EUROCONTROL COAST (Calibration of Optimised Approach Spacing Tool) with use of Machine Learning models

White Paper

Edition: V1.1
Edition date: 16-04-2021
Classification: Green

SUPPORTING EUROPEAN AVIATION

FOUNDING MEMBER SESAR JOINT UNDERTAKING

NETWORK MANAGER

DOCUMENT CONTROL

Document Title	EUROCONTROL COAST (Calibration of Optimised Approach Spacing Tool) with use of Machine Learning models
Document Subtitle	White Paper
Edition Number	V1.1
Edition Validity Date	16-04-2021
Classification	Green
Accessibility	NMD Airport Unit shared drive / 54 TBS-ORD Deployment
Status	Released Issue
Author(s)	Guillaume Stempfel, Ivan De Visscher, Rocío Barragán Montes, Frédéric Rooseleer, Vincent Treve
Contacts	nm.airports@eurocontrol.int

APPROVAL TABLE

The following table identifies all management authorities who have successively approved the present issue of this document.

The approval may also be recorded via electronic workflow, where put in place. Where document approval is made via a meeting or electronic workflow, the details shall be indicated here in place of the approval table.

Authority	Date	Signature
Frédéric Rooseleer Expert, Airport Unit	16/04/2021	
Vincent Treve Deputy Head of Airports Unit	19/04/2021	
Steven Moore Head of Network Operations Management Division	19/04/2021	

EDITION HISTORY

Edition No.	Validity Date	Author(s)	Reason
0.1	16/09/2019	Guillaume Stempfel and Ivan De Visscher (PRISME)	Creation of the document
0.2	18/09/2019	Guillaume Stempfel and Ivan De Visscher (PRISME)	Draft document issue
0.3	18/02/2020	Guillaume Stempfel and Ivan De Visscher (PRISME), Rocío Barragán Montes	Modifications regarding the constant additional buffers, the ML toolbox output and the prediction stage
1.0	30/03/2020	Guillaume Stempfel and Ivan De Visscher (PRISME), Rocío Barragán Montes	Modifications regarding the output file formats, the data validation, the model registry and the general pipeline figures, initial release version
1.1	16/04/2021	Guillaume Stempfel and Ivan De Visscher (PRISME), Rocío Barragán Montes	Update regarding the new conservative models

TABLE OF CONTENT

TABLE OF FIGURES

Executive Summary

Arrival runway capacity/throughput is directly linked to the applicable minimum longitudinal separation between traffic on final approach, as well as the separation delivery performance by the Air Traffic Controllers.

The separation criteria are based on surveillance and on wake turbulence, and typically based on longitudinal distance minima under ATS surveillance service.

In order to safely deliver the separation at the separation delivery point (typically the runway threshold), the Controllers usually take some buffers when spacing the traffic on final. These buffers are taken by experience of the expected traffic speed behaviour (compliance to ATC speed instruction and typical landing speed) and can lead to an average over-separation between 0.5NM and 1NM[1] (above the applicable minima). These buffers also result from the uncertainty related to the compression between traffic on final due to the natural catch-up and the difference in final approach speeds.

In addition, headwind conditions on final approach cause a reduction of the aircraft ground speed which 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 the 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 has a significant impact on the overall network performance and is particularly exacerbated if this occurs on the first rotation of the day because of the impact on all the other rotations throughout the day.

Time-Based Separation (TBS) in final approach is an operational solution, which uses time instead of distance to separate aircraft on their final approach to a runway. In order to apply this concept, approach and tower air traffic controllers need to be supported by a separation delivery tool which

- provides a distance indicator (final target distance – FTD), enabling to visualise surveillance display the distance corresponding to the applicable TBS minima, and taking in account the prevailing wind conditions.
- integrates all applicable separation minima and spacing needs

This separation delivery tool, providing separation indicators between arrival pairs on final approach, also enables an increase in separation performance when providing a second indicator (Intermediate Target Distance – ITD): a spacing indicator to optimise the compression buffers and ensuring optimum runway delivery (ORD).

For being operationally deployed, a separation delivery tool allowing the Time-Based Separation application but also improving ATCO performance and management of complex business rules (separation/spacing), has to be demonstrated as fully reliable.

The move from distance to time-based rules allowing efficient and safe separation management requests to properly model/predict aircraft speed and behaviour in short final and the associated uncertainty. A too conservative definition of buffer can lead to a reduction of efficiency whereas making use of advanced Machine Learning techniques for aircraft behaviour prediction allows improvements of separation delivery compared to today while maintaining or even reducing the associated ATCO workload.

[1] Source: EUROCONTROL Optimum Runway Delivery study, 2017

COAST (Calibration of Optimised Arrival Spacing Tool) developed by EUROCONTROL allows ANSPs to train Machine Learning Models on historical traffic and MET data. Those models, once integrated into a separation delivery tool, allow the computation of the most efficient separation to apply and the associated buffer ensuring a level of safety in line with the TBS Safety Case and the EUROCONTROL TBS-ORD Tool Specification.

With the use of a separation delivery tool and COAST, arrival runway throughput can be increase (up to 4% without effect of TBS, so also under low wind conditions) thanks to the optimisation of the spacing and separation performance.

In addition, TBS operations will deliver mitigation of the headwind effect, and provide full operational resilience and stabilised arrival throughput across headwind conditions.

COAST requires regular updates of the models in order to cover progressively the complete traffic fleet and also to cope with new aircraft and operational evolutions.

1 Introduction

1.1 Time-based Separation concept and separation delivery tool

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 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, the follower speed profile which also depends on the headwind conditions.

In order to allow efficient separation delivery, the support tool also advises the 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 homogenization of the ATCO performances 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 the 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.

COAST (Calibration of Optimised Arrival Spacing Tool) 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 1.

1. 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 Final Target Distance (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. 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 Initial Target Indicator (ITD) is giving the ATCO information about how to separate the aircraft at a pre-agreed speed, typically 160/170kt (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 Delivery Point (DP). It is thus computed using as input the FTD and the leader and follower time-to-fly profiles in the prevailing wind conditions.

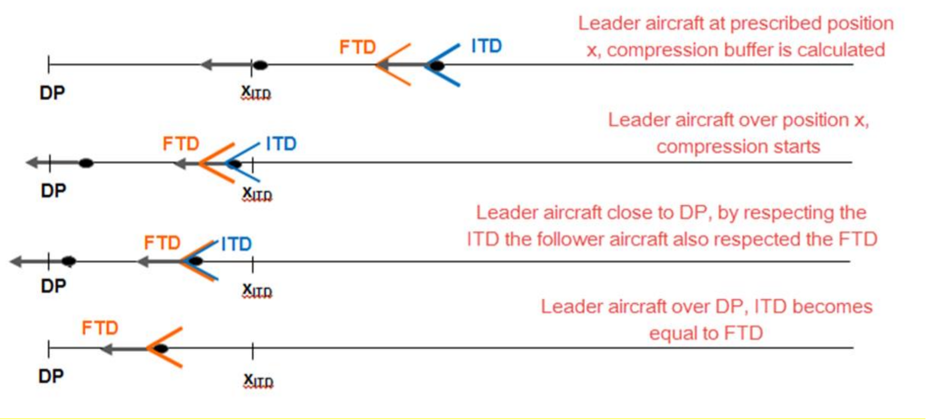

Figure 1: Schematic view of FTD and ITD definition

The benefits related to the use of FTD and ITD are further illustrated in Figure 2 providing three scenarios:

1. Current situation with no optimization: In this scenario there is no FTD (allowing dynamic separation reduction) nor ITD (providing optimizing spacing indication). A conservative spacing buffer is thus applied before leader deceleration in order to cope with compression uncertainty resulting in a separation delivered at threshold showing some margin compared to the minimum.
2. Use of ITD without change of separation minima: In this scenario the use of the ITD allows optimized spacing of the flight before leader deceleration resulting in a separation delivered at threshold with higher accuracy.
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.

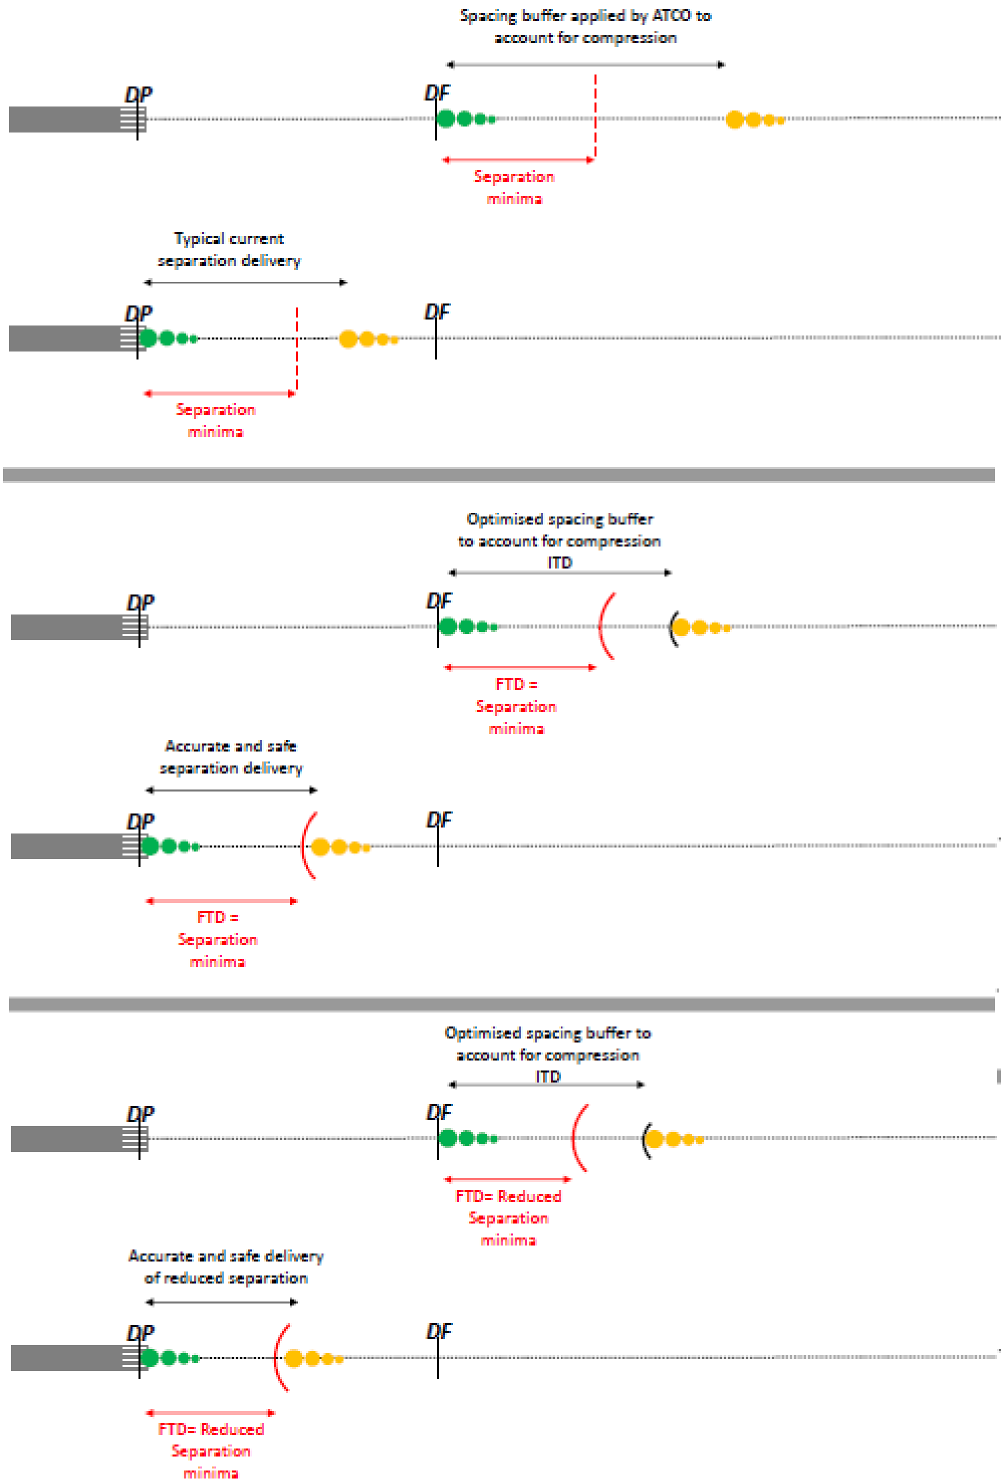

Figure 2: Schematic view of FTD and ITD benefits

However, operationally, uncertainties are found on the time-to-fly profiles of both the leader and follower due to uncertainty on the aircraft airspeed profiles and on the wind conditions. This uncertainty is related to natural variations of aircraft behaviour but also to the wind evolution between the time at which the indicator shall be provided to the controller and the time of aircraft actual landing (for efficient anticipation by the ATCO, this time is typically estimated to 10 minutes). Therefore, buffers are needed both in the FTD and the ITD computation in order to cope with those uncertainties. This is also reflected in the TBS specifications[2]:

TBS-SAF-010 The computed TBS distance representing the time-based wake turbulence separation shall be acceptably safe from a wake risk point of view after considering the effect of uncertainty in the speed profile of the aircraft.

TBS-SAF-020 The computed TBS distance representing the time-based wake turbulence separation shall be acceptably safe from a wake risk point of view after considering the effect of uncertainty in the glide-slope wind conditions.

TBS-SAF-030 If compression indicators are implemented, the computed compression distance shall be acceptably safe after considering the effect of uncertainty in the speed profile of the aircraft.

TBS-SAF-040 If compression indicators are implemented, the computed compression distance shall be acceptably safe after considering the effect of uncertainty in the glide-slope wind conditions.

TBS-MET-020 Uncertainty on the actual glide-slope wind provided to the TBS support tool shall be quantified.

Note that the separation delivery tool can also be used as a first step for optimized distance-based separation application.

1.2 COAST toolchain

COAST produces prediction models and a methodology allowing the computation of FTD and ITD for any pair at a given airport in line with the TBS specification and TBS safety case. It optimizes the model parameterization based on available local data observations. Finally, it also allows maintenance of the models through refinements of the parameterization using additional observation data.

The provided methodology can then be implemented into a separation delivery tool using the provided models as input.

From a technical standpoint, COAST uses local historical flight and meteorological data to model aircraft behaviours in a specific airport environment. The developed models can then be used to compute accurate and safe FTD and ITD indicators for any given situation (i.e., given aircraft pair in a given airport in specific meteorological conditions).

This document will briefly introduce the usage of COAST, then focus on technical aspects, most notably on its inputs and outputs for integration purposes.

[2] EUROCONTROL Specification for Time Based Separation (TBS) for Final Approach

2 Usage of COAST

COAST allows the definition of models and a methodology to use them in order to calculate FTD and ITD for any pair at a given airport.

The use of such Target Distance Indicators allows the ATCOs to:

1. Operate Time-Based Separation concept where separation between aircraft at landing is defined by time instead of distance. Time separations are defined to be equivalent to the distance separations observed in low wind conditions for the same aircraft type.
2. Improve separation conformance and ATCO performance.
3. Deliver a very low under-separation rate compared to today's operations.
4. Account for more complex business rules (e.g., ROT, gap management).

COAST developed by EUROCONTROL allows ANSP to automatically train Machine Learning Models on historical traffic and MET data. Those models, once integrated into a TBS Support Tool, allows the computation of the correct separation indicators (FTD and ITD) to apply (accounting for uncertainties) ensuring a level of safety in line with the TBS Safety Case and the EUROCONTROL TBS-ORD Tool Specification. It also requires regular updates of the models not only to cover progressively the complete traffic fleet, but also to cope with new aircraft and operational evolutions.

COAST is intended to be usable for any airport, provided that at least one-year flights and meteorological data history is available (in order to cover seasonal effects on traffic). TBS targets airports with high runway throughputs, affected by strong wind events which can implicate large delays when a distance-based separation is applied but also to any other airport aiming to improve separation conformance and consequently ATCO performance in delivering consistently optimized runway throughput. This is obtained doing an efficient use of machine learning algorithms allowing the ATCO to better anticipate aircraft behaviour and wind evolution.

COAST is used in two steps, detailed in Figure 3.

1. ANSP:
    - defines the separation minima and TBS criteria,
    - gathers the traffic and meteorological data features available locally (see list in section 5.3.1) and
    - sends the corresponding historical data (at least 1 year or data for the model creation or a few months of data for a model update) to EUROCONTROL.

    With those data, EUROCONTROL uses COAST to train the prediction models.
2. The models developed by COAST are provided to the ANSP together with the description of the methodology on how to use them to compute the FTD and ITD indicators. After local verification, it can be used by the ANSP in a separation delivery tool as an input component to calculate the FTD and ITD indicators.

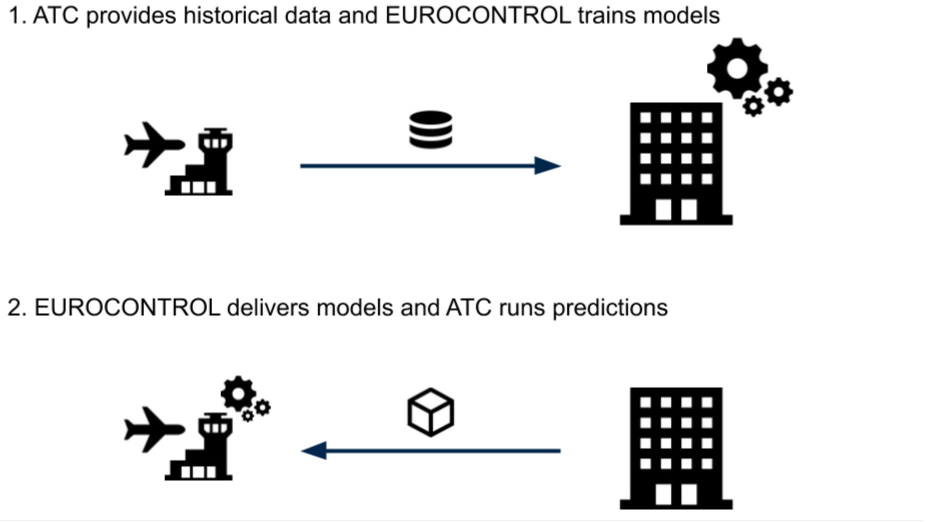

Figure 3: Data and models exchange between ATC and training environment

The models have to be updated on a regular basis in order to check and account for potential evolution in aircraft behaviour, wind condition or aircraft fleet. These updates allow the refinement of the models and the increase of their coverage to cases that were not observed sufficiently frequently (e.g., rare aircraft types and/or rare wind conditions) or not observed at all (e.g., new aircraft type, new airline) in the previously available dataset.

3 Constraints and error rates

A major stake in COAST concept lies in the control of under-separation rates. A set of constraints on predicted separations and their distribution has to be fulfilled in order to consider models output as reliable and acceptable for a regulation authority.

For FTD, the time-separation distribution obtained when applying the FTD computed using COAST must be aligned with the actual time-separation distribution observed in low wind conditions for the same aircraft type, see Figure 4. Then three constraints have been defined to match this target distribution. These constraints are related to three quantiles: TBSp1, TBSp10, TBSp50 (respectively quantiles 1, 10 and 50 of the time separation distribution in low wind conditions). Then, the predicted FTD must respect these constraints with an error rate of at most 1%, 10% and 50%. The quantiles in low wind conditions are obtained based on local experimental data measured in low wind conditions.

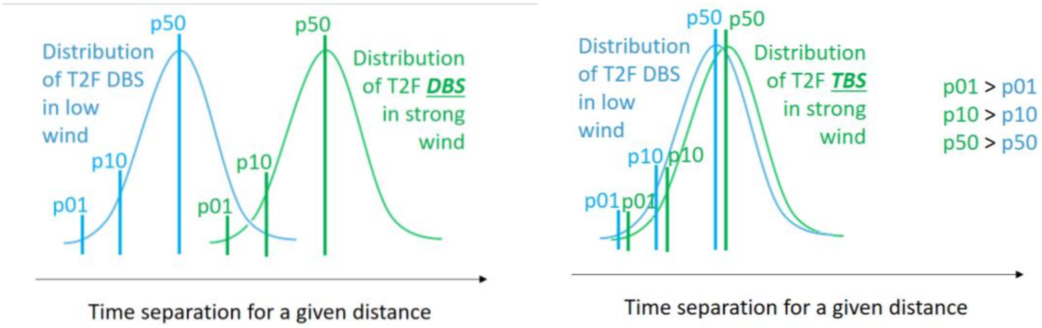

Figure 4: Schematic view of time-to-fly distributions when operating TBS or DBS in low or strong wind condition

Moreover, the ROT constraint shall be respected with a maximal authorized error rate to be defined and agreed locally (typically 2%) or defined using DBS rules as proposed by the ROCAT concept[3].

For ITD definition, we aim at preventing separation compression leading to catch-up of FTD. In particular, if the follower is spaced from the leader by the ITD distance when both aircraft are at a prescribed glide speed (typically 160/170 kts), the follower shall not catch-up the FTD until the leader reaches the runway threshold. Today, ORD Study[4] has shown that 3 to 6% of the delivered separations for constrained pairs are below minima. For the catch-up of the FTD when applying ITD, a maximal error rate of 2.5% is thus lower than what is observed today.

Note that these error rates are defined as if ATC applies exactly the distance computed using COAST methodology and models. Since ATC usually takes some margins with respect to the targeted separation distance either on purpose or just because all aircraft cannot be exactly aligned on the ITD for operational reasons (e.g., vectoring, speed control…), those error rates shall not be taken as absolute numbers of under-separations.

---

[3] EUROCONTROL, "Runway occupancy time data analysis methodology and principles supporting MRS reduction", v1.1, June 2018

[4] Van Baren et al, "The current practice of separation delivery at major European airports", Eleventh USA/Europe Air Traffic Management Research and Development Seminar (ATM2015)

4 General Pipeline of COAST

The pipeline of COAST described in Figure 5 contains two components:

1. The first component is deployed in a EUROCONTROL environment. It is dedicated to the learning process. This component is executed at a regular frequency (typically every 3 to 6 months) or when a threshold on the amount of new data is reached. The update frequency is to be specified locally.

   Initial execution of the COAST pipeline can be summarized as follows: raw flight and meteorological data are first supplied by the ANSP. Data processing operations (denoising, filtering…) 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, a database linked to a file server. The model database also contains its performance in terms of accuracy, coverage, computation details, etc.

   In a maintenance context, when new data are available, the pipeline is similar, except that an automatic data validation stage 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 in the new data compared to the previous ones.

   A more detailed view of the data processing and validation step is depicted in Figure 6. Note that the data validation step is skipped when COAST is carried out for the first time.

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 spec[5]). 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.

[5] EUROCONTROL Specification for Time Based Separation (TBS) for Final Approach

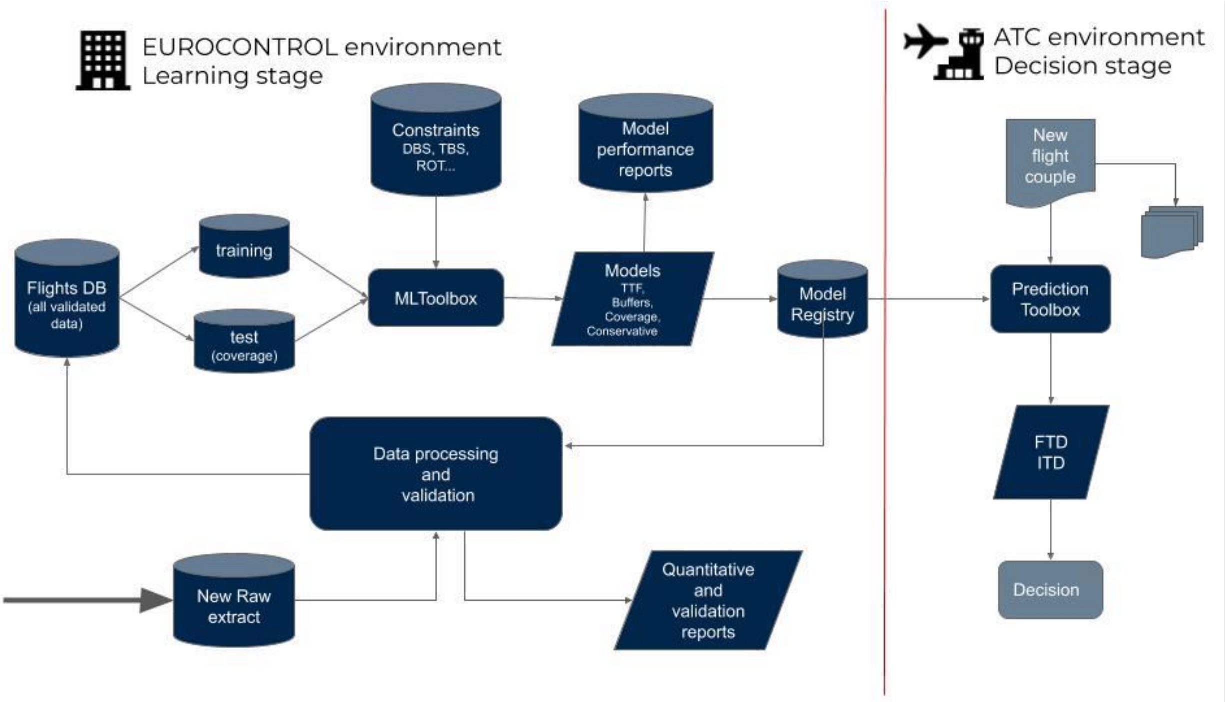

Figure 5: Detailed COAST pipeline

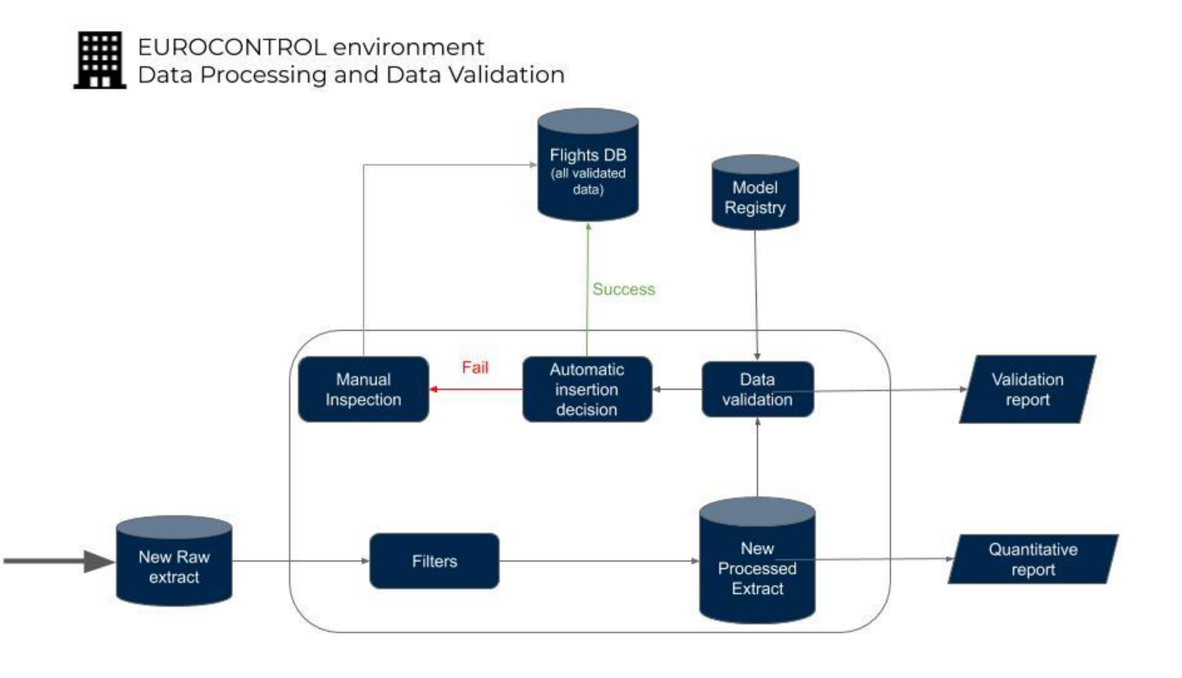

Figure 6: Detailed data processing and data validation pipeline.

5 Technical aspects of COAST - Learning Stage

This section gives a high-level view of COAST learning stage, and of its expected Inputs/Outputs. This part of the pipeline is implemented in EUROCONTROL system.

5.1 Raw input data

COAST needs several sources of data to train reliable models. Data must describe accurately history of flights, time-to-fly profiles and weather. RADAR data, airport data plus meteorological information (such as anemometer, MODE-S and METAR, wind profiler…) are typical inputs of the system.

5.2 Data munging

RAW data must be pre-processed, filtered and denoised before being usable in the learning process. Data reconciliation, removing outliers, data augmentation is performed during this phase.

5.3 ML Toolbox

The ML Toolbox is the Artificial Intelligence core of COAST. From historical data, the toolbox learns and provides the best forecasting models for time-to-fly and safety buffers and a methodology to use them in order to calculate FTD and ITD with a high level of confidence. For flight not (yet) fully covered by the data available (typically less than few hundreds tracks per aircraft type, airline or wind conditions), conservative models are used to increase safety. This conservatism impacts negatively on the efficiency, this is why the models shall be re-trained with more data on a regular basis. This impact is limited since, by definition, “non-covered” flights are a minority.

5.3.1 Inputs

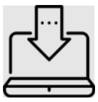

The toolbox requires the following data as inputs:

- **Flight data**:
  - Aircraft type
  - Category (RECAT)
  - Airline
  - Origin airport
  - Runway
  - Landing time
  - Time-to-fly profile
- **Weather data:**
  - Headwind at runway/altitude

  - Crosswind at runway/altitude
  - Temperature at runway/altitude
  - Pressure at runway/altitude
  - ...
- **Constraint information:**
  - DBS
  - ROT
  - TBS quantiles
  - TBS margins

Note that all weather data, except headwind and crosswind at threshold, are not mandatory. Though, it was shown experimentally that having access to reliable altitude wind data allows large improvement in predictions and therefore reduction of buffer to be taken.

5.3.2 ML Toolbox flow

The ML Toolbox is based on:

- a set of machine learning models:
  - Time-to-fly models
  - Buffer models
- A set of coverage functions and the associated decision trees
- A set of conservative models:
  - Conservative models for time-to-fly
  - Conservative models for buffers

5.3.2.1 DATASET SPLIT

The input dataset is split into two subsets at the beginning of the process:

- The training subset is used for learning time-to-fly models, buffer models, conservative models
- The test subset is used to calculate the quality of the models and then to produce coverage functions

5.3.2.2 TIME-TO-FLY MODEL

The first model computed is focused on predicting the time-to-fly profile of a flight. The output of the model is a set of 80 passage times at predefined points.

5.3.2.3 BUFFER MODELS

The time-to-fly model predicts an average expected profile for an aircraft. Its predictions allow the computation of expected FTD and ITD. However, uncertainties exist on the time-to-fly predicted profiles of both the leader and follower and also due to uncertainty on the aircraft airspeed profiles and on the wind conditions. This uncertainty is also related to the time horizon between the indicator computation and their display to the ATCO and the actual time when the follower is landing. This time horizon is typically estimated to 10 minutes.

In order to meet the error rates constraints, buffer models are computed to take into account the time-to-fly model prediction uncertainty. These buffer predictions are added in the calculation of the expected FTD or ITD. One buffer is computed for each constraint: 4 for FTD and 2 for ITD.

For the sake of safety, since the buffer models are designed to be tight with respect to the constraints, a constant additional distance is added at the end of the process. Two constant additional distances are defined: one for the FTD, one for the ITD. This is the final conservative means that needs to be discussed and agreed with regulatory authority. Typically, some hundred meters buffer on top of values calculated by the ML Buffer model are usually sufficient for ensuring an insignificant error rate.

5.3.2.4 COVERAGE FUNCTIONS

Time-to-fly and buffer models are estimated on a training dataset. In order to check the quality of the output, and then to ensure confidence in the models, their performances are evaluated on another independent dataset. This independence is fundamental to avoid introducing bias in the evaluation.

Coverage is defined by feature, for aircraft type, airline, runway, categories, wind band…

A feature is covered if the error rate (obtained on FTD and ITD calculation using the developed models) is guaranteed to meet the design criteria.

A flight is covered if its airline and aircraft type are covered.

A couple is covered if both flights, runway, and the combination between leader categories, follower categories and wind conditions are covered.

5.3.2.5 CONSERVATIVE MODELS

When a flight or a couple is not covered, a fallback is needed. To do so, conservative models are defined for both time-to-fly and buffers. Time-to-fly conservative models are built as the average TAS behaviour per type and surface headwind band on the training set. In the cases where a type is too rare to compute an average, we resort to the worst-case behaviour of the category. Conservative buffers are computed as high quantiles of the buffer distributions observed on the training set. To take in account the multiple cases for conservative ITD (conservative leader and/or conservative follower), a conservative buffer is estimated for each case. All conservative buffers are defined per leader type, follower category and surface headwind band.

These conservative models can also be used for previously unseen cases. This conservative approach guarantees that safety is preserved with a limited operational cost in term of over-spacing. The uncovered pairs are by definition rare enough for not impacting dramatically on the overall performance.

5.4 Outputs

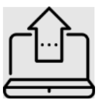

- **Time-to-fly model**
- **Buffer models**

- **Coverage functions**
- **Conservative models**

5.5 Data validation

In order to detect any deviation in the flight normal behaviour, and to prevent any lack of accuracy from the models, a validation process is carried out for each new batch of acquired data.

This validation process monitors two aspects:

- Model appraisal: assessment of the latest model performance using the more recent data (not used to train or validate the model);
- New data validation: detection of inconsistency between old and new data in flight behaviour.

The data validation is performed on acquired data. It is executed before the insertion of new data in the database. It raises a warning if new data are not in line with the rest of the database and related trained models.

In more detail, the data validation is performed on new data extract which has already been pre-processed and filtered. This new data extract must have the same format than the data used to build the original dataset. The deviation is detected through an alignment test between the extract and the latest deployed models. Using the new data, the performance achieved by the latest deployed COAST models and coverage functions is assessed. For this new dataset, if the use of the latest COAST models and coverage functions leads to satisfy the design criteria for FTD and ITD, the data are considered consistent with the rest of the data that were used to train and validate the latest COAST models. The data are then added to the flights database. Otherwise, if the automatic data validation fails (e.g., for certain aircraft or airlines), the new data are considered inconsistent with the previous data, an alert is raised and a manual inspection is performed to understand the cause of the failure. Such cases could for instance occur when:

- A new weather acquisition device is deployed
- An acquisition device is faulty
- An airline changes its approach policy
- ATC modifies its procedure
- ….

The difficulty to detect such drift in the data depends on:

- The number of data affected by the drift. A drift on a very frequent type will be detected with much less history than one on a rare type
- The magnitude of the drift. Few dozens of flights may be sufficient to detect a major event, whereas a light drift might stay undetectable with such an amount of data

The detection of both types of drifts following this methodology guarantees that the safety criteria is met throughout time.

5.6 Final database

The final database contains all previously validated data. This database is used as input of the model training.

6 Packaging of COAST outputs

The models produced by the learning pipeline will be delivered in text-like formats in order to be used by an FTD and ITD calculation tool inside an ATC environment. The exact specifications still have to be defined precisely. Typical foreseen formats are:

- onnx: Open Neural Network Exchange for machine learning models
- json

6.1 COAST outputs supporting FTD prediction calculation

For FTD prediction, COAST provides:

- The trajectory models
  - M/L prediction (encapsulated in onnx format)
  - Conservative trajectory models (provided in json format)
- The FTD buffer models
  - M/L prediction (encapsulated in onnx format)
  - Conservative models (provided in json format)
- FTD computation decision tree logic (pseudo-code detailed in a document) and corresponding coverage functions (in json format)

6.2 COAST outputs supporting ITD prediction calculation

For ITD prediction, COAST provides:

- The trajectory models
  - M/L prediction (encapsulated in onnx format)
  - Conservative leader trajectory models (provided in json or csv format)
  - Conservative follower trajectory models (provided in json format)
- The ITD buffer models
  - M/L prediction (encapsulated in onnx format)
  - Conservative models (provided in json format)
- ITD computation decision tree logic (in pseudo-code detailed in a document) and corresponding coverage functions (in json format)

7 Integration of COAST models into FTD and ITD calculation tool

This section is dedicated to a high-level description of the integration of the models produced by COAST into a calculation tool deployed locally on airport IT.

7.1 Calculation implementation

The calculation must be implemented by the IT provider of the airport. A complete documentation with accurate guidelines will be provided by EUROCONTROL. The calculation makes use of the models produced by COAST allowing locally optimised prediction of safe separation distances between aircraft given a set of inputs about both flights and weather. It also considers the decision tree to determine if, for each arriving flight, the calculation of the indicators should be based on the machine learning predictive models, on the result coming from the conservative models, or on a combination of both.

7.1.1 Inputs

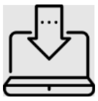

The calculation process requires the following data as input:

- **Flight data**:
  - Aircraft type
  - Category (RECAT)
  - Airline
  - Origin airport
  - Runway
  - Date and Time
- **Weather data:**
  - Headwind at runway/altitude
  - Crosswind at runway/altitude
  - Temperature at runway/altitude
  - Pressure at runway/altitude
  - ...
- **Constraint information:**
  - DBS
  - ROT
  - TBS quantiles
  - TBS margins

- **Models:**
    - Time-to-fly
    - Buffers
    - Conservative
- **Decision tree logic with associated coverage functions**

7.1.2 Calculation flow

The first step is based on a prediction model for the **time-to-fly** of an incoming aircraft.

Then a **buffer model** to predict the distance between a couple of aircraft in the ITD scenario and the distance between a given aircraft and the runway threshold in the FTD scenario.

Recall that all of the above process is used only if we have a sufficient amount of consistent data to train the model and if it leads us to ensure the targeted error rates. If this is not the case, we fallback to the use of the **conservative models**. This approach allows the indicators calculation to be able to cover any aircraft in any situation approaching the airport under safety conditions.

7.1.2.1 FTD

The computation of FTD is shown in Figure 7 while the decision tree to check the most appropriate coverage function is shown in Figure 8.

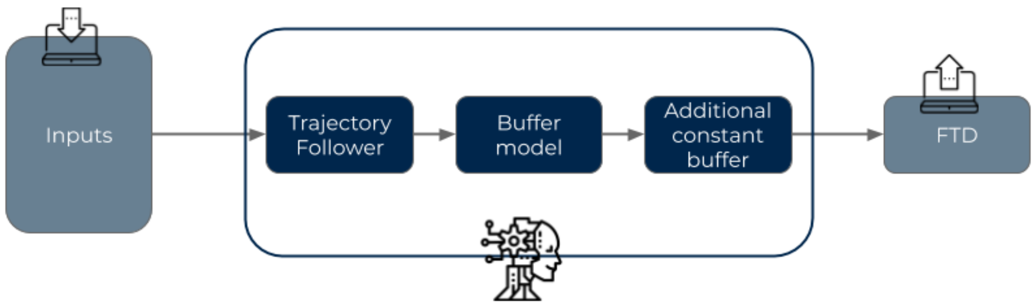

Figure 7: FTD computation

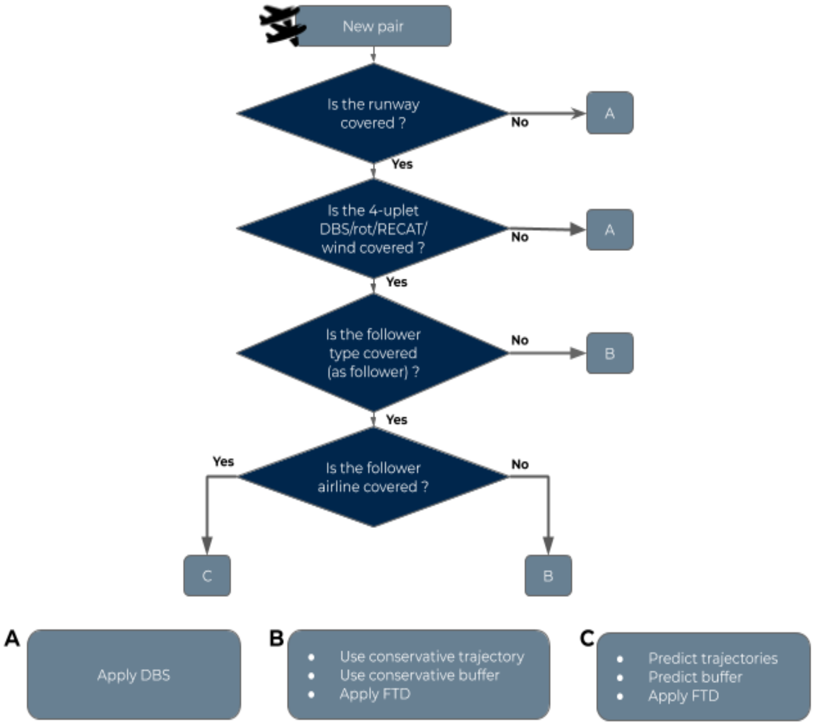

Figure 8: FTD coverage decision tree

7.1.2.2 ITD

The computation of ITD is shown in Figure 9 while the decision tree to check the most appropriate coverage function is shown in Figure 10.

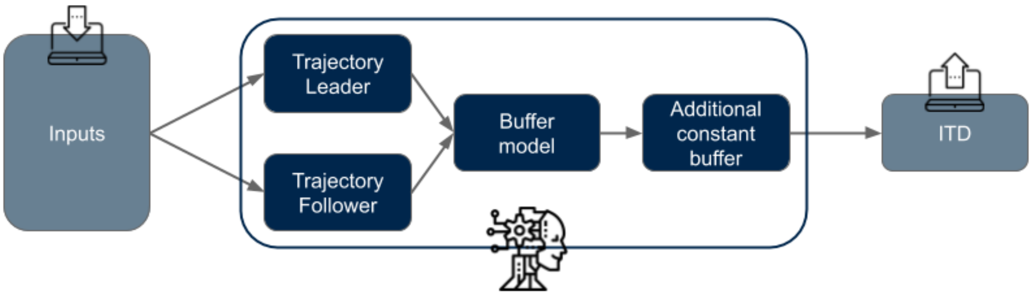

Figure 9: ITD computation

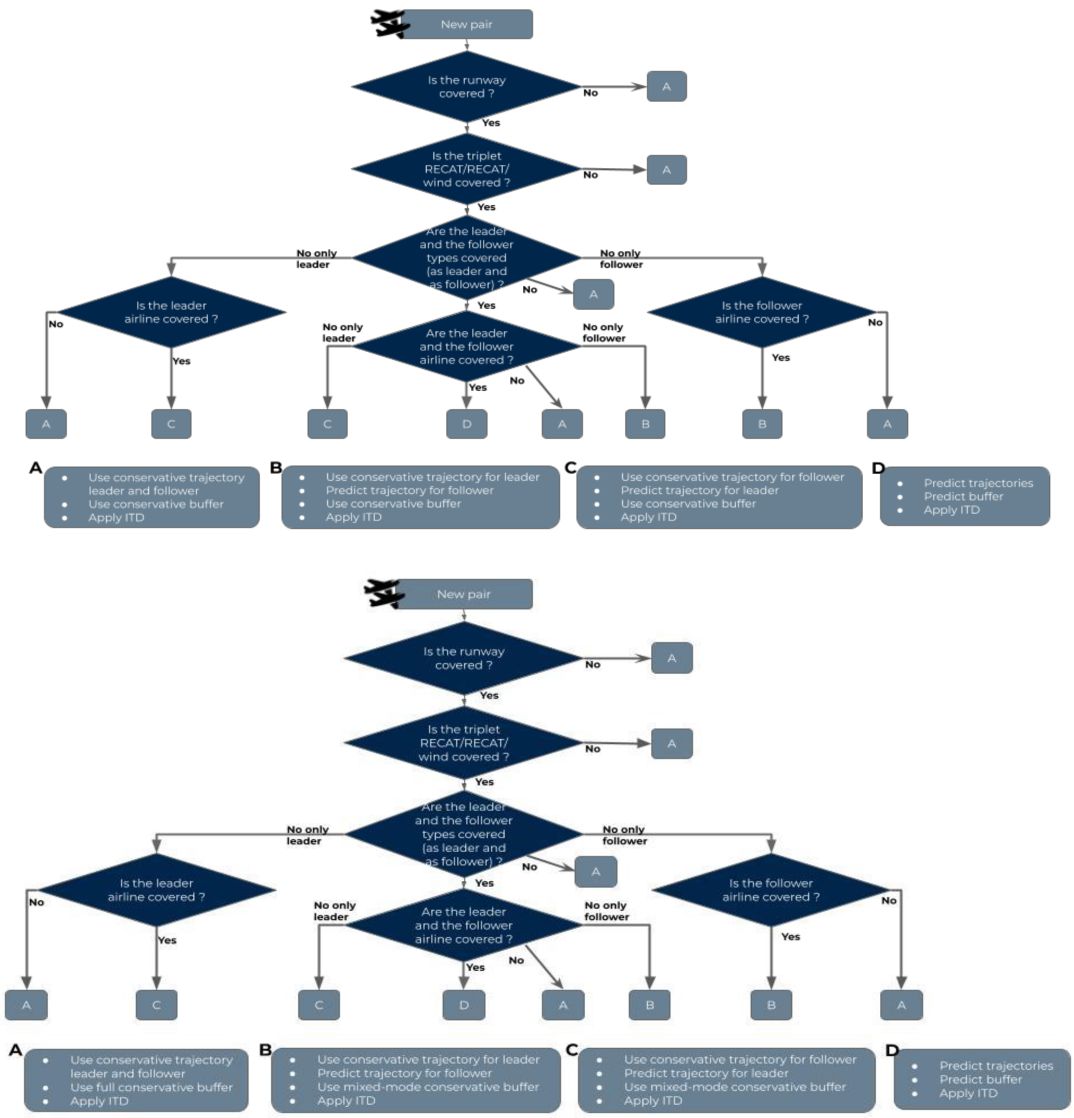

Figure 10: ITD coverage decision tree

7.2 Outputs

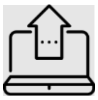

The format of the outputs is dependent on the requirements of the software in which COAST models are implemented. It should contain at least the following information:

- **FTD**
- **ITD**
- **Prevailing compensation constraint used**[6]
- **Are conservative models used?**

[6] When several constraints (e.g., wake minima, MRS, ROT) are applicable to a certain pair, the tool provides the user with the FTD and ITD values of the prevailing one (i.e., that leading to largest distance separation), The tool then also outputs the type of the prevailing constraint, allowing the system to e.g., change the display of the separation indicator (shape, colour) depending on the separation/spacing minimum type.

8 Observed performance on an airport data set

Compared to current situation, the use of a separation delivery tool with the indicators computed using COAST models and methodology has led to

- a safety improvement through reduction of under-spacing,
- a reduction of the separation buffer (by 20% for the pairs covered by the model)
- about 2-4% arrival throughput increase (without effect of TBS, so also in low wind) once a good coverage of the traffic fleet is achieved

It also allows further capacity improvements as it considers TBS operations, which improves throughput in strong head wind conditions.

9 Abbreviations and Definitions

9.1 Abbreviations

Term	Definition
ANSP	Air Navigation Service Provider
COAST	Calibration of Optimised Approach Spacing Tool
DBS	Distance-Based Separation
DP	Delivery Point
FTD	Final Target Distance
ITD	Initial Target Distance
ML	Machine Learning
MRS	Minimum Radar Separation
ORD	Optimum Runway Delivery
ROT	Runway Occupancy Time
TAS	True Air Speed
TBS	Time-Based Separation

9.2 Definitions

Term	Definition
Time-based Separation (TBS)	It consists in improving the landing rate resilience to headwind conditions on final approach by recovering the lost landing rate currently experienced when applying distance-based separation. This is achieved by stabilising the delivered time separation across a range of headwind conditions between aircraft in sequence on final approach to the same runway.
COAST	Calibration tool using Machine Learning techniques built in EUROCONTROL and calibrating models and coverage functions that can be implemented in a separation delivery Tool to be deployed by an ANSP
Final Target Distance (FTD)	The FTD is the distance separation applicable at the delivery point (DP) (e.g., runway threshold) such that all separation and spacing constraints are fulfilled in the prevailing wind conditions. This includes: • the applicable wake time- or distance- based separation minimum (TBS or DBS) • the leader runway occupancy time (ROT) • the minimum surveillance spacing (MRS).

Initial Target Distance (ITD)	The ITD provides an indication of the spacing distance to be applied when the leader is at a prescribed speed (e.g., 160 kts) and before it is decelerating to its final approach speed such that the FTD separation will be observed when the leader will be at the Delivery Point (DP).
ML Toolbox	The ML Toolbox is the core of COAST AI. From historical data, the toolbox learns a set of models supporting the computation of FTD and ITD with large enough confidence.
Error rate	It is the percentage of aircraft for which we do not ensure the targeted constraints.
Separation Delivery Tool	ATC delivery tool that provides guidance to the air traffic controller to also enable time-based separation of aircraft during final approach that considers the effect of the headwind. This is achieved through the provision of indicator (FTD and/or ITD) support for visualising the applicable separation or spacing constraint, where the indicator support is displayed on the extended runway centreline of the Final Approach controller and the Tower Controllers surveillance display.

.

SUPPORTING EUROPEAN AVIATION

www.eurocontrol.int