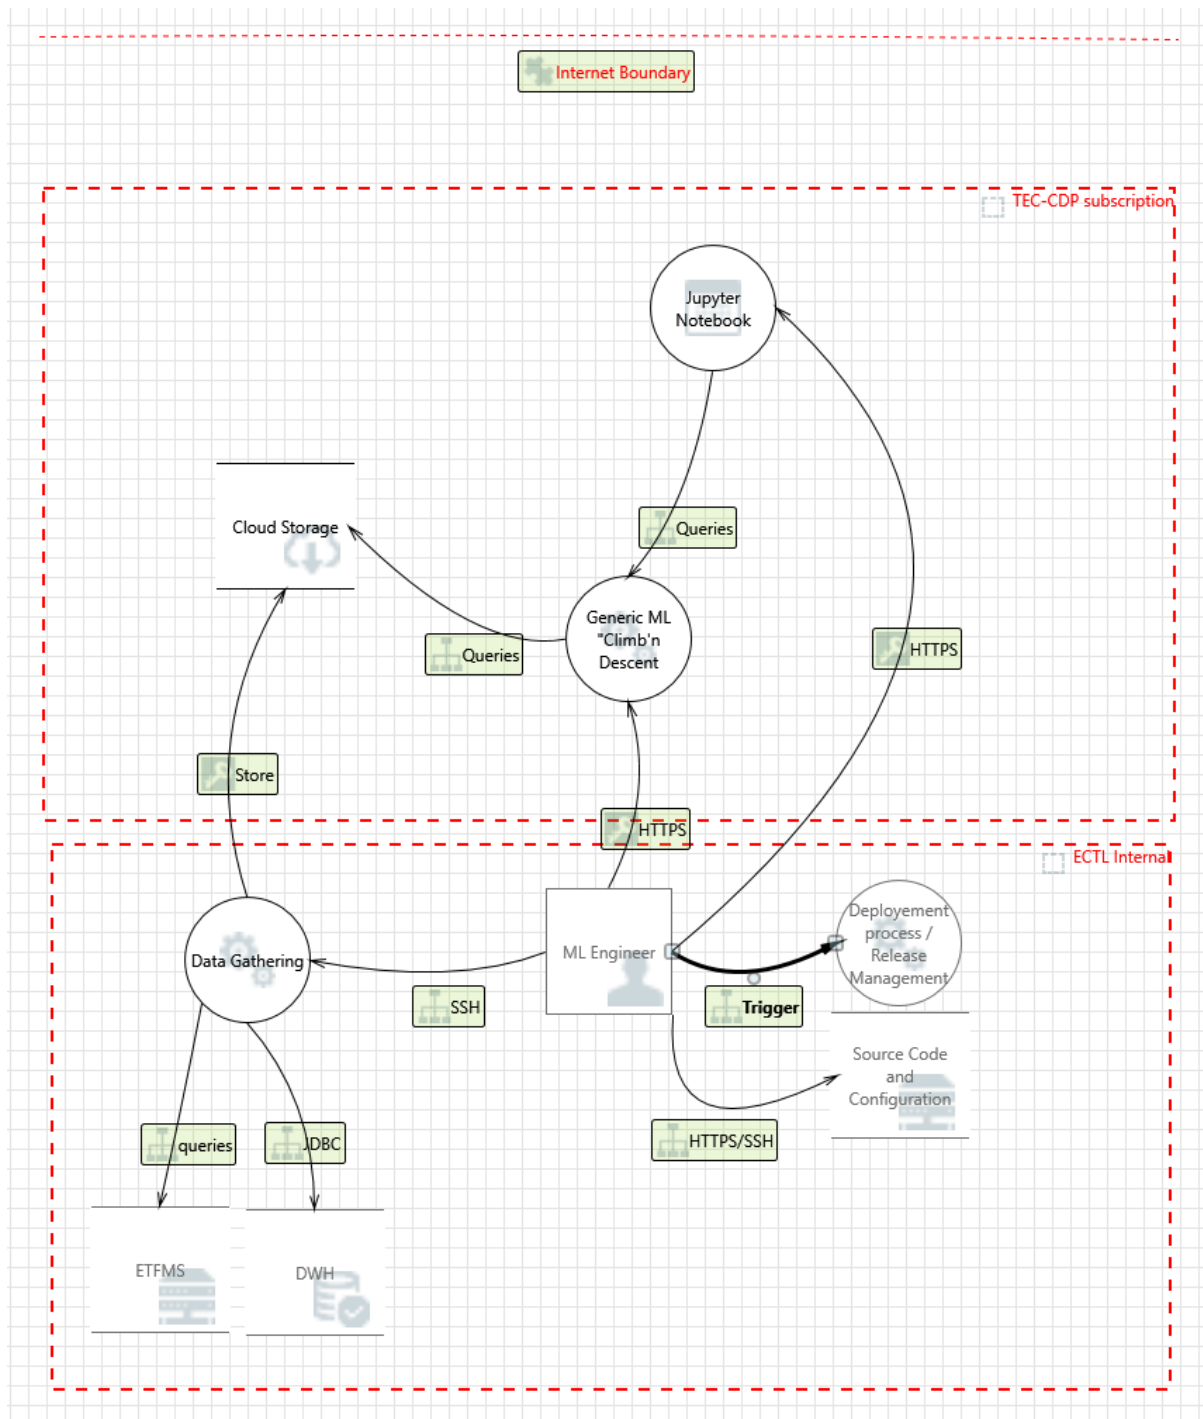


Figure 16 — Modelisation of the 'climb and descent' ML system

#### — Assumptions for security assessment

After an interview with the team in charge of the use case and considering the safety support case, the following considerations apply:

- The system considered is limited to the development phase of the model. The transfer to operations follows a dedicated workflow outside the scope.

- All data processed is post operations (no data confidentiality requirements, traffic light protocol (TLP):GREEN)
- The system is not considered as an operational system and does not present time-sensitive information.
- Safety support requirements and mitigations are in place, including the non-regression test.
- All involved communication networks are considered private with no interactive access to/from the internet.

Security and risks that are not inherent to the activities relating to the learning process are not considered in this assessment. Therefore, the applicable ratings for confidentiality, integrity and availability are:

- Confidentiality: Low
  - Integrity: High
  - Availability: Low
- Specific risks assessed
- Model poisoning: the threat was considered as mitigated by the assumptions: the isolation of the ML system vis-a-vis any external component whether from network or access permissions is considered sufficient mitigation.
  - Data poisoning of training data: the threat was considered as mitigated by the assumptions: the isolation of the ML system vis-a-vis any external component whether from network or access permissions as well as the controlled source for all training data is considered sufficient mitigation.
  - Model stealing: the threat was considered as mitigated by risk management: while there is no specific mitigation in place against the threat, it would not harm the organisation if it was to occur (no value loss).
  - Denial of service on any component: the threat was considered as mitigated by the operational model: unavailability of the training data or ML environment has no operational impact and only results in limited financial costs.

Other risks have been considered during the analysis but are not considered pertinent in view of the operational model in place (for example, defacement, data exfiltration, model backdoor, etc.).

#### 3.1.4. Learning assurance (in particular data management considerations)

**Objective DA-01:** The applicant should describe the proposed learning assurance process, taking into account each of the steps described in Sections C.3.3 to C.3.12, as well as the interface and compatibility with development assurance processes.

Most of the activities expected to be performed as per the 'learning assurance' have been executed. The following will make the demonstration of this statement.

### 3.1.4.1. Data preparation

#### a. Data collection

**Objective DM-03:** To enable the data collection step, the applicant should identify explicitly and record the input space and the operating parameters that drive the selection of the training, validation and test data sets.

##### — Data sources

Almost 3 years of data (from 01/01/2018 until 30/09/2020) was extracted from the NM Ops data warehouse from the ARU<sup>22</sup> schema. This contains basically all flights in the NM area for the last 3 years, and these were taken into the data set.

Weather information was taken from the UK Met office Sadis source, stored in the operational FTP server under the Met directory. EUROCONTROL has had a long-standing contract with the UK Met office to provide this data.

**Objective DM-04:** Once data sources are collected, the applicant should make sure that the data set is correctly annotated or labelled.

##### — Data labelling

The data labels<sup>23</sup> are also extracted from the ARU data set.

##### — Rates of climb and descent by performance slice

In a first step, the rate of climb between consecutive points of the point profile was calculated.

For a given flight phase, the time  $T$  for which a flight arrives at the flight level  $F$ , if there is no point at this flight level in the profile, can be approximated by linear interpolation:

$$T = T_{prev} + \frac{T_{next} - T_{prev}}{F_{next} - F_{prev}}(F - F_{prev})$$

where *prev* and *next* stand for the point of the profile respectively before and after the flight level.

If there is a point at the requested flight level, we simply use its time over.

<sup>22</sup> Due to the mainframe phase-out, this system was converted to Unix under the heading of the ARU System (Archive System on Unix). Once most functions were migrated to Unix, the system was renamed to Data Warehouse System (DWH).

<sup>23</sup> Data labelling is a key part of data preparation for machine learning because it specifies which parts of the data the model will learn from.




Figure 17 – Climb and descent rates by performance slice

### — Removing high frequency noise

It was observed that the calculated climb rates appear to have a lot of high-frequency noise overlaid on the signal and so we removed it by applying a low-pass filter to that in the form of a simple moving average window function of width 5.

### b. Data pre-processing

**Objective DM-06:** The applicant should ensure a sequence of operations on the collected data in preparation of the training.

#### — Data cleaning

Several data cleaning operations were performed, including the removal of yo-yo flights<sup>24</sup> (polluting the quality of the model), and the removal of the data corresponding to the cruise phase of the flight.

#### — Outliers

All data samples with climb rates that were calculated to be greater than 1 000 ft/min (likely to be not physically realistic and related to inaccuracy in the radar plots) were removed from the data set. Around 0.1 % of the 400 million samples were removed during this operation.

**Objective DM-08:** The applicant should ensure that the data is effective for the stability of the model and the convergence of the learning process, possibly via normalisation.

#### — Data normalisation

All data was normalised by centring on zero by subtracting the mean and given similar ranges by dividing by the standard deviation of that feature.

### c. Feature engineering

**Objective DM-07:** When applicable, the applicant should transform the pre-processed data from the input space into features which are effective for the performance of the selected ML algorithm.

<sup>24</sup> Yo-yo flight: flight with multiple climb and descent phases in the flight profile.

Feature engineering was managed via a pipeline. The pipeline's purpose is to enrich the data with various calculated features required for the subsequent operation.

Firstly, the SID and STAR are extracted from the flown route and attached to separate fields to the flight information so that they can be used as independent features.

The representations of coordinates in the database was string format rather than decimal format and these were converted into decimal degrees.

Several operations were made on the weather forecast data source. For more information, please refer to the full report available by EUROCONTROL.

Several additional calculated weather-forecast-related features were then produced, namely wind speed and wind direction relative to the aircraft.

Some further features were then added. It was discovered that using the latitude and longitude of the aerodrome of departure and destination as well as the first and last point of the climb and descent was more effective than any other encoding of these values. For example, an embedding layer was used to encode the categorical values e.g. the ICAO names for aerodromes of departure and destination, but this was not nearly as effective as the vector encoding as latitude and longitude.

This resulted in a model with some 40 features which was saved in a parquet file which when loaded was around 100 gigabytes in RAM.

The permutation importance (a similar method is described in Breiman, 'Random Forests', Machine Learning, 45(1), 5-32, 2001) for these features was then calculated. This was a very heavy calculation taking several days on a GPU to complete.

Permutation importance:

Climb		Descent	
Weight	Feature	Weight	Feature
494468.7164 ± 269.1501	PERF_CAT_LOWER_FL	392129.5391 ± 248.7002	PERF_CAT_LOWER_FL
217568.8688 ± 138.2701	FTFM_CLIMB_RATE	211356.3405 ± 95.6282	FTFM_DESC_RATE
138494.9605 ± 44.0213	FTFM_MAX_FL	133131.4453 ± 68.8156	FTFM_DESC_FIRST_PT_LAT
114020.7645 ± 86.3738	FLT_DEP_AD	85637.1216 ± 64.1071	FTFM_DESC_LAST_PT_PT_LAT
109271.3590 ± 243.7783	FLT_DEP_AD_LAT	85262.9041 ± 138.5218	FLT_FTFM_ADES_LAT
105701.0231 ± 96.9098	FTFM_CLIMB_FIRST_PT_LAT	80916.0368 ± 71.9405	FLT_FTFM_ADES
95154.7142 ± 86.0832	ICAO_ACFT_TY_ID	72740.5408 ± 34.9251	FTFM_DESC_FIRST_PT_LNG
86846.6291 ± 88.8068	FTFM_CLIMB_FIRST_PT_LNG	70372.2655 ± 109.2796	FTFM_DESC_LAST_PT_LNG
86710.6489 ± 193.9731	FLT_DEP_AD_LNG	69247.5777 ± 83.0451	FLT_FTFM_ADES_LNG
23296.1818 ± 26.1849	FTFM_CLIMB_DURATION	43342.9997 ± 56.8700	FTFM_MAX_FL
21731.4291 ± 59.1714	AO_ICAO_ID	37916.0572 ± 130.2117	FTFM_DESC_DURATION
20337.5237 ± 73.7881	FTFM_CLIMB_FIRST_PT	32727.9660 ± 55.2942	FTFM_DESC_LAST_PT
18971.2889 ± 22.4656	FLT_FTFM_ADES_LAT	12746.5049 ± 19.2558	ETA_DAYOFYEAR

Climb		Descent	
Weight	Feature	Weight	Feature
18136.2638 ± 26.9874	FLT_FTFM_ADES_LNG	11355.1165 ± 65.0552	AIRAC_CYCL
18026.4043 ± 22.2186	FTFM_DESC_LAST_PT_PT_LAT	9524.1099 ± 37.4795	ICAO_ACFT_TY_ID
16417.4972 ± 20.0458	FTFM_DESC_LAST_PT_LNG	6437.3164 ± 30.2539	AO_ICAO_ID
15343.8757 ± 44.8245	ETA_DAYOFYEAR	5731.4322 ± 19.5940	FLT_REG_MARKING
15176.5899 ± 32.8208	FLT_REG_MARKING	5658.8823 ± 21.7385	FTFM_CLIMB_FIRST_PT_LAT
15034.2075 ± 24.5128	FTFM_CLIMB_LAST_PT_LNG	5400.5508 ± 40.4232	FTFM_CLIMB_LAST_PT_LNG
14964.0634 ± 29.0470	FTFM_CLIMB_LAST_PT_LAT	5119.9972 ± 15.9033	FTFM_CLIMB_FIRST_PT_LNG

Table 9 – Extract of candidate features by importance (20 out of 40)

When the permutation importance of a feature is low, this means the feature is not very decisive for obtaining a result.

#### d. Hosting for data preparation and model training

Data preparation was hosted under Microsoft Azure. The model training was hosted in a Cloudera Machine Learning (CML) environment. This is Cloudera's cloud-native ML service, built for CDP. The CML service provisions clusters, also known as *ML workspaces*, that run natively on Kubernetes.

ML workspaces support fully-containerised execution of Python, R, Scala, and Spark workloads through flexible and extensible *engines*.

This facility allows automating analytics workloads with a job and pipeline scheduling system that supports real-time monitoring, job history, and email alerts.

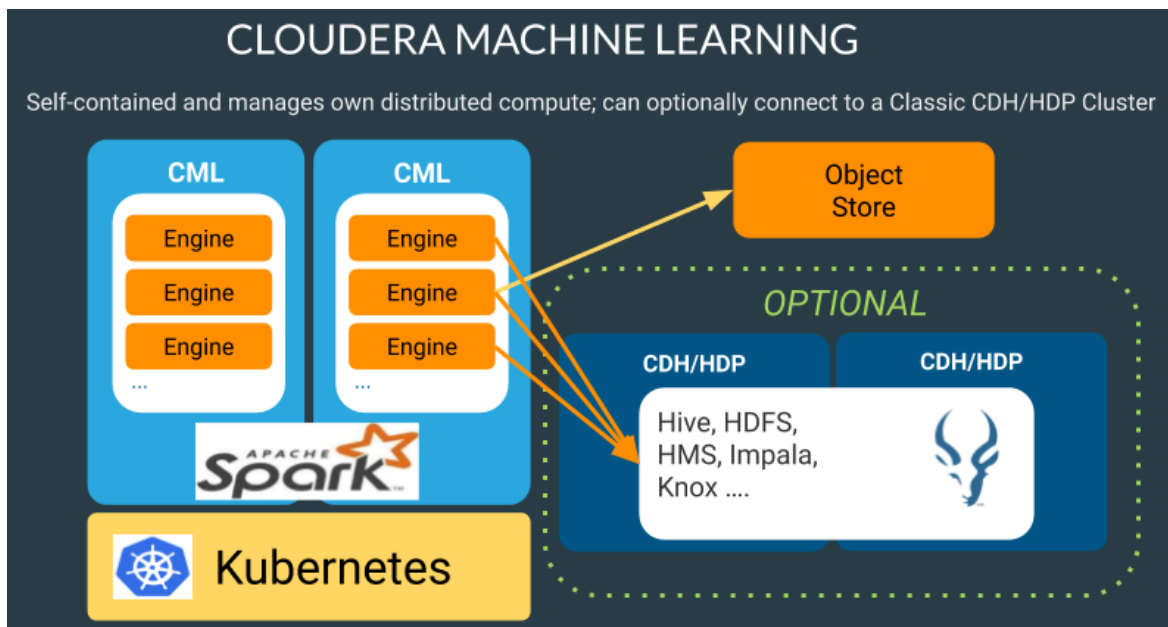


Figure 18 – Cloudera machine learning environment

For more information, please refer to the full report available by EUROCONTROL, or contact the teams at EUROCONTROL in charge of such an environment.

### 3.1.4.2. Data validation

#### a. Data completeness

**Objective DM-10:** The applicant should ensure validation and verification of the data all along the data management process so that the DQRs are addressed.

The period which has been considered for the data in the data set (3 years of archived data from the DWH), and the inherent quality of the DWH via its usage by thousands of stakeholders on a daily basis, ensure the completeness of the data for the use case.

#### b. Data accuracy

Data accuracy has been established through the different activities performed during the data management phase. In particular, incorrect or non-representative data has been removed from the data set during data cleaning (e.g. removal of yo-yo flights), or when identifying outliers (flights with unrealistic climb or descent rates).

#### c. Data traceability

All operations performed on the source data set extracted from the DWH were orchestrated via scripting and pipelining in different python modules. All code is under configuration management, ensuring full traceability and capability to reproduce featured input and labelled data for subsequent training.

#### d. Data representativeness

The 4D trajectory applies to the ECAC area. The DWH archives all information which has been processed by IFPS/ETFMS, then ensuring that the data set fully covers this geographical area.

#### e. Data allocation — data independence

**Objective DM-09:** The applicant should distribute the data into three separate and independent data sets which will meet the specified DQRs:

- the training and validation data sets, used during the training and learning process validation;
- the test data set used during the learning process verification, and the inference model verification.

There are roughly 370 million data samples in the data set. The test set was chosen at random and had 5 % set-aside.

The validation set was a further 20 % of the remaining.

Considering the large amount of data samples, keeping 5 % of all data for the test set represents 25 million samples in the test data set, which is enough to provide a statistically valid result. The same remark applies to the validation data set.

### 3.1.4.3. Learning process management

**Objective LM-01:** Objective LM-01: The applicant should describe the AI/ML components and model architecture (including computational graph and activation functions).

**Objective LM-02:** The applicant should capture the requirements pertaining to the learning management and training processes.

#### a. Model selection

A DNN was selected.




Figure 19 – DNN structure

Multiple architectures were tested during hyper-parameter tuning. The most successful architecture for the hidden layers was as follows.

Layer number	number of neurons
1	512
2	512
3	256
4	256
5	128
6	64

Table 10 – Internal architecture of the DNN



The table below summarises the main decisions/configurations made/applied at the end of the training process:

Title	Information / Justification
<b>Activation function</b>	The <b>PreLU</b> activation function was chosen for a number of its advantages in DNNs; particularly, avoidance of the vanishing gradients problem as was the case with standard ReLU, but in addition the avoidance of the dying neuron problem
<b>Loss function selection</b>	Several loss function strategies were studied during the learning and training process. Finally, it was decided to use ' <b>mean absolute error</b> ' which appears to give the best results on the test set
<b>Initialisation strategy</b>	The <b>Glorot</b> initialisation technique was chosen for initialising the values of the weights before training
<b>Hyper-parameter tuning</b>	Hyper-parameter tuning was a recurrent activity all along the learning process management and the model training

Table 11 — Key elements of the DNN

#### b. Hosting the model predictor

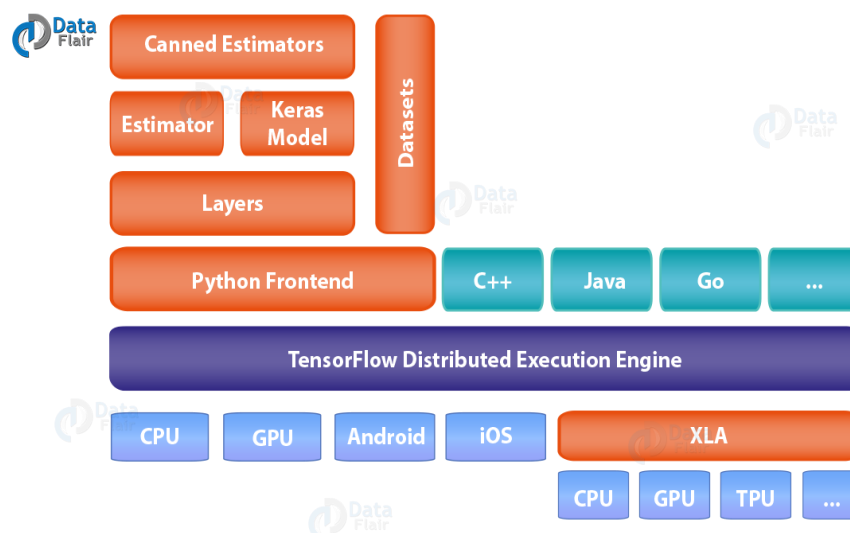


Figure 20 – Tensorflow component model and dependencies

The above diagram represents the TensorFlow component model and dependencies. The predictive models were developed using Keras Python interfaces to TensorFlow — see above on the left side.

The model training pipeline based on Python and Keras produces a saved model in protobuf format and associated model weights files. This is done in the cloud as described above.

### 3.1.4.4. Model training

#### a. Feature set

The following table represents the current list of features which were used for the training:

Feature	Feature
AO_ICAO_ID	float32
ETA_DAYOFYEAR	float32
FLT_DEP_AD_LAT	float32
FLT_DEP_AD_LNG	float32
FLT_FTFM_ADES_LAT	float32
FLT_FTFM_ADES_LNG	float32
FLT_REG_MARKING	float32
FTFM_CLIMB_RATE	float32
ICAO_ACFT_TY_ID	float32
PERF_CAT_LOWER_FL	float32

Table 12 — List of features as an input to model training

**Objective LM-05:** The applicant should document the result of the model training.

#### b. Learning curves

The figure below depicts a learning curve when using the feature set and the labelled data:




Figure 21 — Model training (mean absolute error)

### 3.1.4.5. Learning process verification

**Objective LM-09:** The applicant should perform an evaluation of the performance of the trained model based on the test data set and document the result of the model verification.

#### a. 3D histogram plot of the predicted values

The below figures show two plots for all of the test data of climb rate predictions against actually observed climb rates.

The dispersion is greatly reduced with the trained ML model.




Figure 22 – Predicted climb rate (with BADA) v actual from CTFM




Figure 23 - Predicted climb rate (with ML) v actual from CTFM

#### b. Comparison of error rates between current (FTFM) and new ML calculation




Figure 24 — Mean square error on actual climb rates (with low-pass filter)

### 3.1.4.6. Implementation

**Objective IMP-03:** For each transformation step, the environment (i.e. software tools and hardware) necessary to perform model transformation should be identified and any associated assumptions or limitations captured and validated.

#### a. System architecture

Depending on the context where the 4D trajectory calculation is performed, the AI/ML library could be called from different processes. The following is the logical architecture of ETFMS. The 4D trajectory is calculated within the ‘profiler process’:

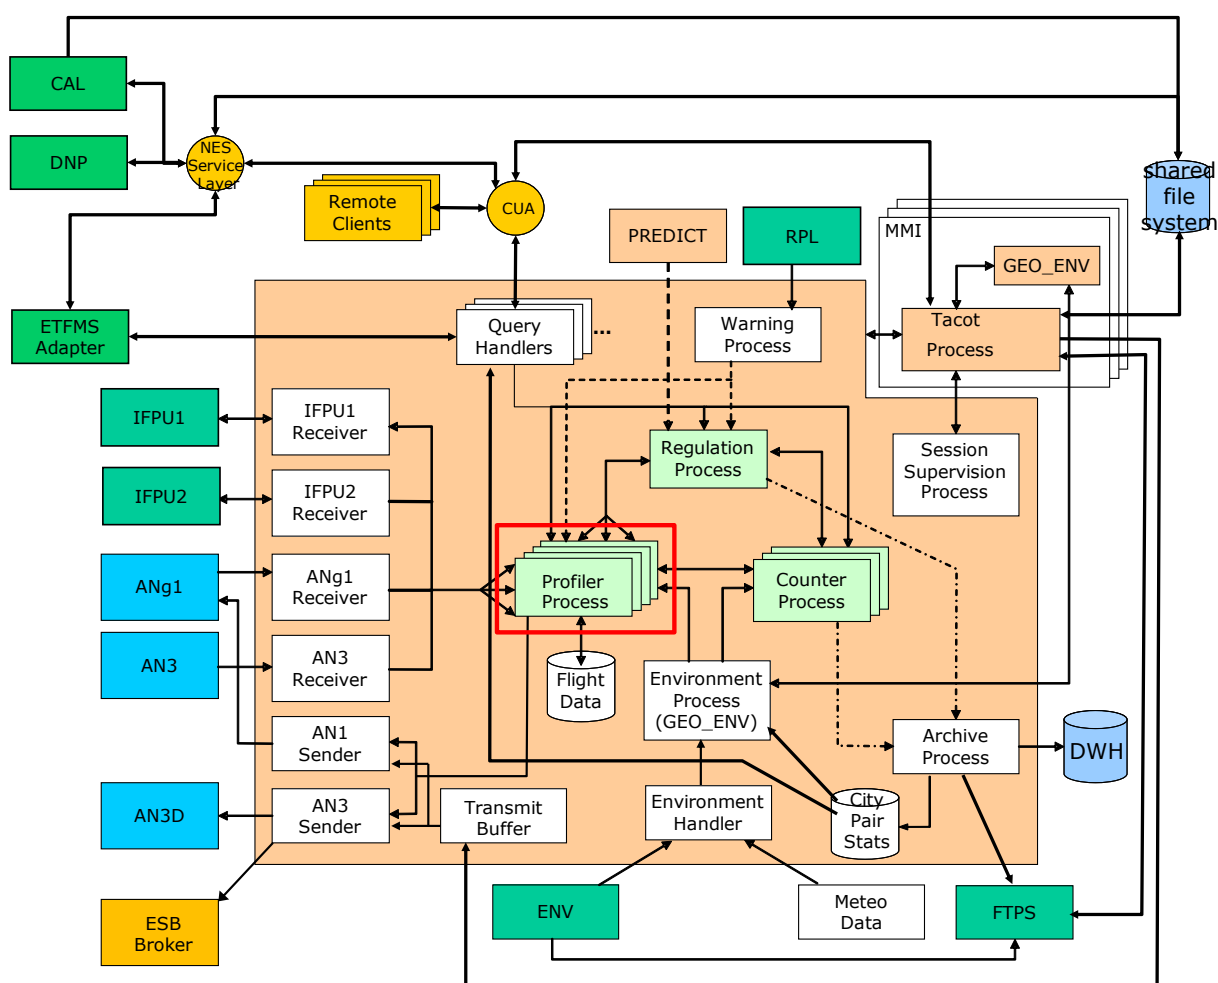


Figure 25 — ETFMS logical architecture

The ‘profiler process’ computes the flight profile or 4D trajectory. For performance reasons, several processes can co-exist in ETFMS. An algorithm statically associates a flight with a ‘profiler process’ to allow parallelism.

The ‘profiler process’ is mission-critical. Its failure induces an ETFMS failure.

The flight load is distributed equally by a hashing algorithm amongst the number of ‘profiler processes’. Once a flight has been associated with a given instance of a process, for the sake of data consistency, this instance is the only one that manages the flight; all messages relating to the flight are directed to it.

The ‘profiler process’ embeds the Curtain software package.

The Curtain software package has been adapted to use the AI/ML component.

## b. AI/ML component as a library

### — General information

A prediction is a numerical value provided by a TensorFlow model. The inputs are an ordered list of fields and, usually, after transformation and normalisation, are passed to the model which returns a value, the prediction. The library should be supported with additional information: the TensorFlow model resulting from training, the statistics from the training data (mainly mean and standard deviation) used by the normalisation, and the conversion from categorical value to numerical value used to include categories in the prediction. The library is also configured with a description of the fields, categories, eventual ways to validate the input and output, and, in the case of invalid input, how to replace them by acceptable values.

A prediction is provided by a predictor. The API lets the user create and register one or more predictors with a given name. It is possible to remove an existing predictor but also to swap two predictors (they exchanged their names) as a shortcut to remove and re-create. Creation implies moving in memory several lookup tables, so swapping can improve performance in some cases.

Each predictor is linked to one or more TensorFlow models, provided as TensorFlow .pb and checkpoint files.

As a lot is triggered by configuration, there is a function in the API to print the global configuration (input data and pre-computed lookup tables) from a predictor. Another function will try to analyse the predictor in order to see if it is consistent (at least one model, at least one field, etc.).

The API is a C API and will provide different functions, structures to represent input data and enumerations for code values.

### — Workflow

The library implemented a workflow which is generic and can be reused for different AI/ML use cases.

The figure below depicts the workflow for prediction which was implemented:

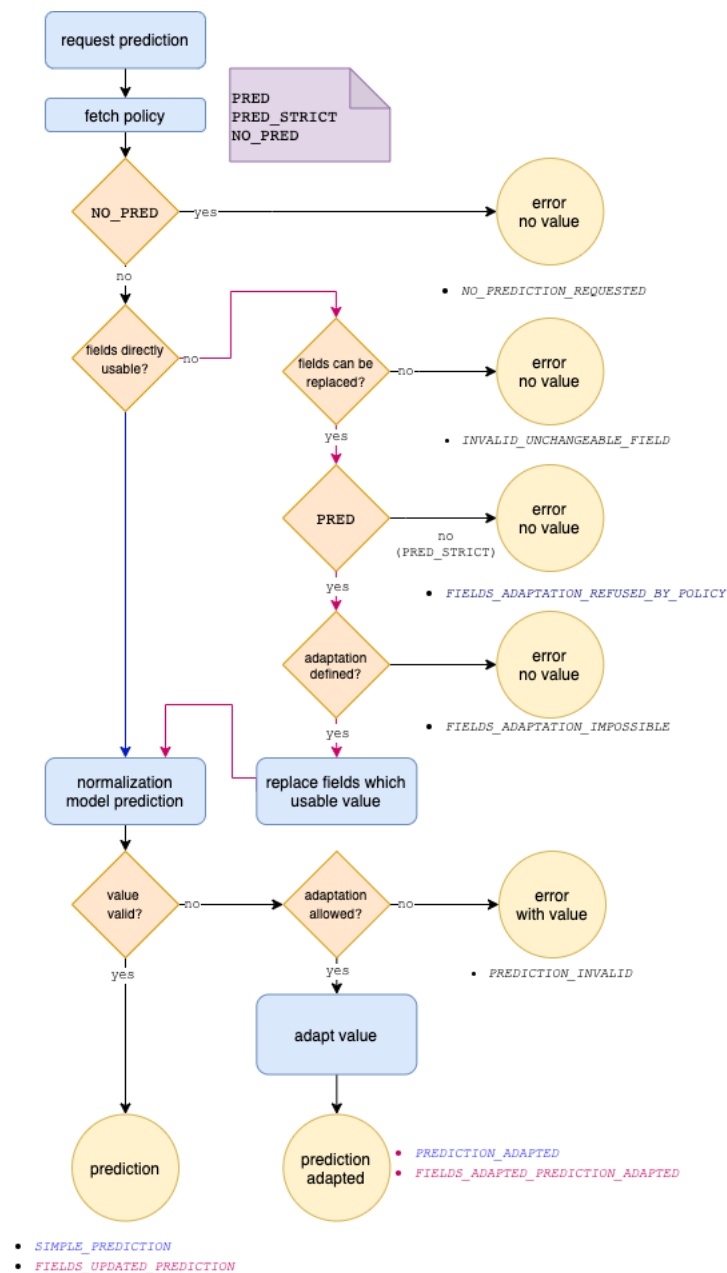


Figure 26 — Workflow for prediction delivered by the AI/ML library

The saved models were used in the ETFMS operational context via the C/C++ API.

This library was delivered to the ETFMS, and an Ada binding is produced so that the predictions could be provided by a simple in-process call in the same address space.

The reason for this is the need for very low latency and high bandwidth to ML predictions as the trajectory calculations in ETFMS are particularly performance-sensitive. It is not feasible or desirable to use a traditional technique of providing a ML REST-based server to provide the predictions as the latency of the network connection would make the predictions useless in this context.

### c. Executable model architecture

The following depicts part of the NM operational infrastructure running the ETFMS.

ETFMS is located on-premise. It is part of a mission-critical cluster (called RED). ETFMS operational instance (ptacop1 or ptacop3) is part of the sub cluster RED\_03, which contains four virtual machines (red011 to red014).

These virtual machines are based on **Linux Red Hat Enterprise Server** Operating System.

All the virtual machines of the RED Cluster are spread over the 6 **HP DL560 G10** Servers (rambo, rocky, rufus, romeo, rusty, roger), all based on 36 CPS and 1,5 TB of RAM.

The hypervisor used to manage the virtual machines is **VMWare ESXi**.

Storage for this cluster is spread over NASes (rednas10 to 13).

#### 3.1.4.7. Inference model verification

**Objective IMP-06:** The applicant should perform an evaluation of the performance of the inference model based on the test data set and document the result of the model verification.

##### a. Verification of improvements at network level

The most appropriate way to assess the performance of the AI/ML component was to analyse the impact on the network situation. This analysis is possible based on some tools capable of replaying specific situations which have occurred in the past (also known as PREQUAL). For more information, please refer to the full report available by EUROCONTROL.

The table below demonstrates significant improvements on the network for two separate dates in 2020:

	10/08/2020		14/02/2020		18/12/2020	
	Avg	Max	Avg	Max	Avg	Max
<i>BL</i>	283 051	1 860 046	544 428	3 226 165	285 194	1 783 651
<i>ML</i>	265 889	1 655 747	514225	2 880 420	272 071	1 632 486
<i>Improv.</i>	6,06 %	10,98 %	5,54 %	10,71 %	4,60 %	8,48 %

Table 13 — Improvements on the network

**Objective IMP-07:** The applicant should perform a requirements-based verification of the inference model behaviour and document the coverage of the ML component requirements by verification methods.

In addition to verification of the improvement brought at network level, verification activities have taken place from various perspectives, including system resilience.

## b. Robustness

**Objective IMP-08:** The applicant should provide an analysis on the robustness (or stability) of the inference model.

At the date of this report, the robustness of the AI/ML component remains to be investigated. It will be progressively assessed via additional testing at the limits (e.g. how will the model perform when being faced to abnormal data like an unknown airport or unknown aircraft type).

## c. Resilience

Based on the system requirements identified for Curtain, and the target architecture, should the model face robustness limitations, then the legacy climb and descent computation would continue to deliver the service even in a less performant mode of operation. All these measures ensure resilience at system level.

## 3.2. Time-based separation

### 3.2.1. Trustworthiness analysis — description of the system and ConOps

#### 3.2.1.1. Description of the system

**Objective CO-02:** Define the AI-based system taking into account domain-specific definitions of 'system'.

Headwind conditions on final approach cause a reduction of the aircraft ground speed which for distance-based separation results in increased time separation for each aircraft pair, a reduction of the landing rate, and a lack of stability of the runway throughput during arrival operations. This has a negative impact not only on the achieved capacity, but also on the predictability of operations, time and fuel efficiency, and environment (emissions). The impact on predictability for core hubs is particularly important at the network level. The service disruption caused by the reduction in achieved runway throughput compared to declared capacity in medium and strong headwinds on final approach has a significant impact on the overall network performance. It is also particularly exacerbated if this occurs on the first rotation of the day because of the impact on all the other rotations throughout the day.

This service disruption also has a significant impact on airline operations due to delayed and cancelled flights. The time-based separation (TBS) solution addresses this problem by defining procedures and specifying user and high-level system requirements to allow stable arrival runway throughput in all headwind conditions on final approach. However, TBS application entails the use of a support tool providing the separation distance indicator depending on the applicable time separation minimum, and the follower speed profile which also depends on the headwind conditions.

In order to allow efficient separation delivery, the support tool also advises the air traffic controller on the expected compression (i.e. spacing reduction because of the catch-up of the follower still at a high glide slope speed when the leader decelerates to its final approach speed) which depends on both the leader and follower speed profiles. This compression indicator enables better separation conformance with homogenisation of the ATCO performance in terms of separation delivery and workload. This is needed when dealing with variable separation minima as it is the case when applying



TBS, but is also of great interest when applying distance-based separation; for instance, in challenging wind conditions where compression effects may be difficult to anticipate for the ATCOs.

Finally, the separation delivery tool is able to account for and to provide indication of more complex business rules to ATCOs than what they use to have today: e.g. runway occupancy time, spacing gap for departure, pair-wise wake separation minima, weather-dependent separation, separation increase/reduction in case of enhanced approach procedure operation.

### 3.2.1.2. Concept of operations

**Objective CO-03:** The applicant should define and document the ConOps for all AI-based systems. A focus should be put on the definition of the operational design domain (ODD) and on the capture of specific operational limitations and assumptions.

1. Final target distance (FTD) — Whatever the mode in use (time- or distance-based), time separations or spacing need to be translated into distance for being visualised on the ATCO screen. This translation of time into distance is based on aircraft true airspeed profile and headwind. The FTD indicator corresponds to the minimum distance separation to be applied between leader and follower, when the leader is overflying the separation delivery point (DP) (e.g. the runway threshold). The FTD shall account for all applicable separations and spacing constraints in the prevailing wind conditions. This includes the applicable wake time- or distance-based separation minimum (DBS or TBS), the leader runway occupancy time (ROT), and the minimum surveillance spacing (minimum radar separation (MRS)). It is computed based on the previously mentioned: TBS (or DBS), ROT and MRS minima and the time-to-fly profile of the follower aircraft in the prevailing wind conditions.
2. Initial target indicator (ITD) — TBS allows for applying optimised and consequently more variable separations (i.e. when applying TBS, the separation minima are indeed defined on a pair-wise basis and, depending on the wind, leading to more variability in the minima compared to a category-based DBS scheme). The compression effect is therefore more complicated to anticipate. The ITD is giving the ATCO information about how to separate the aircraft at a pre-agreed speed, typically 160/170 kt (TAS), in order to have separation reserve for anticipating on compression effect. It is the distance separation applicable when the leader aircraft is at a prescribed glide speed before deceleration to final approach speed such that the FTD will be obtained at the separation DP. It is thus computed using as input the FTD and the leader and follower time-to-fly profiles in the prevailing wind conditions.

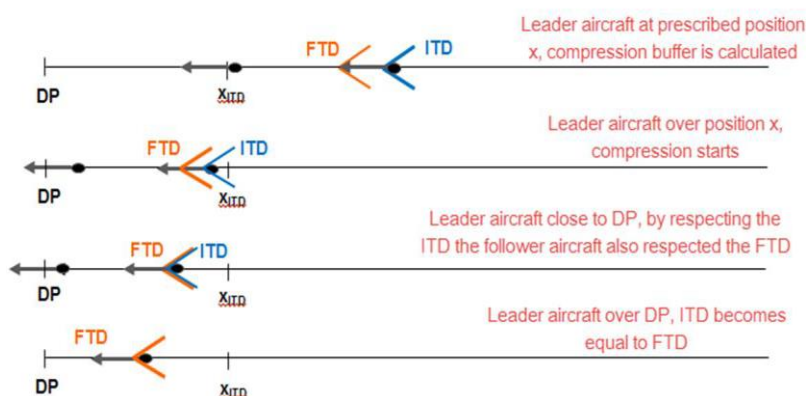


Figure 27 — Schematic view of FTD and ITD definition

The Calibration of Optimised Approach Spacing Tool (COAST) is a service supported by EUROCONTROL which provides the models and a methodology to use them in order to compute the values of two indicators as inputs to a separation delivery tool as illustrated in Figure 27 — Schematic view of FTD and ITD definition. A deeper explanation is provided in EUROCONTROL COAST (Calibration of Optimised Approach Spacing Tool), White Paper, Ed V1.1, 16/04/2021 (EUROCONTROL, 2021).

### 3.2.1.3. Description of the system(s) involved (inputs, outputs, functions)

The detailed COAST pipeline is shown in Figure 28 — COAST development cycle.

1. Initial execution of the COAST pipeline can be summarised as follows: raw flight and meteorological data is first supplied by the ANSP. Data processing operations (denoising, filtering, etc.) are applied before storing the transformed data in the flights database. The database is split in two data subsets, called training and test, which are given as an input to the ML toolbox. The toolbox computes time-to-fly and safety buffers models using the training set and calculates the coverage on the test set. For every model, a model performance report is generated. It includes information about:
  - time-to-fly estimation quality;
  - buffers estimation accuracies;
  - coverage functions explanation.

A full history of the computed models and of their accuracies is maintained for the sake of performance monitoring. All models are stored in a registry that is a database linked to a file server. The model database also contains the performance in terms of accuracy, coverage, computation details, etc.

In a maintenance context, when new data is available, the pipeline is similar; except that an automatic data validation is performed before effectively adding the extract to the flights database. This data validation process, further detailed in Section 5.5, allows a maintenance of the model validation and the detection of any deviation of the new data compared to the previous ones.

2. The second component is deployed on the airport side. It is dedicated to the decision process where the models produced by COAST are used. Using this methodology, the separation delivery tool provides the computation of the FTD and ITD indicators accounting for flights and weather data in the operational environment (as specified in the spec5). The separation delivery tool invokes the prediction models with information about the arriving couples of flights and weather conditions. The models allow the tool to compute the time-to-fly profile and the safety buffers supporting the computation of FTD and ITD separation indicators.

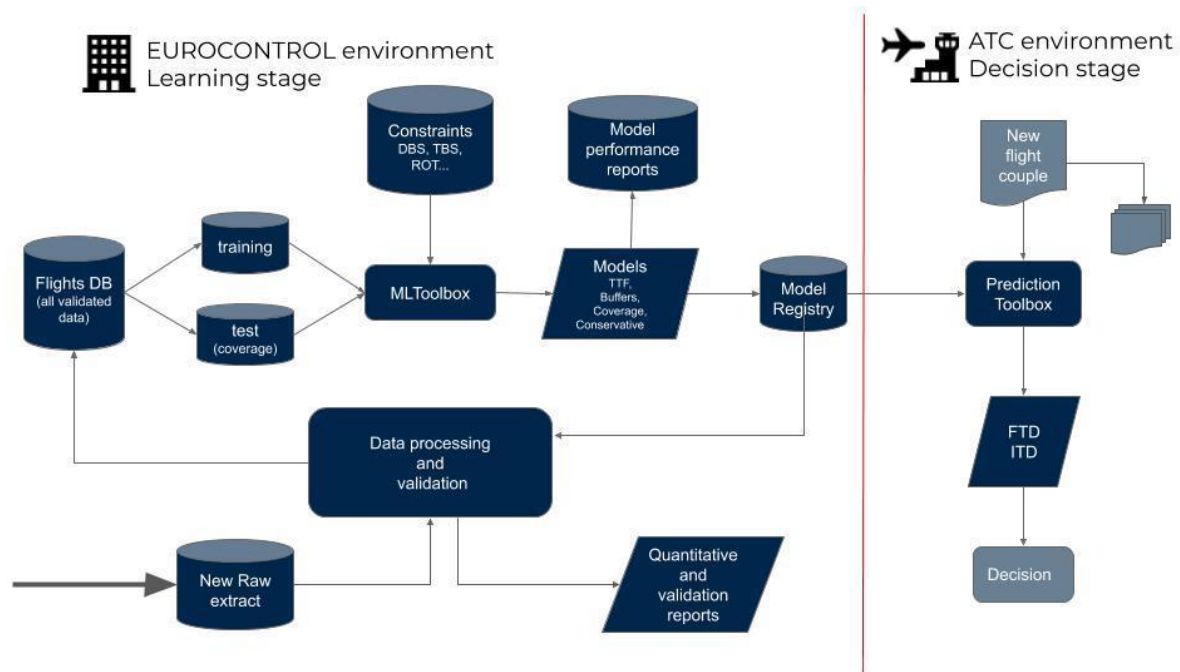


Figure 28 — COAST development cycle

### 3.2.1.4. Expected benefits and justification for Level 1B and Level 2A

The use of ML algorithms to calculate time-based separation and spacing indicators makes it possible to use statistical behaviour in their computation and hence harmonising error rates, enhancing prediction accuracy, which results in improved operational efficiency.

The benefits related to the use of FTD and ITD are further illustrated by providing three scenarios:

1. **Current situation with no optimisation:** In this scenario, there is no FTD (allowing dynamic separation reduction) nor ITD (providing optimised spacing indication). A conservative spacing buffer is thus applied before leader deceleration (starting at deceleration fix (DF)) in order to cope with compression uncertainty resulting in a separation delivered at threshold showing some margin compared to the minimum.

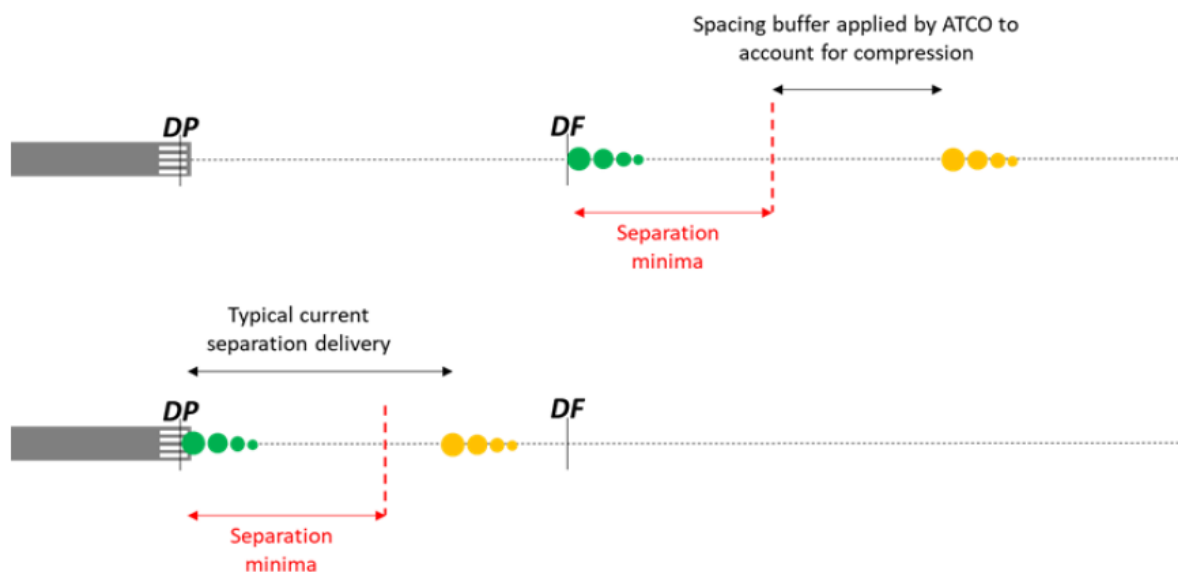


Figure 29 — Current ATCO support tool for separation and spacing

2. Use of ITD without change of separation minima mode: In this scenario, the use of the ITD allows optimised spacing of the flight before leader deceleration (starting at DF) resulting in a DBS delivered at threshold with higher accuracy.

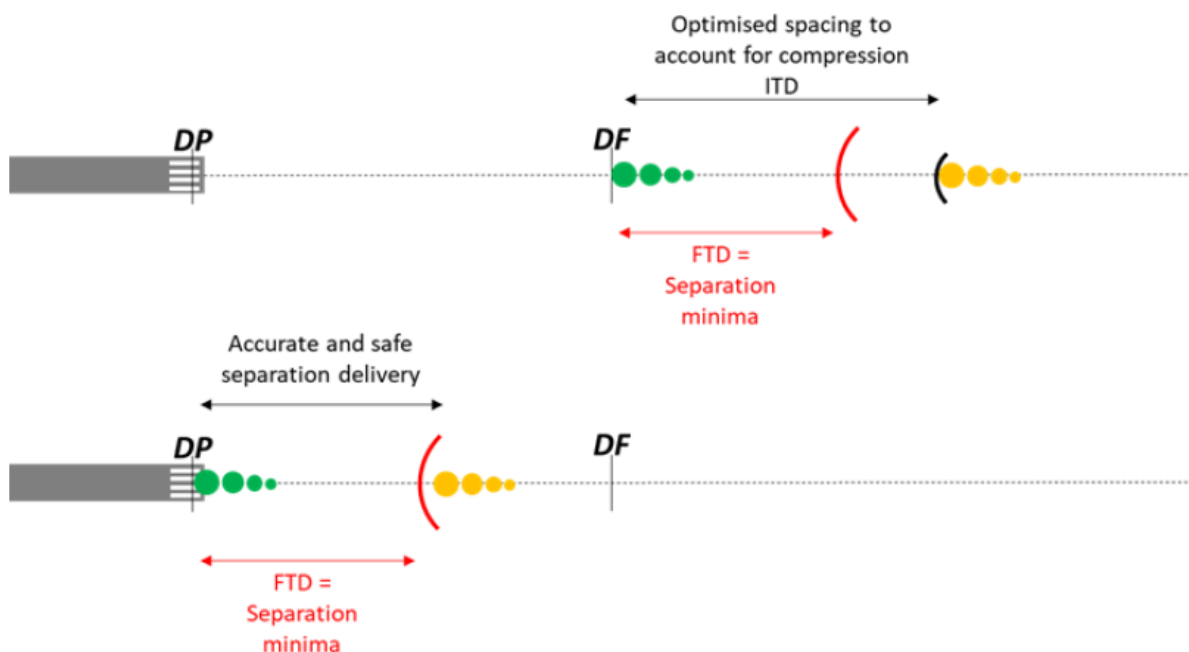


Figure 30 — DBS mode + ATC support tool for separation and spacing

3. Use of ITD with reduced separation minima: In this scenario, the use of FTD, allowing dynamic separation reduction (e.g. applying TBS) combined with ITD, improving the separation delivery accuracy, shows significant decrease in the delivered separation.

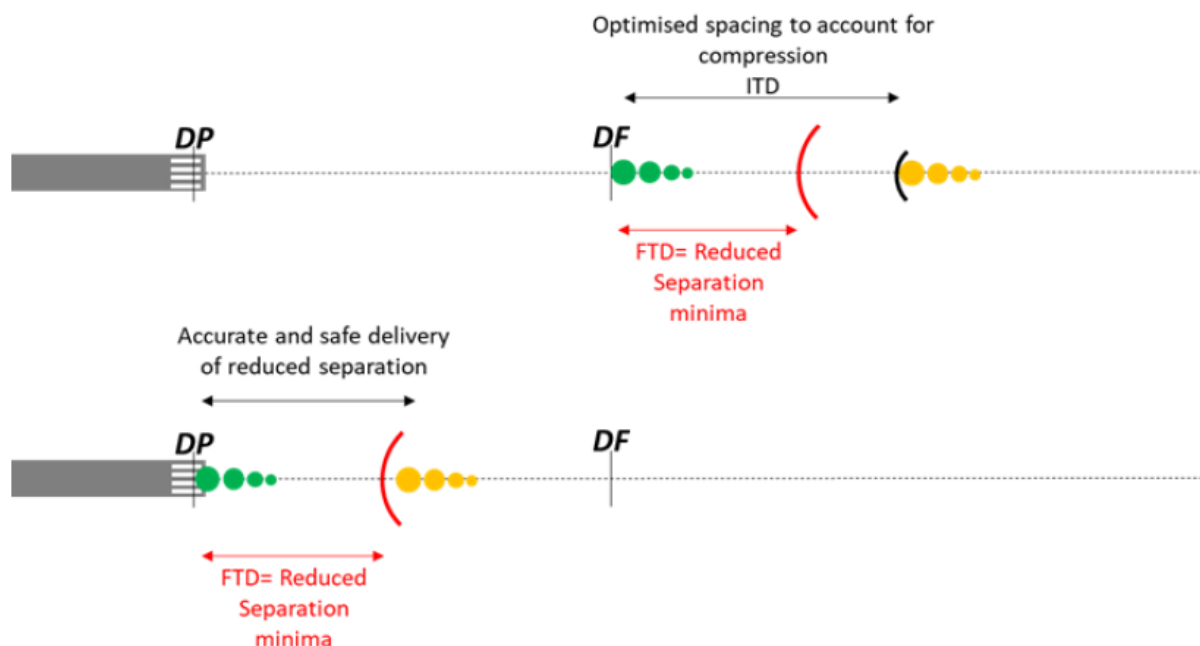


Figure 31 — TBS mode + ATC support tool for separation and spacing

**Objective CL-01:** The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.

Three applications were identified during use case development:

*a. Distance-based separation — optimum runway delivery (DBS — ORD)*

In this mode of application, the separation is based on distance as usual, but the tool provides an ITD indicator to the air traffic controller allowing the optimal spacing of aircraft on final approach.

The system assists the human in this case. The decision is solely the task and responsibility of the air traffic controller. Therefore, this application is **Level 1B**.

*b. Time-based separation (TBS)*

In this mode of application, the separation is based on time. The separation to be applied by the air traffic controller is indicated on the CWP as a final target distance (FTD). The controller has to follow the indicator, as (s)he has no other means to verify the distance between two aircraft on final approach.

*c. Time-based separation — optimum runway delivery (TBS — ORD)*

In this mode of application, the separation is based on time, similarly as in the previous mode, but the tool provides an ITD indicator to the air traffic controller allowing the optimal spacing of aircraft on final approach.

The system in both TBS cases provides the human with the sole information on which the decision is based. Strictly speaking, the decision task and responsibility is still lies on the air traffic controller, but as (s)he cannot override the system because of the lack of information that would allow such an override, these applications are **Level 2**.

#### 4. Use cases — aircraft production and maintenance

It should be noted that maintenance to assure continuing airworthiness of products is divided into two fundamentally different levels of activity:

- **Planning and scheduling of maintenance tasks:** this is typically done in by continuing airworthiness management organisations (CAMOs).  
In the generic wording of GM M.A.708(b)(4) ‘the CAMO is responsible for determining what maintenance is required, when it has to be performed, by whom and to what standard in order to ensure the continued airworthiness of the aircraft.’, to determine *what* and *when* is currently decided based on fixed maintenance schedules and monitoring mainly simple usage parameters of the aircraft (e.g. flights, flight hours, calendar time), also including a regular update of the maintenance schedule taking into account in-service experience.  
Modern aircraft providing an enormous amount of data in service and other information available (e.g. environmental data) does now provide a data pool which would allow scheduling maintenance much more appropriately and individually; however, to evaluate such big amount of data, sophisticated algorithms are required potentially containing AI/ML elements.
- **Performance of maintenance:** this is typically done by approved maintenance organisations (often also referred to as Part-145 organisations, as they are covered in Part-145).  
During performance of more complex maintenance tasks, it is normal to make use of special test equipment, today often including software. The use of test equipment containing AI/ML has a high potential to improve the quality of tests and inspections, while also improving efficiency.

In both domains, AI systems could be used to augment, support or replace human action, hence two examples are given.

##### 4.1. Controlling corrosion by usage-driven inspections

###### 4.1.1. Trustworthiness analysis

###### 4.1.1.1. Description of the system

Currently the so-called corrosion prevention and control programmes (CPCP) managed at fleet level do control corrosion by scheduled inspections implemented at a fixed threshold and performed at fixed intervals, which are from time to time adjusted depending on the severity of corrosion found during previous inspections.

Today we have detailed data about where the aircraft has been at which point in time, which temperature, rainfall, de-icing agents, corrosion-critical pollutants, etc. it has been exposed to, how it has been utilised, which corrosion findings have been made on other aircraft, and a lot of other usage, utilisation, maintenance, repair, events etc. it has experienced. From this huge data pool an AI algorithm could be trained to evaluate the individual corrosion risk of all relevant



locations within each individual aircraft, to allow the CAMO to schedule focused inspections for corrosion at the most appropriate time (when airworthiness is not at risk, the probability of findings is high and repair is still economic).




Figure 32 – General philosophy of CPCA by utilisation of data and AI

**Objective CO-02:** Define the AI-based system taking into account domain-specific definitions of 'system'.

A system at the CAMO would constantly receive operational data from the aircraft, either directly through satellite data link (e.g. ACARS), or indirectly as download by the operator or a contracted service provider. Additional data (e.g. weather data, whether de-icing has been performed, occurrences, repairs) would be constantly acquired as well creating a database covering the full-service history of all individual aircraft under the control of the CAMO.

This does already happen today, but to a lower extent and not specifically focusing on corrosion, but is typically more related to system components (which do provide more specific data easily processed by conventional deterministic algorithms).

A special system would then analyse the data collected, making use of AI and an algorithm trained on similar data of other aircraft in the past to predict the level of corrosion which is probably present at specific areas within individual aircraft.




Figure 33 — Example of a possible system output: predicted corrosion in specific areas

#### 4.1.1.2. Description of the system(s) involved (inputs, outputs, functions)

##### Input:

- Usage data of individual aircraft
- Environmental data (covering the location at the time of operation)
- Operational information (e.g. type of cargo loaded, seafood?)
- Findings from inspections (in all of the fleet)

##### Output:

- Corrosion risk level at individual locations of individual aircraft  
(output could be in the form of an alert or regular status information)

##### Type of AI:

- Pattern detection in large databases

#### 4.1.1.3. Expected benefits and justification for Level 1

The application is expected to improve corrosion control by identifying areas of specific aircraft which have been exposed to increased corrosion risk and require an earlier inspection to limit the severity of structural degradation, or to identify areas of specific aircraft which have not been exposed to high corrosion justifying a later inspection reducing cost, downtime and the risk of access induced damage. This would allow the increase of safety while reducing cost at the same time.



For the maintenance planning activity, it is not so easy to determine the role of humans. Whereas the actual inspection at the aircraft is still performed by humans, the planning of such physical human interference with the aircraft could be implemented at a high level of automation.

Maintenance planning is done today already using computers. Even if performed by humans, all maintenance work at the aircraft is scheduled through computer tools. There is however also always a certain level of human involvement; for example, humans decide which mechanic/inspector should perform which of the scheduled tasks. As such all physical human interference with the aircraft requested by the system can always be overridden by humans (they can always inspect an aircraft although not requested, they can always reject the request to inspect).

In a first application, the system would only support the maintenance planning engineer in deciding when to perform a corrosion inspection at a certain area of an individual aircraft, which would make it a Level 1B system. As the decision to perform a specific maintenance task is always following several considerations (e.g. aircraft availability at the place of the maintenance organisation, availability of hangar space, access requirements and the possibility to perform several tasks at the same time), the final decision is always complex, so the system may also be understood as being only Level 1A and only supporting the maintenance engineer by providing and analysing information.

It could however be possible to upgrade the system up to Level 3A, if all those practical and economical aspects of maintenance planning could be ignored, and the system could automatically schedule inspections without any human interference at CAMO level.

The system could be set up with two types of fundamentally different output:

- Providing the maintenance engineer with regular (e.g. weekly) reports of the aircraft status
- Providing the maintenance engineer with a warning if an area reaches a selected alert threshold

This is similar to the concept of installing either an indication or a warning on the flight deck to either allow monitoring by the flight crew or to alert them when required. There are advantages and disadvantages for both concepts and a combination is also possible.

This will finally make the difference between a Level 1A or 1B system.

**Objective CL-01:** The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.

The *AI Level 1A 'Human augmentation'* classification is justified by only providing additional/advisory information (*support to information analysis*) to the maintenance engineer without any suggestion for action or decision-making.

## 4.2. Damage detection in images (X-Ray, ultrasonic, thermography)

### 4.2.1. Trustworthiness analysis — description of the system and ConOps

Visual inspections and non-destructive testing (NDT) are typical methods to detect damage of aircraft structure.

Those tasks rely on specifically trained inspectors visually detecting damages either by directly inspecting items or by evaluating pictures (e.g. X-Ray pictures). With today's technology, most pictures are no longer produced physically but digitally, detection of damages is already typically performed on computer screens, either 'offline' in offices after taking them at the aircraft or even 'online' directly at the aircraft using portable test equipment with displays.

This use case could be similarly applied to a variety of images, from optical pictures (photographs) taken by humans, fixed cameras or programmed or potentially autonomously acting machines, through sophisticated imaging technology like X-ray or thermography (infrared) up to fully synthetic pictures generated by scanning an area with ultrasonic or eddy current probes. All these inspection methods finally produce digital images which have to be checked for showing damages or defects. Recognising damage shown on digital pictures would be a typical application of AI, similar to some other applications currently widely discussed (runway detection, 'see-and-avoid'). The ML algorithm and the training of the algorithm of course would be individually different for the different types of image to be evaluated.

To be able to address specific issues, the example chosen is the analysis of thermographic images, a method when pictures of the aircraft are taken by digital optical means in the infrared range of the light spectrum in combination with production of a temperature difference (typically heating up the according test item and then inspect it in a room colder than the item), allowing the detection of several types of typical damage in composites structures by visualising the local thermal capacity and conductivity of the item. Infrared cameras have advanced enormously in the last two decades and are now as easy to use as any other optical camera.




Figure 34 — Thermographic images of a fighter aircraft rudder showing water ingress in honeycomb cells




Figure 35 — Optical and thermographic image of a GFRP sandwich panel

#### 4.2.1.1. Description of the system

**Objective CO-02:** Define the AI-based system taking into account domain-specific definitions of 'system'.

A system supporting thermographic inspections of aircraft could be integrated in portable test equipment to be used at the aircraft.

Such a system would not only show, but also analyse the digital image produced with the infrared camera and provide the inspector with additional information and classification of the details seen in the picture by damage type and criticality. A data link to the operator/CAMO/manufacture databases could be envisaged in the future.



Figure 36 — Portable thermographic test equipment, potentially including an image evaluation system




Figure 37 — Example of how the system could mark some areas in images to support inspection of honeycomb sandwich

#### 4.2.1.2. Concept of operations

**Objective CO-03:** The applicant should define and document the concept of operation for all AI-based systems. A focus should be put on the definition of the operational design domain (ODD) and on the capture of specific operational limitations and assumptions.

The terms operation and limitation are not typical in the maintenance domain.

The AI system is intended to be used for NDT to inspect aircraft structures. The system needs to be trained on specific types of structures (e.g. monolithic composites, bonded metal), specific materials (e.g. CFRP, aluminium) and specific failures/damages/defects (e.g. delaminations, disbond, water ingress). Each specific system configuration is strictly limited to be used on the according type of structure.

This is comparable to the situation today with human inspectors, who are also just qualified to perform certain NDT methods on certain types of structure. Training the AI algorithm is comparable to the requirements for human inspectors to be specifically trained for the NDT they perform.

Additionally M.A.608 requires that ‘Tools and equipment shall be controlled and calibrated to an officially recognised standard.’ Specifically for NDT equipment, the individual tools and equipment used have individual sensitivity and detection characteristics. It is therefore normal practice that those are adjusted in line with equipment and aircraft manufacturer instructions in order to be calibrated. To this purpose, defects (type, size) are predefined by the manufacturer by use of a ‘standard’ (i.e. one or more test pieces with an artificial defect as defined by the aircraft manufacturer). This very same philosophy is applicable for ML. The end user needs to train (calibrate) the algorithm (equipment) with a data set (standard) defined by the aircraft manufacturer. Then the end user needs to demonstrate that the trained algorithm is able to correctly classify all the standard samples.

M.A.608 also covers ‘verified equivalents as listed in the maintenance organisation manual’ to ‘the equipment and tools specified in the maintenance data’, meaning it is allowed and normal practice not to use the specific NDT method and/or equipment required by the manufacturer, but an



alternative method/equipment verified to be equivalent. This implicitly allows the use of equipment making use of AI/ML if it is verified to provide equivalent detection capability. This of course needs to be demonstrated to the approving authority.

#### 4.2.1.3. Description of the system(s) involved (inputs, outputs, functions)

Input:

Digital image from an infrared camera

Output:

Digital picture with highlighted areas of interest

Information about the type and severity of damage found

Type of AI:

Image recognition

#### 4.2.1.4. Expected benefits and justification for Level 1

The application is expected to reduce workload and improve the quality of inspection. A major issue of human performance is the change in attention over the day as a lot of maintenance is performed at night either as line maintenance given that the aircraft flies during the day, or in a 24-hour activity to keep the downtime short. The use of AI systems would allow for a more consistent quality of inspections reducing the impact of human factors.

Additionally, the use of an image assessment based on a computer system allows the inspector to be provided with additional information derived from databases, e.g. by recognising which exact location of the aircraft is shown in the picture, to highlight the location of previous repairs or to show modifications and to provide additional information such as the allowable damage size in that area, information which today has to be manually produced by the inspector using the according handbooks.

In a first step, the system would be classified as a Level 1B, as the system would support the inspector to take the decision whether:

- the inspected structure is free of defects;
- it only contains allowable damage; or
- a deeper inspection or a repair is required before the aircraft can return to service.

The final decision and the need to sign off the inspection would remain with the human; the system would just support this.

In a later stage, higher levels would be technically possible but would require a change of the current philosophy about how maintenance is performed, also requiring changes to regulatory requirements.

**Objective CL-01:** The applicant should classify the AI-based (sub)system, based on the levels presented in Table 1 — EASA AI typology and definitions, with adequate justifications.

The *AI Level 1B 'Human assistance'* classification is justified by providing information to *support decision/action selection* to the maintenance engineer.

#### 4.2.1.5. Potential safety impact

A risk of complacency and over-reliance on the applications exists. Inspectors may be biased in their final decision if the system would classify a detail in the image to not show a defect and they may not check as thoroughly as today when being very confident in the performance of the system.

As many inspections are intended to prevent catastrophic failure by detecting damages or defects before they grow to a critical size (damage tolerance concept), non-detection of existing damage can have a safety impact. As long as the final decision is still with the human and the system just provides support, these safety risks exist in combination with human factors, for which safety management systems are already in place.

### 5. Training / FSTD

#### 5.1. Assessment of training performance

This use case will be developed in a future revision of this document.

### 6. Aerodromes

It needs to be made clear that the scope of the European rules for aerodrome safety address the aviation activities and operational processes on the airside only; and that the so-called landside is not covered by these rules. It is however inside the terminal and in relation to passenger services and passenger management where AI has manifold application areas. For example, AI is integrated with airport security systems such as screening, perimeter security and surveillance since these will enable the Aerodrome operator to improve the safety and security of the passengers. Furthermore, border control and police forces use facial recognition and millimetre-wave technologies to scan people walking through a portable security gate. ML techniques are used to automatically analyse data for threats, including explosives and firearms, while ignoring non-dangerous items — for example, keys and belt buckles — users may be carrying. In addition, ML techniques are used by customs to detect prohibited or restricted items in luggage.

On the airside, there are by comparison fewer use cases of AI/ML in the service of aerodrome safety. The most well-known ones are:

#### 6.1. Detection of foreign object debris (FOD) on the runway

The presence of FOD on the runways can end up damaging aircraft, vehicle and equipment, and ultimately can even cause accidents. FOD prevention and the inspection of movement area for the presence of FOD is a core activity of aerodrome operators. Because physical inspections of runways are time-consuming and reduce capacity and are also not free of human detection error, the use of technological solutions for FOD detection has long been attempted. More recently the application of ML by such systems has been included, as this way the detection of FOD and the related alerts would be more reliable. Since there is a considerable market for FOD detection systems and not all systems are of the desired reliability and maturity, it is not advised to single any of them out.

This use case may be further developed in a future revision of this document. EASA would welcome if it could be alerted of any impediments to the evolution of such systems in today's rules for aerodrome safety.

## 6.2. Avian radars

At airports, the prevention of bird strikes to aircraft is an ongoing challenge. Avian radars can track the exact flight paths of both flocks and individual birds up to 10 km. They automatically detect and log hundreds of birds simultaneously, including their size, speed, direction, and flight path. Bird radar tracks may be presented to tablets of the bird control vehicles in real time, thereby creating situational awareness and allowing for better response by bird control staff. Collection of data related to bird activities may be used to predict future problematic areas, identify specific patterns and support decision-making. Since there is a considerable market for avian radar systems and as not all systems are of the desired reliability and maturity, it is not advised to single any of them out.

This use case may be further developed in a future revision of this document. EASA would welcome if it could be alerted of any impediments to the evolution of such systems in today's rules for aerodrome safety.

## 6.3. Drone detection systems

Similar to the situation with birds, the surroundings of aerodromes may be affected by the unlawful operation of commercially available drones. These represent a hazard to aircraft landing and taking off from the runways and can even be operated with malicious intent. Drone detection, tracking and classification, in conjunction with alert and even neutralisation functions by reliable technological solutions will one day provide the desired safety and security for the airport environment; however, as today's technology-based C-UAS solutions are mostly multi-sensor-based, no single technology can perform several functionalities satisfactorily. The improvement of such technologies with ML appears to be the logical evolution.

Since there is a considerable market for such drone detection systems and as not all systems are of the desired reliability and maturity, it is not advised to single any of them out.

This use case may be further developed in a future revision of this document. EASA would welcome if it could be alerted of any impediments to the evolution of such systems in today's rules for aerodrome safety.

## 7. Environmental protection

### 7.1. Engine thrust and flight emissions estimation

This use case will be developed in a future revision of this document.

## 8. Safety management

### 8.1. Quality management of the European Central Repository (ECR)

This use case will be developed in a future revision of this document.

### 8.2. Support to automatic safety report data capture

This use case will be developed in a future revision of this document.

### 8.3. Support to automatic risk classification

This use case will be developed in a future revision of this document.

## G. Annex 3 — Definitions and acronyms

### 1. Definitions

**Accessibility** — The extent to which products, systems, services, environments and facilities can be used by people from a population with the widest range of user needs, characteristics and capabilities to achieve identified goals in identified contexts of use (which includes direct use or use supported by assistive technologies)<sup>25</sup>.

**Accountability** — This term refers to the idea that one is responsible for their action – and as a corollary their consequences – and must be able to explain their aims, motivations, and reasons. Accountability has several dimensions. Accountability is sometimes required by law. For example, the General Data Protection Regulation (GDPR) requires organisations that process personal data to ensure that security measures are in place to prevent data breaches and report if these fail<sup>26</sup>.

**Adaptivity** — The ability to improve performance by learning from experience. [In the machine learning context,] a system is considered **adaptive** when it continues to learn in real-time operation.

**Artificial intelligence (AI)** — Technology that appears to **emulate human performance** typically by learning, coming to its own conclusions, appearing to understand complex content, engaging in natural dialogues with people, enhancing human cognitive performance (also known as cognitive computing) or replacing people on execution of non-routine tasks<sup>27</sup>.

**Artificial neural network (ANN) or neural network (NN)** — A computational graph which consists of connected nodes ('neurons') that define the order in which operations are performed on the input. Neurons are connected by edges which are parameterised by weights (and biases). Neurons are organised in layers, specifically an input layer, several intermediate layers, and an output layer. This document refers to a specific type of neural network that is particularly suited to process image data: convolutional neural networks (CNNs) which use parameterised convolution operations to compute their outputs.

Commonly used types of neural networks are to be highlighted:

- **Convolutional neural networks (CNNs)** — a specific type of deep neural networks that are particularly suited to process image data, based on convolution operators. (Daedalean, 2020)
- **Recurrent neural networks (RNNs)** — A type of Neural Network that involves directed cycles in memory.

**Auditability** — Refers to the ability of an AI system to undergo the assessment of the system's algorithms, data and design processes. This does not necessarily imply that information about business models and intellectual property related to the AI system must always be openly available. Ensuring traceability and logging mechanisms from the early design phase of the AI system can help enable the system's auditability<sup>28</sup>.

<sup>25</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

<sup>26</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

<sup>27</sup> Source: <https://www.gartner.com/en/information-technology/glossary/artificial-intelligence>

<sup>28</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)



**Automation** — The use of control systems and information technologies reducing the need for human input, typically for repetitive tasks.

**Autonomy** — The ability to perform tasks in complex environments without input by a human.

**Bias** — Different definitions of bias have to be considered depending on the context:

- **Bias (in the data)** — The common definition of data bias is that the available data is not representative of the population or phenomenon of study.
- **Bias (in the ML model)** — An error from erroneous assumptions in the learning [process]. High bias can cause an algorithm to miss the relevant relations between attributes and target outputs (= underfitting).

**Big Data** — A recent and fast evolving technology, which allows the analysis of a big amount of data (more than terabytes), with a high velocity (high speed of data processing), from various sources (sensors, images, texts, etc.), and which might be unstructured (not standardised format).

**Completeness** — A data set is complete if it sufficiently covers the entire space of the operational design domain for the intended application.

**Cost function** — A function that measures the performance of an AI/ML model for given data and quantifies the error between predicted values and expected values.

**Critical maintenance task** — A maintenance task that involves the assembly or any disturbance of a system or any part on an aircraft, engine or propeller that, if an error occurred during its performance, could directly endanger the flight safety.

**Data-driven AI** — An approach focusing on building a system that can learn a function based on having trained on a large number of examples.

**Data governance** — A data management concept concerning the capability of an organisation to ensure that high data quality exists throughout the complete life cycle of the data, and data controls are implemented that support business objectives. The key focus areas of data governance include data availability, usability, consistency, integrity, and sharing. It also regards establishing processes to ensure effective data management throughout the enterprise such as accountability for the adverse effects of poor data quality and ensuring that the data which an enterprise has can be used by the entire organisation<sup>29</sup>.

**Data protection impact assessment (DPIA)** — Evaluation of the effects that the processing of personal data might have on individuals to whom the data relates. A DPIA is necessary in all cases in which the technology creates a high risk of violation of the rights and freedoms of individuals. The law requires a DPIA in case of automated processing, including profiling (i), processing of personal data revealing sensitive information like racial or ethnic origin, political opinions, religious or philosophical beliefs (ii), processing of personal data relating to criminal convictions and offences (iii) and systematic monitoring of a publicly accessible area on a large scale (iv)<sup>30</sup>.

**Data Protection Officer (DPO)** — This denotes an expert on data protection law. The function of a DPO is to internally monitor a public or private organisation's compliance with GDPR. Public or private organisations must appoint DPOs in the following circumstances: (i) data processing activities are

<sup>29</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

<sup>30</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

carried out by a public authority or body, except for courts acting in their judicial capacity; (ii) the processing of personal data requires regular and systematic monitoring of individuals on a large scale; (iii) the processing of personal data reveals sensitive information like racial or ethnic origin, political opinions, religious or philosophical beliefs, or refers to criminal convictions and offences. A DPO must be independent of the appointing organisation<sup>31</sup>.

**Data for safety (EASA)** — Data4Safety (also known as D4S) is a data collection and analysis programme that supports the goal of ensuring the highest common level of safety and environmental protection for the European aviation system.

The programme aims at collecting and gathering all data that may support the management of safety risks at European level. This includes safety reports (or occurrences), flight data (i.e. data generated by the aircraft via the flight data recorders), surveillance data (air traffic data), weather data — but those are only a few from a much longer list.

As for the analysis, the programme's ultimate goal is to help to 'know where to look' and to 'see it coming'. In other words, it will support the performance-based environment and set up a more predictive system.

More specifically, the programme will facilitate better knowledge of where the risks are (safety issue identification), determine the nature of these risks (risk assessment) and verify whether the safety actions are delivering the needed level of safety (performance measurement). It aims to develop the capability to discover vulnerabilities in the system across terabytes of data [Source: EASA].

**Deep learning (DL)** — A specific type of machine learning based on the use of large neural networks to learn abstract representations of the input data by composing many layers.

**Determinism** — A system is deterministic if when given identical inputs produces identical outputs.

**Development assurance** — All those planned and systematic actions used to substantiate, to an adequate level of confidence, that errors in requirements, design, and implementation have been identified and corrected such that the system satisfies the applicable certification basis.

**Development error** — A mistake in requirements, design, or implementation.

**Domain** — Operational area in which a system incorporating a ML subsystem could be implemented/used. Examples of domains considered in the scope of this guideline are ATM/ANS, air operations, operator/flight crew training, environmental protection or Aerodromes.

**End user** — An end user is the person that ultimately uses or is intended to ultimately use the AI system. This could either be a consumer or a professional within a public or private organisation. The end user stands in contrast to users who support or maintain the product<sup>32</sup>.

**Failure** — An occurrence which affects the operation of a component, part, or element such that it can no longer function as intended (this includes both loss of function and malfunction). Note: Errors may cause failures, but are not considered to be failures.

---

<sup>31</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

<sup>32</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

**Fairness** — Fairness refers to a variety of ideas known as equity, impartiality, egalitarianism, non-discrimination and justice. Fairness embodies an ideal of equal treatment between individuals or between groups of individuals. This is what is generally referred to as ‘substantive’ fairness. But fairness also encompasses a procedural perspective, that is the ability to seek and obtain relief when individual rights and freedoms are violated.

**Feature** (in computer science) — A feature is any piece of information which is relevant for solving the computational task related to a certain application.

- **Feature** (in machine learning in general) — A feature is an individual measurable property or characteristic of a phenomenon being observed.
- **Feature** (in computer vision) — A feature is a piece of information about the content of an image; typically about whether a certain region of the image has certain properties.

**General Data Protection Regulation (GDPR)** — EU’s data protection law, refer to <https://gdpr.eu> for more details.

**Human agency** — Human agency is the capacity of human beings to make choices and to impose those choices on the world.

**Hyper-parameter** — A parameter that is used to control the algorithm’s behaviour during the learning process (e.g. for deep learning with neural networks, the learning rate, the batch size or the initialisation strategy). Hyper-parameters affect the time and memory cost of running the algorithm, or the quality of the model obtained at the end of the training process. By contrast, other parameters, such as node weights or biases, are the result of the training process<sup>33</sup>.

**Inference** — The process of feeding the machine learning model an input and computing its output. See also related definition of **Training**.

**Input space** — Given a set of training examples of the form  $\{(x_1, y_1) \dots (x_N, y_N)\}$  such that  $x_i$  is the feature vector of the  $i$ -th example and  $y_i$  is its label (i.e. class), a learning algorithm seeks a function  $g : X \rightarrow Y$ , where  $X$  is the input space and  $Y$  is the output space.

**Machine learning (ML)** — The branch of AI concerned with the development of algorithms that allow computers to evolve behaviours based on observing data and making inferences on this data.

*Machine learning* strategies include three methods:

- **Supervised learning** — The process of learning a function that maps an input to an output based on example input-output training samples.
- **Unsupervised learning** — The process of learning a function from a non-labelled data set by adapting the model to increase accuracy of the algorithm based on a given cost function.
- **Reinforcement learning** — The process of learning in which the algorithm rewards positive results and punishes for a negative result, enabling it to improve over time.

**Machine learning model** — A parameterised function that maps inputs to outputs. The parameters are determined during the training process.

- **Trained model** — the machine learning model which is obtained at the end of the learning/training phase.
- **Inference model** — the machine learning model obtained after transformation of the trained model, so that the model is adapted to the target platform.

---

<sup>33</sup> Source: adapted from (Goodfellow-et-al, 2016)

**Predictability** — The degree to which a correct forecast of a system's state can be made quantitatively. Limitations on predictability could be caused by factors such as a lack of information or excessive complexity.

**Redress by design** — Redress by design relates to the idea of establishing, from the design phase, mechanisms to ensure redundancy, alternative systems, alternative procedures, etc. in order to be able to effectively detect, audit, rectify the wrong decisions taken by a perfectly functioning system and, if possible, improve the system<sup>34</sup>.

**Reliability** — The probability that an item will perform a required function under specified conditions, without failure, for a specified period of time<sup>35</sup>.

**Representativeness** (of a data set) — A data set is representative when the distribution of its key characteristics is similar to the actual input state space for the intended application.

**Residual risk** — risk remaining after protective measures have been taken<sup>36</sup>. In the context of this guidance, residual risk designates the amount of risk remaining due to a partial coverage of some objectives. Indeed, it may not be possible in some cases to fully cover the learning assurance building block objectives or the explainability block objectives. In such cases, the applicant should design its AI/ML system to first minimise the residual risk and then mitigate the remaining risk using the SRM concept defined in this guidance.

**Resilience** — In the context of this guidance, the resilience definition is derived from DEEL White Paper on Machine learning in Certified System (DEEL Certification Workgroup, 2021) where resilience is defined as the ability of a system to continue to operate while an error or a fault has occurred.

**Robustness** — for an input varying in a region of the [input] state space, system producing same outputs.

In DEEL White Paper on Machine learning in Certified System (DEEL Certification Workgroup, 2021), robustness is defined as the ability of the system to perform the intended function in the presence of abnormal or unknown inputs, and to provide equivalent response within the neighbourhood of an input.

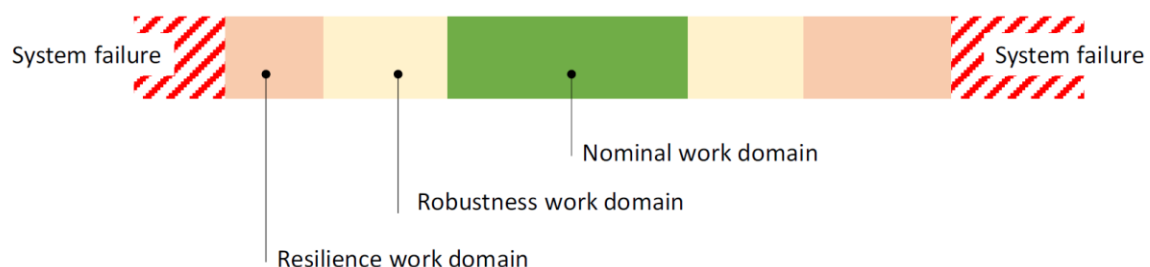


Figure 38 — Illustration of resilience and robustness as defined by DEEL white paper

<sup>34</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

<sup>35</sup> Source: ARP 4761 Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment, 1996

<sup>36</sup> Source: IEC ref 903-01-11 - [https://std.iec.ch/iev/iev.nsf/ID\\_xref/en:903-01-11](https://std.iec.ch/iev/iev.nsf/ID_xref/en:903-01-11)

**Safety criteria** — This term is specific to the ATM/ANS domain and is defined in point ATS.OR.210 of Regulation (EU) 2017/373. This Regulation does not have the notion of safety objective for non-ATS providers; it instead uses the notion of safety criteria. Although the two notions are not fully identical, they are used in an equivalent manner in this document.

**Safety objective** — A qualitative and/or quantitative attribute necessary to achieve the required level of safety for the identified failure condition, depending on its classification.

**Safety requirement** — A requirement that is necessary to achieve either a safety objective or satisfy a constraint established by the safety process.

This term is used in various domains with domain-specific definitions. For the ATM/ANS domain, according to GM1 to AMC2 ATS.OR.205(a)(2), safety requirements are design characteristics/items of the functional system to ensure that the system operates as specified.

**Safety support requirement** — Safety support requirements are characteristics/items of the functional system to ensure that the system operates as specified. This term is used in the ATM/ANS domain for non-ATS providers and is defined in GM1 to AMC2 ATM/ANS.OR.C.005(a)(2).

**Subject** — A subject is a person, or a group of persons affected by the AI system<sup>37</sup>.

**Synthetic data** — Any production data applicable to a given situation that is not obtained by direct measurement.

**System** — A combination of inter-related items arranged to perform a specific function(s) [ED-79A/ARP4754A]

**Safety science** — A broad field that refers to the collective processes, theories, concepts, tools and technologies that support safety management.

**Traceability** — The ability to track the journey of a data input through all stages of sampling, labelling, processing and decision-making<sup>38</sup>.

**Training** — The process of optimising the parameters (weights) of a machine learning model given a data set and a task to achieve on that data set. For example, in supervised learning the training data consists of input (e.g. an image) / output (e.g. a class label) pairs and the machine learning model 'learns' the function that maps the input to the output, by optimising its internal parameters. See also the related definition of **Inference**.

**User** — A user is a person that supports or maintains the product, such as system administrators, database administrators, information technology experts, software professionals and computer technicians<sup>39</sup>.

**Variance** — An error from sensitivity to small fluctuations in the training set. High variance can cause an algorithm to model the random noise in the training data, rather than the intended outputs (=overfitting).

---

<sup>37</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

<sup>38</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

<sup>39</sup> Source: adapted from (EU High-Level Expert Group on AI, 2020)

## 2. Acronyms

AI	artificial intelligence
ALTAI	Assessment List for Trustworthy AI
ALS	airworthiness limitation section
AMC	acceptable means of compliance
AMO	approved maintenance organisation
ANN	artificial neural network
ANS	air navigation services
ANSP	air navigation service provider
ATC	air traffic control service
ATM	air traffic management
ATO	approved training organisation
ATS	air traffic service
CAMO	continuing airworthiness management organisation
CBT	Computer-based training
CMRs	certification maintenance requirements
CNN	convolutional neural network
CNS	communication navigation and surveillance systems
ConOps	concept of operation
CRI	certification review item
CS	certification specification
D4S	Data for safety
DAL	development assurance level
DL	deep learning
DNN	deep neural network
DOA	design organisation approval
DPPIA	data protection impact assessment
DPO	Data Protection Officer
DQRs	data quality requirements
EASA	European Union Aviation Safety Agency
EU	European Union
EUROCAE	European Organisation for Civil Aviation Equipment
FSTD	flight simulation training device
GDPR	General Data Protection Regulation
GM	guidance material
GPU	graphics processing unit
HIC	human-in-command
HITL	human-in-the loop
HLEG	AI High-Level Expert Group
HMI	human-machine interface
HOTL	human-on-the-loop
HOOTL	human-out-of-the-loop
ICA	instructions for continued airworthiness
ICAO	International Civil Aviation Organization
ISM	independent system monitoring
IR	implementing rule
IUEI	intentional unauthorised electronic interaction

JAA	Joint Aviation Authorities
LOAT	level of automation
ML	machine learning
MOA	maintenance organisation approval
MOC	means of compliance
NLP	natural language processing
NN	neuronal network
ODD	operational design domain
PISRA	product information security risk assessment
RPAS	remotely piloted aircraft system
RTCA	Radio Technical Commission for Aeronautics
SMS	safety management system
SRM	safety risk mitigation
SAE	Society of Automotive Engineering
SWAL	software assurance level
VTOL	vertical take-off and landing
WG	working group



## H. Annex 4 — References

**Daedalean. 2020.** *Concepts of Design Assurance for Neural Networks (CoDANN)*. Cologne : EASA, 2020.

**DEEL Certification Workgroup. 2021.** *White Paper - Machine Learning in Certified Systems*. Toulouse : IRT StExupery, 2021.

**ECATA Group. 2019.** *ECATA Technical Report 2019 - The exploitation of Artificial Intelligence in future Aircraft Systems*. 2019.

*Enhancing the reliability of out-of-distribution image detection in neural networks.* **Liang-et-al, Shiyu. 2018.** Vancouver : s.n., 2018. ICLR 2018.

**EU Commission. 2018.** *Communication AI for Europe*. 2018.

**EU High Level Expert Group on AI. 2020.** *Assessment List for Trustworthy AI (ALTAI)*. s.l. : European Commission, 2020.

— **2019.** *Ethics Guidelines for Trustworthy AI*. s.l. : European Commission, 2019.

**EUROCONTROL. 2020.** *ATFCM Users Manual, Edition: 24.0 - Validity Date: 23/06/2020*. s.l. : (accessible at: <https://www.eurocontrol.int/publication/atfcm-users-manual>), 2020.

— **2021.** *Calibration of Optimised Approach Spacing Tool; ED V1.1; Date: 16/04/2021*. s.l. : (accessible at: <https://www.eurocontrol.int/publication/eurocontrol-coast-calibration-optimised-approach-spacing-tool-use-machine-learning>), 2021.

— **2020.** *IFPS Users Manual, Edition: 24.1 - Validity Date: 01/12/2020*. s.l. : (accessible at: <https://www.eurocontrol.int/publication/ifps-users-manual>), 2020.

*Function Allocation Considerations in the Era of Human Autonomy Teaming.* **Emilie M. Roth, Christen Sushereba, Laura G. Militello, Julie Diulio, Katie Ernst. December 2019.** 4 page(s): 199-220, s.l. : Journal of Cognitive Engineering and Decision Making , December 2019, Vol. 12.

**Goodfellow-et-al. 2016.** *Deep Learning*. s.l. : MIT Press, 2016.

**Javier Nuñez et al. 2019.** *Harvis D1.1 State of the Art Review*. s.l. : Clean Sky 2 JU, 2019.

**Kenji Kawaguchi, Leslie Pack Kaelbling, and Yoshua Bengio. 2018.** *Generalization in Deep Learning. Mathematics of Deep Learning*. s.l. : Cambridge University Press, 2018, p. Proposition 5.

**Liu, Qiang & Li, Pan & Zhao, Wentao & Cai, Wei & Yu, Shui. 2018.** *A Survey on Security Threats and Defensive Techniques of Machine Learning: A Data Driven View*. s.l. : IEEE Access. 6. 12103-12117. 10.1109/ACCESS.2018.2805680, 2018.

*On the Uniform Convergence of Relative Frequencies of Events to their Probabilities.* **Chervonenkis, V.N. Vapnik and A.Ya. 1971.** 2, 1971, Theory of Probability and its Applications, Vol. 16, pp. 264-280.

**Parasuraman-et-al, Raja. 2000.** *A Model for Types and Levels of Human Interaction with Automation. IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans, Vol 30, No. 3. May 2000, pp. 286-297.*



*Stronger generalization bounds for deep nets via a compression approach.* **Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang.** 2018. s.l. : Andreas Krause and Jennifer Dy. International Machine Learning Society (IMLS), 2018. 35th International Conference on Machine Learning (ICML). Vol. 35th International Conference on Machine Learning (ICML), pp. pp. 390–418.



## J. Annex 5 — Full list of questions from the ALTAI adapted to aviation

Adaptations from the ALTAI are in *italics*.

### 1. Gear #1 — Human agency and oversight

#### Human agency in aviation applications

- a. **Is the AI-based system designed to interact, guide or take decisions by end users that affect humans or society?**
  - Could the AI-based system generate confusion for ~~some or all~~ end users *and/or* subjects on whether a decision, content, advice or outcome is the result of an algorithmic decision?
  - Are end users *and/or* other subjects adequately made aware that a decision, content, advice or outcome is the result of an algorithmic decision?
- b. **Could the AI-based system generate confusion for ~~some or all~~ end users *and/or* subjects on whether they are interacting with a human or AI system?**
  - Are end users *and/or* subjects informed that they are interacting with an AI-based system?
- c. **Could the AI-based system affect human autonomy by generating over-reliance by end users?**
  - Did you put in place procedures to avoid that end users over-rely on the AI system?
- d. **Could the AI-based system affect human autonomy by interfering with the end user's decision-making process in any other unintended and undesirable way?**
  - Did you put in place any procedure to avoid that the AI system inadvertently affects human autonomy?
- e. **Does the AI-based system simulate social interaction with or between end users *and/or* subjects?**

#### Human oversight in aviation applications

- f. **Please determine whether the AI-based system is overseen by a Human-in-the-Loop, Human-on-the-Loop, Human-in-Command, considering the definitions below.**
  - *From a design perspective, the 'capability for the human to oversee the full design cycle of the AI-based system' is a prerequisite in the aviation domain. It is therefore a strong requirement for any AI application, irrespective of the planned oversight mechanism. This also means that 'self-learning' approaches are not supported by these guidelines.*
  - *From an operational perspective, the EU Commission Guidelines on Trustworthy AI (EU High-Level Expert Group on AI, 2019), introduce definitions for the governance mechanisms Human-in-command (HIC), Human-in-the-loop (HITL) and Human-on-the-loop (HOTL). Those definitions would require refinement for the aviation domain, as these terminologies are already used. As these mechanisms are not further used in these guidelines and as the definitions may vary from one domain of aviation to another, it was not judged necessary to provide a different set of definitions at this stage. Note: Applicants may find necessary to answer this question and characterize the*

*functions/tasks of the AI-based system(s) with such oversight mechanisms. In such a case, the applicant should clarify the definitions used.*

- g. Have the humans (human-in-the-loop, human-on-the-loop, human-in-command) been given specific training on how to exercise *human oversight*?**
- *From an AI-based system design perspective, this question is to be analysed and answered by the applicant, including specific assumptions made on the necessary training provided to the humans performing oversight.*
  - *From an operational perspective, this question pertains to the organisations of the end users of the system. Therefore, coordination with the end users' organisations is required to fully answer this question to ensure adequate level of awareness and competence when interacting with the AI-based system.*
- h. Did you establish any detection and response mechanisms for undesirable adverse effects of the AI system for the operator?**
- *This question is answered through compliance with the objectives of the safety and security assessments, the learning assurance, the explainability and the safety risk mitigation.*
- i. Did you ensure a 'stop button' or procedure to safely abort an operation when needed?**
- *This question is answered through compliance with the objectives of the safety and security assessments and the safety risk mitigation.*
- j. Did you take any specific oversight and control measures to reflect the self-learning or autonomous nature of the AI system?**
- *The 'capability for the human to oversee the full design cycle of the AI-based system' is a strong requirement for any AI application, irrespective of the planned oversight mechanism. Therefore, application for 'self-learning' AI/ML items will not be accepted by EASA at this stage.*
  - *More autonomous systems are considered to be covered under Level 3 AI applications and will be addressed in a future revision of these guidelines.*

## **2. Gear #2 — Technical robustness and safety**

### **Resilience to attack and security in aviation applications**

- a. Could the AI system have adversarial, critical or damaging effects (e.g. to human or societal safety) in case of risks or threats such as design or technical faults, defects, outages, attacks, misuse, inappropriate or malicious use?**
- b. Is the AI system certified for information security (e.g. the certification scheme created by the Cybersecurity Act in Europe) or is it compliant with specific security standards?**
- c. How exposed is the AI system to cyberattacks?**
- Did you assess potential forms of attacks to which the AI system could be vulnerable?

- Did you consider different types of vulnerabilities and potential entry points for attacks such as:
  - data poisoning (i.e. manipulation of training data),
  - model evasion (i.e. classifying the data according to the attacker's will),
  - model inversion (i.e. infer the model parameters).
- d. Did you put measures in place to ensure the integrity, robustness and overall security of the AI system against potential attacks over its life cycle?**
- e. Did you red-team/pentest the system?**
- f. Did you inform end users of the duration of security coverage and updates?**
  - What length is the expected time frame within which you provide security updates for the AI system?

#### **General safety in aviation applications**

- g. Did you define risks, risk metrics and risk levels of the AI system in each specific use case?**
  - Did you put in place a process to continuously measure and assess risks?
  - Did you inform end users and/or subjects of existing or potential risks?
- h. Did you identify the possible threats to the AI system (design faults, technical faults, environmental threats) and the possible consequences?**
  - Did you assess the risk of possible malicious use, misuse or inappropriate use of the AI system?
  - Did you define safety-criticality levels (e.g. related to human integrity) of the possible consequences of faults or misuse of the AI system?
- i. Did you assess the dependency of a critical AI system's decisions on its stable and reliable behaviour?**
  - Did you align the reliability/testing requirements with the appropriate levels of stability and reliability?
- j. Did you plan fault tolerance via, e.g. a duplicated system or another parallel system (AI-based or 'conventional')?**
- k. Did you develop a mechanism to evaluate when the AI system has been changed to merit a new review of its technical robustness and safety?**

#### **Accuracy in aviation applications**

- l. Could a low level of accuracy of the AI system result in critical, adversarial or damaging consequences?**
- m. Did you put in place measures to ensure that the data (including training data) used to develop the AI system is up to date, of high quality, complete and representative of the environment the system will be deployed in?**

- n. Did you put in place a series of steps to monitor, and document the AI system's accuracy?
- o. Did you consider whether the AI system's operation can invalidate the data or assumptions it was trained on, and how this might lead to adversarial effects?
- p. Did you put processes in place to ensure that the level of accuracy of the AI system to be expected by end users and/or subjects is properly communicated?

#### **Reliability, fallback plans and reproducibility in aviation applications**

- q. Could the AI system cause critical, adversarial, or damaging consequences (e.g. pertaining to human safety) in case of low reliability and/or reproducibility?
  - Did you put in place a well-defined process to monitor if the AI system is meeting the intended goals?
  - Did you test whether specific contexts or conditions need to be taken into account to ensure reproducibility?
- r. Did you put in place verification and validation methods and documentation (e.g. logging) to evaluate and ensure different aspects of the AI system's reliability and reproducibility?
  - Did you clearly document and operationalise processes for the testing and verification of the reliability and reproducibility of the AI system?
- s. Did you define tested fail-safe fallback plans to address AI system errors of whatever origin and put governance procedures in place to trigger them?
- t. Did you put in place a proper procedure for handling the cases where the AI system yields results with a low confidence score?
- u. Is your AI system using (online) continual learning?
  - Did you consider potential negative consequences from the AI system learning novel or unusual methods to score well on its objective function?

### **3. Gear #3 — Privacy and data governance**

#### **Privacy in aviation applications**

- a. Did you consider the impact of the AI system on the right to privacy, the right to physical, mental and/or moral integrity and the right to data protection?
- b. Depending on the use case, did you establish mechanisms that allow flagging issues related to privacy concerning the AI system?

#### **Data governance in aviation applications**

- c. Is your AI system being trained, or was it developed, by using or processing personal data (including special categories of personal data)?

- d. **Did you put in place any of the following measures some of which are mandatory under the General Data Protection Regulation (GDPR), or a non-European equivalent?**
- Data protection impact assessment (DPIA);
  - Designate a Data Protection Officer (DPO) and include them at an early state in the development, procurement or use phase of the AI system;
  - Oversight mechanisms for data processing (including limiting access to qualified personnel, mechanisms for logging data access and making modifications);
  - Measures to achieve privacy-by-design and default (e.g. encryption, pseudonymisation, aggregation, anonymisation);
  - Data minimisation, in particular personal data (including special categories of data).
- e. **Did you implement the right to withdraw consent, the right to object and the right to be forgotten into the development of the AI system?**
- f. **Did you consider the privacy and data protection implications of data collected, generated or processed over the course of the AI system's life cycle?**
- g. **Did you consider the privacy and data protection implications of the AI system's non-personal training-data or other processed non-personal data?**
- h. **Did you align the AI system with relevant standards (e.g. ISO25, IEEE26) or widely adopted protocols for (daily) data management and governance?**

#### 4. Gear #4 — Transparency

##### Traceability

- a. **Did you put in place measures that address the traceability of the AI system during its entire life cycle?**
- b. **Did you put in place measures to continuously assess the quality of the input data to the AI system?**
- c. **Can you trace back which data was used by the AI system to make a certain decision(s) or recommendation(s)?**
- d. **Can you trace back which AI model or rules led to the decision(s) or recommendation(s) of the AI system?**
- e. **Did you put in place measures to continuously assess the quality of the output(s) of the AI system?**
- f. **Did you put adequate logging practices in place to record the decision(s) or recommendation(s) of the AI system?**

### Explainability in aviation applications

- g. Did you explain the decision(s) of the AI system to the *end users*?
- h. Do you continuously survey the *end users* if they understand the decision(s) of the AI system?

### Communication in aviation applications

- i. In cases of interactive AI systems, do you communicate to users that they are interacting with an AI system instead of a human?
- j. Did you establish mechanisms to inform users about the purpose, criteria and limitations of the decision(s) generated by the AI system?
  - Did you communicate the benefits of the AI system to users?
  - Did you communicate the technical limitations and potential risks of the AI system to users, such as its level of accuracy and/ or error rates?
  - Did you provide appropriate training material and disclaimers to users on how to adequately use the AI system?

## 5. Gear #5 — Diversity, non-discrimination and fairness

*This gear may not be applicable to all use cases. Therefore, in a first analysis, applicants should check whether the AI-based system could have any impact on diversity, non-discrimination or fairness.*

*If no impact exists, the record of this analysis should be added to the ethical assessment documentation.*

*In case of an impact, please consider the questions from the ALTAI related to Gear #5. The original questions can be seen in the ALTAI (EU High-Level Expert Group on AI, 2020).*

### Avoidance of unfair bias

- a. Did you establish a strategy or a set of procedures to avoid creating or reinforcing unfair bias in the AI system, both regarding the use of input data as well as for the algorithm design?
- b. Did you consider diversity and representativeness of end users and/or subjects in the data?
  - Did you test for specific target groups or problematic use cases?
  - Did you research and use publicly available technical tools, that are state-of-the-art, to improve your understanding of the data, model and performance?
  - Did you assess and put in place processes to test and monitor for potential bias during the entire life cycle of the AI system (e.g. bias due to possible limitations stemming from the composition of the used data sets (lack of diversity, non-representativeness)?
  - Where relevant, did you consider diversity and representativeness of end users and/or subjects in the data?
- c. Did you establish a strategy or a set of procedures to avoid creating or reinforcing unfair bias in the AI system, both regarding the use of input data as well as for the algorithm design?

- d. **Did you ensure a mechanism that allows for the flagging of issues related to bias, discrimination or poor performance of the AI system?**
- *Did you establish clear steps and ways of communicating on how and to whom such issues can be raised?*
  - *Did you identify the subjects that could potentially be (in)directly affected by the AI system, in addition to the (end) users and/or subjects?*
- e. **Is your definition of fairness commonly used and implemented in any phase of the process of setting up the AI system?**
- *Did you consider other definitions of fairness before choosing this one?*
  - *Did you consult with the impacted communities about the correct definition of fairness, i.e. representatives of elderly persons or persons with disabilities?*
  - *Did you ensure a quantitative analysis or metrics to measure and test the applied definition of fairness?*
  - *Did you establish mechanisms to ensure fairness in your AI system?*

#### **Accessibility and universal design**

- f. **Did you ensure that the AI system corresponds to the variety of preferences and abilities in society?**
- g. **Did you assess whether the AI system's user interface is usable by those with special needs or disabilities or those at risk of exclusion?**
- *Did you ensure that information about the AI system is also accessible to users of assistive technologies (such as screen readers)?*
  - *Did you ensure that the user interface of the AI system is also usable by users of assistive technologies (such as screen readers)?*
  - *Did you involve or consult with end users and/or subjects in need for assistive technology during the planning and development phase of the AI system?*
- h. **Did you ensure that universal design principles are taken into account during every step of the planning and development process, if applicable?**
- i. **Did you take the impact of the AI system on the potential end users and/or subjects into account?**
- *Did you assess whether the team involved in building the AI system engaged with the possible target end users and/or subjects?*
  - *Did you assess whether there could be groups who might be disproportionately affected by the outcomes of the AI system?*
  - *Did you assess the risk of the possible unfairness of the system onto the end users' and/or subjects' communities?*

## **6. Gear #6 — Societal and environmental well-being**

### **Environmental well-being**





- a. **Did you identify potential negative impacts of the AI system on the environment or human health throughout its life cycle (development, deployment, use, end of life)?**
  - *Does the AI system require additional energy and/or generates additional emissions?*
  - *Does the AI system have adverse effects on the product's environmental compatibility, in particular on aircraft/engine noise and emissions and emissions arising from the evaporation or discharge of fluids?*
  - *Does the AI system have adverse effects on the product's environmental performance in operation?*
  - *Could the use of the AI system have rebound effects, e.g. lead to an increase in traffic, which in turn could become harmful for the environment or human health?*
- b. **Did you assess the impacts of the AI system on the environment or human health throughout its life cycle?**
- c. **Did you define measures to reduce or mitigate these impacts?**

#### **Work and skills, and impact on society at large or democracy**

*This sub-gear may not be applicable to all use cases. Therefore, in a first analysis, applicants should check whether the AI-based system could have any impact on work and skills.*

*If no impact exists, the record of this analysis should be added to the ethical assessment documentation.*

*In case of an impact, please consider the questions from the ALTAI related to Gear #6 'Work and skills' and 'Impact on society at large or democracy'. Those questions can be found in the ALTAI (EU High-Level Expert Group on AI, 2020).*

- d. **Does the AI system impact human work and work arrangements?**
- e. **Did you pave the way for the introduction of the AI system in your organisation by informing and consulting with impacted workers and their representatives (trade unions, (European) work councils) in advance?**
- f. **Did you adopt measures to ensure that the impacts of the AI system on human work are well-understood?**
  - *Did you ensure that workers understand how the AI system operates, which capabilities it has and which it does not have?*
- g. **Could the AI system create the risk of de-skilling of the workforce?**
  - *Did you take measures to counteract de-skilling risks?*
- h. **Does the system promote or require new (digital) skills?**
  - *Did you provide training opportunities and materials for re- and up-skilling?*
- i. **Could the AI system have a negative impact on society at large or democracy?**

- Did you assess the societal impact of the AI system's use beyond the (end) user and/or subject, such as potentially indirectly affected stakeholders or society at large?
- Did you take action to minimise potential societal harm of the AI system?
- Did you take measures that ensure that the AI system does not negatively impact democracy?

## 7. Gear #7 — Accountability

### Anticipated MOC ET-02-11: Accountability / Auditability

*The AI system should be auditable by internal and external parties, including the approving authorities.*

- a. Did you establish mechanisms that facilitate the AI system's auditability (e.g. traceability of the development process, the sourcing of training data and the logging of the AI system's processes, outcomes, positive and negative impact)?
- b. Did you ensure that the AI system can be audited by independent third parties?

### Anticipated MOC ET-02-12: Accountability / Monitoring

*Questions c., d., e. and f. of the accountability gear may not be applicable to all use cases. Therefore, in a first analysis, applicants should check whether the AI-based system could have any impact on the monitoring of ethical concerns from an organisation's perspective.*

*If no impact exists, the record of this analysis should be added to the ethical assessment documentation.*

*In case of an impact, please consider the questions from the ALTAI related to Gear #7 'Accountability' 3rd, 4th and 5th and 8th bullet points. (Those questions correspond to the items c, d, e and f below and can be found in the ALTAI (EU High-Level Expert Group on AI, 2020).*

- c. Did you foresee any kind of external guidance or third-party auditing processes to oversee ethical concerns and accountability measures?
  - Does the involvement of these third parties go beyond the development phase?
- d. Did you organise risk training and, if so, does this also inform about the potential legal framework applicable to the AI system?
- e. Did you consider establishing an AI ethics review board or a similar mechanism to discuss the overall accountability and ethics practices, including potential unclear grey areas?
- f. For applications that can adversely affect individuals, have redress-by-design mechanisms been put in place? *Note: This is the last item in the ALTAI but has been moved to group it with the three preceding questions).*

### **Anticipated MOC ET-02-13: Accountability / Risk Management**

*The AI system should foresee a risk management process, including interaction with third parties.*

- g. Did you establish a process to discuss and continuously monitor and assess the AI system's adherence to this Assessment List for Trustworthy AI (ALTAI)?**
- Does this process include identification and documentation of conflicts between the six aforementioned *gears* or between different ethical principles and explanation of the 'trade-off' decisions made?
  - Did you provide appropriate training to those involved in such a process and does this also cover the legal framework applicable to the AI system?
- h. Did you establish a process for third parties (e.g. suppliers, end users, subjects, distributors/vendors or workers) to report potential vulnerabilities, risks or bias in the AI system?**
- Does this process foster revision of the risk management process?

Stay informed:  
**[easa.europa.eu/ai](http://easa.europa.eu/ai)**

**European Union  
Aviation Safety Agency**

***Postal address***

Postfach 10 12 53  
50452 Cologne  
Germany

***Visiting address***

Konrad-Adenauer-Ufer 3  
50668 Cologne  
Germany

**Tel.** +49 221 89990-000

**Fax** +49 221 89990-999

**Web** [www.easa.europa.eu](http://www.easa.europa.eu)