



(12) **DEMANDE DE BREVET CANADIEN  
CANADIAN PATENT APPLICATION**

(13) **A1**

(86) Date de dépôt PCT/PCT Filing Date: 2020/03/13  
 (87) Date publication PCT/PCT Publication Date: 2020/10/08  
 (85) Entrée phase nationale/National Entry: 2021/07/02  
 (86) N° demande PCT/PCT Application No.: GB 2020/050640  
 (87) N° publication PCT/PCT Publication No.: 2020/201693  
 (30) Priorités/Priorities: 2019/03/29 (EP19275035.4);  
 2019/03/29 (GB1904380.1)


(51) Cl.Int./Int.Cl. *G01S 7/41* (2006.01),  
*G01S 13/78* (2006.01), *G01S 13/91* (2006.01),  
*G08G 5/00* (2006.01)

(71) Demandeur/Applicant:  
 BAE SYSTEMS PLC, GB

(72) Inventeurs/Inventors:  
 CHEHADE, BENJAMIN THOMAS, GB;  
 DEITTERT, MARKUS, GB;  
 JOSEPH, MATTHEW NEIL, GB;  
 METTRICK, SIMON JONATHAN, GB;  
 RIBEIRO, YOHAHN ALEIXO HUBERT, GB;  
 TAYLOR, FREDERIC FRANCIS, GB

(74) Agent: SMART & BIGGAR IP AGENCY CO.

(54) Titre : SYSTEME ET PROCEDE DE CLASSIFICATION DE COMPORTEMENT DE VEHICULE  
 (54) Title: SYSTEM AND METHOD FOR CLASSIFYING VEHICLE BEHAVIOUR



*Fig. 1a*

(57) **Abrégé/Abstract:**

The present invention relates generally to a method and system(10)for classifying vehicle behaviour, particularly abnormal behaviour of civil aircraft(12). The method may comprise receiving aircraft data from an aircraft (12) which is to be classified; and determining whether the received aircraft data comprises identification information for the aircraft (12).In response to a determination that the received aircraft data comprises identification information, the method may comprise using said identification information to classify the behaviour of the aircraft(12). In response to a determination that the received aircraft data does not comprises identification information, the method may comprise obtaining the position of the aircraft and comparing the obtained position to an expected route for the aircraft to classify the behaviour of the aircraft (12).

## (12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property  
Organization  
International Bureau

(43) International Publication Date  
08 October 2020 (08.10.2020)



(10) International Publication Number  
**WO 2020/201693 A1**

## (51) International Patent Classification:

G01S 7/41 (2006.01) G08G 5/00 (2006.01)  
G01S 13/91 (2006.01) G01S 13/78 (2006.01)

## (21) International Application Number:

PCT/GB2020/050640

## (22) International Filing Date:

13 March 2020 (13.03.2020)

## (25) Filing Language:

English

## (26) Publication Language:

English

## (30) Priority Data:

1904380.1 29 March 2019 (29.03.2019) GB  
19275035.4 29 March 2019 (29.03.2019) EP

(71) Applicant: **BAE SYSTEMS PLC** [GB/GB]; 6 Carlton Gardens, London SW1Y 5AD (GB).

(72) Inventors: **CHEHADE, Benjamin Thomas**; BAE Systems, Blue Fin Building (4th floor), 110 Southwark Street, London Greater London SE1 0TA (GB). **DEITERT, Markus**; BAE Systems, Maritime Naval Ships, Golf Course Lane, Bristol South Gloucestershire BS34 7QW (GB). **JOSEPH, Matthew Neil**; BAE Systems, Marabout Industrial Estate, Dorchester Dorset DT1 1YA (GB). **METTRICK, Simon Jonathan**; BAE Systems, Apex Tow-

er, High Street, New Malden Surrey KT3 4LH (GB). **RIBEIRO, Yohahn Aleixo Hubert**; BAE Systems, Apex Tower, High Street, New Malden Surrey KT3 4LH (GB). **TAYLOR, Frederic Francis**; BAE Systems, Building 1, Broad Oak, The Airport, Portsmouth Hampshire PO3 5PQ (GB).

(74) Agent: **BAE SYSTEMS PLC, GROUP IP DEPT**; PO Box 87, Farnborough Aerospace Centre, Farnborough Hampshire GU14 6YU (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,

(54) Title: SYSTEM AND METHOD FOR CLASSIFYING VEHICLE BEHAVIOUR




Fig. 1a

(57) Abstract: The present invention relates generally to a method and system(10)for classifying vehicle behaviour, particularly abnormal behaviour of civil aircraft(12). The method may comprise receiving aircraft data from an aircraft (12) which is to be classified; and determining whether the received aircraft data comprises identification information for the aircraft (12).In response to a determination that the received aircraft data comprises identification information, the method may comprise using said identification information to classify the behaviour of the aircraft(12). In response to a determination that the received aircraft data does not comprises identification information, the method may comprise obtaining the position of the aircraft and comparing the obtained position to an expected route for the aircraft to classify the behaviour of the aircraft (12).

[Continued on next page]



WO 2020/201693 A1

**WO 2020/201693 A1** 

---

TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

**Published:**

- *with international search report (Art. 21(3))*

## SYSTEM AND METHOD FOR CLASSIFYING VEHICLE BEHAVIOUR

The present invention relates generally to a system for classifying vehicle behaviour, particularly abnormal behaviour of civil aircraft.

### BACKGROUND

5 Aircraft typically include a transponder which emits signals carrying identifying information about the aircraft. Two common signals that are used to identify aircraft are an automatic dependent surveillance broadcast (ADS-B) and a response to an identification friend or foe (IFF) interrogation. A problem with these signals is that they can be “spoofed” by other vehicles and/or aircraft so  
10 that the other vehicles can conceal their true identity from surveillance systems. For example, a military aircraft can spoof the signal to appear incorrectly on radars as a civilian aircraft.

It is an example aim of example embodiments of the present invention to at least partially overcome or avoid one or more problems of the prior art, whether  
15 identified herein or elsewhere, or to at least provide an alternative to existing systems and related methodologies.

### SUMMARY OF INVENTION

According to a first aspect of the present invention, there is provided a method for classifying aircraft behaviour, the method comprising: receiving  
20 aircraft data from an aircraft which is to be classified; determining whether the received aircraft data comprises identification information for the aircraft; in response to a determination that the received aircraft data comprises identification information, using said identification information to classify the behaviour of the aircraft; and in response to a determination that the received  
25 aircraft data does not comprises identification information, obtaining the position of the aircraft and comparing the obtained position to an expected route for the aircraft to classify the behaviour of the aircraft.

The different kinds of aircraft data may include some or all of radar data, ship data, radar performance data and identification data, e.g. IFF data. Some or  
30 all of the indicated data may not be available. The identification data may include

the mode S address and/or the flight ID, and may also include a time as well as positional information for the aircraft, e.g. slant range and bearing.

The method may further comprise obtaining the position of the aircraft which is to be classified; comparing the position of the aircraft to known routes to  
5 determine whether the aircraft is within any known route; and in response to determining that the aircraft is not on any known route, classifying the behaviour as suspicious. In this way, spoofing aircraft may be detected.

The method may comprise obtaining the position of the aircraft from the received data or from radar data. The position may include positional information  
10 such as latitude, longitude and height.

The method may further comprise receiving identification information from a plurality of aircraft at the same time as receiving the aircraft data from the aircraft to be classified. Using said identification information to classify the behaviour of the aircraft may then comprise determining whether at least part of  
15 the identification information duplicates identification information received from one of the plurality of aircraft. In response to determining that at least part of the identification information is a duplicate, the method may comprise classifying the behaviour as suspicious. In this way, spoofing aircraft may be detected.

The method may further comprise storing historical data for a plurality of  
20 aircraft. Using said identification information to classify the behaviour of the aircraft may then comprise determining whether at least part of the identification information matches identification information in the historical data. In response to determining that at least part of the identification information matches identification information in the historical data, the method may further comprise  
25 obtaining a route associated with the matching identification information in the historical data; obtaining the position of the aircraft; comparing the obtained position to the associated route to determine whether the aircraft is travelling along the associated route; and classifying the behaviour as suspicious when it is determined that the aircraft is not travelling along the associated route.  
30 Furthermore, when it is determined that the aircraft is travelling along the associated route, the method may comprise determining an expected position along the associated route; comparing the obtained position to the expected

position along the associated route; and classifying the behaviour as suspicious when it is determined that the aircraft is not at the expected position. In these ways, spoofing aircraft may be detected.

The method may further comprise selecting the expected route by  
5 comparing the position data for the aircraft to be classified with a plurality of routes which have been generated by a model which has been trained using historical data received from a plurality of aircraft.

The method may further comprise issuing an IFF request and receiving the received data in response to the IFF request.

10 The identification information may comprise at least one of an aircraft identification which uniquely identifies the aircraft and a flight identification which identifies a flight path along which the aircraft is travelling. The method may then further comprise determining that the identification information comprises both the aircraft identification and the flight identification. Using said identification  
15 information to classify the behaviour of the aircraft may then comprise comparing the aircraft identification and the flight identification to historical data to determine whether they are a matching pair, and classifying the behaviour as suspicious when it is determined that the aircraft identification and the flight identification are not a matching pair. This is another way that spoofing aircraft  
20 may be detected.

The identification information may comprise an identification of the type of aircraft. Using said identification information to classify the behaviour of the aircraft may then comprise receiving radar data about the aircraft, determining a  
25 cross-section value for the aircraft from the radar data, comparing the determined cross-section value with an expected cross-section value for the type of aircraft, and classifying the behaviour as suspicious when it is determined that the determined cross-section value and the expected cross-section value do not match. This is another way that spoofing aircraft may be detected.

The method may comprise determining the cross-section value from the  
30 signal to noise ratio and observed radar cross-section for the radar data. The method may further comprise calculating the expected cross-section value using

a stored value for the type of aircraft and adjusting based on at least a distance to the aircraft and an orientation of the aircraft.

According to another aspect of the invention, there is also provided a system for classifying aircraft behaviour, the system comprising: at least one  
5 sensor for receiving aircraft data from an aircraft which is to be classified; and a processor which is configured to: determine whether the received aircraft data comprises identification information for the aircraft; in response to a determination that the received aircraft data comprises identification information, use said identification information to classify the behaviour of the aircraft; and in response  
10 to a determination that the received aircraft data does not comprises identification information, obtain the position of the aircraft and comparing the obtained position to an expected route for the aircraft to classify the behaviour of the aircraft.

The processor may be further configured to carry out the method steps described above.

15 According to another aspect of the invention, there is a non-transitory computer-readable medium comprising processor control code which when running on a system causes the system to carry out the method described above.

It will be appreciated that any one or more features described in relation to an aspect of the invention may replace, or be used in combination with, any one  
20 or more features described in relation to another aspect of the invention, unless such replacement or combination would be understood by the skilled person as mutually exclusive, after a reading of this disclosure. In particular, any features described in relation to apparatus-like aspects may be used in combination with, or in place of, any features described in relation to method-like aspects. For  
25 example, each of the ways of determining or detecting a spoofing aircraft which are described above may be used alone or in conjunction with one another. They may also be carried out in any appropriate order.

## FIGURES

For a better understanding of the invention, and to show how embodiments of the same may be carried into effect, reference will now be made, by way of example, to the accompanying diagrammatic figures in which:

5           Figure 1a is a schematic illustration of the classifying system;

          Figure 1b is an illustration of a plurality of flight paths which may be derived by the classifying system of Figure 1a;

          Figures 2a to 2d depict a flowchart setting out example steps carried out by the classifying system of Figure 1a; and

10           Figure 3 is a schematic block diagram of components of the classifying system of Figure 1a.

          Figure 1a shows a schematic illustration of the classifying system 10. A vehicle which in this example is an aircraft 12 is periodically sending a signal from a transponder 14 mounted to the underside of the aircraft 12. A common format  
15 for such a periodic signal is an automatic dependent surveillance broadcast (ADS-B) signal, although it will be appreciated that other formats can be used. Such a signal may be detected by the classifying system 10, for example by one or more sensors 16 within the system. The transponder 14 may also send a signal in response to a query from the classifying system 10 and the response may be  
20 detected by the same or a different sensor 16. Such a response signal may be a response to an identification friend or foe (IFF) interrogation.

          The format of these two signals (e.g. the periodic and call-response signals) may be different which means that it may be difficult to compare the signals. For example, an ADS-B signal may comprise the Mode S address, the  
25 flight ID, the position (e.g. latitude and longitude) and time (e.g. day and hour). The IFF response signal typically comprises the Mode S address and time. The IFF response may further comprise the flight ID depending on the IFF interrogation. The position associated with the aircraft sending the IFF response can be derived from the position data in the IFF together with the position data  
30 from the classifying system itself.

The mode S address (or ICAO – International Civil Aviation organisation - address) is a unique identification for the aircraft unless the aircraft is sold to another country. The flight ID (e.g. EZY1234) is an identifier which is set by the crew upon commencing a flight. The flight ID is typically associated with a particular route and thus as explained below, the classifying system can learn the departure and arrival locations associated with the flight ID.

Figure 1b is an illustration of a plot of aircraft tracks in a European airspace. Each track represents a different aircraft. The tracks were plotted based on the ADS-B data which was received in an hour window from multiple aircraft. As shown, the flight paths typically cluster around airports. Furthermore, there is a significant amount of data and thus it is typically difficult to identify an aircraft which is in an unexpected location or which is attempting to spoof another aircraft.

Figures 2a to 2d illustrate an example method of detecting inconsistent aircraft behaviour, for example to identify spoofing of aircraft.

In a first step, data is received from a vehicle by the classifying system (step S100) and this effectively identifies a new track to be analysed by the system. The received data may for example be an IFF response, or an ADS-B signal. The classifying system then checks the data which has been received to determine whether to follow the classification paths shown in Figures 2a, 2b, 2c or 2d. For example, as shown in step S102, there is a check to see whether both the mode S address and the flight ID have been received. If one or both of these pieces of information are missing, the method proceeds to the next step (S202) shown in Figure 2b. However, if both of these pieces of information are present, the data can then be analysed, and the method proceeds as shown in Fig.2a. However, it will be appreciated that the following steps of the analysis can be carried out in any order not just the one described herein. Furthermore, some of the steps could be carried out simultaneously or some steps could be omitted.

In one example, the next step in the analysis is to determine whether or not one or both of the received mode S address and the flight ID have already been received at the classifying system (S104). In other words, a duplicity check is carried out to compare the received data to all current data, for example

received ADS-B data, to determine whether or not there is a match. If one or both of the pieces of data is duplicated, it is possible that one aircraft is spoofing the other matching aircraft. Accordingly, the next step is to issue an alert to a user (step S114). The alert can be in any audio/visual format and may simply  
5 recommend that a user interaction is required to further analyse the received data.

The next example step which may be included in the analysis is to check whether the received mode S address and flight ID match each other. In other words, a data pairing check is carried out to see whether they would be expected  
10 to form a pair (step S106). This matching step could be done by comparison to a database of stored pairings or based on a comparison with historical data to derive a pattern of pairings. If the pieces of data do not form a pair, the next step is to issue an alert to a user (step S114) as described above.

The classifying system may also have received or derived position data for  
15 the aircraft. Accordingly, another example check that could be included is an air lane check to determine whether the aircraft is within a recognised air lane (step S108). The recognised air lanes may be published by aviation authorities or can be learnt by analysing the routes taken by aircraft previously from data recordings. If the aircraft is not within any recognised air lane, or is flying too high  
20 or low, the classifying system 10 issues an alert to a user (step S114) as described above.

The position data may also be used to perform a current position check to check the current position of the aircraft against historical data which has been stored for the aircraft (step S110). The mode S address may be used to identify  
25 the historical data. For example, the air lane in which the aircraft is flying may be identified. There may then be a check to see whether the aircraft has flown in the identified air lane before, and if not the next step is to issue an alert to a user (step S114) as described above.

The position data may be used along side received time data for a check  
30 on the expected position of the aircraft as shown in step S112. In this check, the time and day of the current position of the aircraft is compared to the historic data

to determine whether this flight is likely to be another instance of a regularly scheduled route that the aircraft has flown previously. Accordingly, the combination of time and position data can be used to determine if the aircraft is exceptionally early or late to reach this location. Alternatively, the historical data  
5 can be used to determine where the aircraft was last seen and whether it is feasible for the aircraft to travel to the current location without any intervening reports. If not, the next step is to issue an alert to a user (step S114) as described above.

In another example step, radar data may be used for a cross-section check  
10 (step S113). A cross-section value for the aircraft may be derived from the radar data, which may be data which is included in the radar data or alternatively, the cross-section value may be inferred from the signal to noise ratio in the radar data. The cross-section value may be a classification that the aircraft has a large, medium or small cross-section or may be a numerical value. The cross-section  
15 value determined from the radar data may be compared to a cross-section value within a target database. To make this comparison, it is necessary to know the reported type of aircraft. This aircraft type may be determined using the identification information, e.g. the ADB-S data which was received from the aircraft (at step S102). An adjustment, e.g. for position, orientation and other  
20 factors, to the stored cross-section value may be made to calculate an expected cross-section value based on several parameters. Once the expected value of the cross-section has been calculated, this is compared with the cross-section value derived from the radar data. A confidence value which is indicative of the likelihood that the two values are matched may then be determined. The  
25 confidence value may be compared to a threshold value which may be (pre)determined by training a machine learning algorithm module using historical data and determining which confidence value is sufficient to give an accurate identification. In general it is noted that some types of aircraft (e.g. military) have significantly smaller cross-sections than other types of aircraft (e.g. commercial  
30 aircraft). Accordingly, the matching step does not need to be fine grained and detailed. If there is no match (i.e. the radar cross-section of the aircraft is

significantly smaller/larger than the expected value) an alert can be issued as shown in step S114.

If none of the checks described in the example steps above (S104; S106; S108; S110; S112; S113) results in an alert (S114), the track which has been  
5 analysed can be classified as “neutral”, and the method proceeds to wait for the next data to be received (S116), before returning to step S100 for analysis of the next signal. A neutral classification is indicative that the aircraft has been recognised and is not exhibiting any unusual behaviour. Such a classification may be displayed to a user for example by changing a colour of the track as displayed  
10 to the user on a user interface. Other classifications may include pending (e.g. where the track has not yet been classified), or unknown (e.g. where the track cannot be identified – for example within a given time limit, or owing to a breakdown in the process). Tracks which are pending or unknown may be flagged to a user, e.g. with the alert above. The system may also include a user interface  
15 to allow a user to change the classification and thus override the decision of the system. The method will be then loop back to the beginning (S100 – Fig.2a). It will be appreciated that not all of the checks/steps listed above need to be carried out. Moreover, the checks could be carried out in a different order to that listed above.

20 If the method determines at step S102 that at least one of the mode S address and/or the flight ID data is missing, then the method may proceed to Figure 2b. At step S102, the alternative classification path determines whether only the mode S address is present. If the mode S address is not present on its own at step S202, the method may progress to the process shown in Figure 2c.  
25 However, if only the mode S address is present at step S202, the data may then be analysed as follows. Again, the order of the steps can be changed, steps can be taken at the same time or steps can be omitted.

In a first example step (S204) in the analysis, it is determined whether the received mode S address has already been received by the classifying system.  
30 In other words, a duplicity check is carried out to compare the received data to all current data to determine whether or not there is a match. This step is similar to

step S104 in Figure 2a. If the mode S address is duplicated by another aircraft, the next step is to issue an alert to a user (step S114). The step of issuing an alert is the same as described above in relation to Figure 2a and thus the same reference number is used.

5           As described in relation to Figure 2a, the position data for the aircraft can be used to determine whether the aircraft is within a recognised air lane (step S108) and/or to check the current position of the aircraft against historical data which has been stored for the aircraft (step S110) and/or to determine whether this flight is likely to be another instance of a regularly scheduled route that the  
10 aircraft has flown previously (step S112). Similarly, although it is not depicted, step S113 relating to the radar cross-sectional check can also be carried out if there is sufficient information to identify the type of aircraft. These steps are the same as described above in relation to Figure 2a and thus the same reference numbers are used. If none of the checks described above prompts an alert, the  
15 track which has been analysed can be classified as neutral and the method can simply wait until the next data is received (step S116) before returning to the first step (S100) shown in Figure 2a.

Figure 2c shows an alternative classification path when only the flight ID is present, as determined in step S302. If the flight ID is not present, the method  
20 may progress to the steps shown in Figure 2d. However, if the flight ID is present, the data can then be analysed as follows. Again, the order of the steps can be changed, steps can be taken at the same time or steps can be omitted.

In an example first step in the analysis, it is determined whether the flight ID has already been received at the classifying system (step S304). In other  
25 words, a duplicity check is carried out to compare the received data to all current data to determine whether or not there is a match. This step is similar to steps S104 in Figure 2a and S204 in Figure 2b. If the flight ID is duplicated by another aircraft (within a prescribed time period, or at the same time), the next step is to issue an alert to a user (step S114). The step of issuing an alert is the same as  
30 described above in relation to Figure 2a and thus the same reference number is used.

As described in relation to Figures 2a and 2b, the position data for the aircraft can be used to determine whether the aircraft is within a recognised air lane (step S108) and/or to check the current position of the aircraft against historical data which has been stored for the aircraft (step S110) and/or to  
5 determine whether this flight is likely to be another instance of a regularly scheduled route that the aircraft has flown previously (step S112). Similarly, although it is not depicted, step S113 relating to the radar cross-section check can also be carried out if there is sufficient information to identify the type of aircraft. These steps are the same as described above in relation to Figure 2a  
10 and thus the same reference numbers are used. If none of the checks described above prompts an alert, the track which has been analysed can be classified as neutral and the method can simply wait until the next data is received (step S116). The method will be then loop back to the first step (S100) shown in Figure 2a.

Figure 2d shows an alternative classification path when neither the flight  
15 ID nor the mode S address is present in the received signal, as determined in step S402. It will be appreciated that each of steps S102, S202, S302 and S402 can be carried out simultaneously to select the correct classification path. If neither of the flight ID nor mode S address are present, none of the comparison to historical data which is described above can take place. However, the position  
20 data can be analysed as follows. Again, the order of the steps can be changed, steps can be taken at the same time or steps can be omitted.

As described in relation to Figures 2a to 2c, the position data for the aircraft can be used to determine whether the aircraft is within a recognised air lane (step S108) and if not an alert can be output (step S114). These steps are  
25 the same as described above in relation to Figure 2a and thus the same reference numbers are used.

Unlike the classification paths of Figures 2a to 2c, there is no mode S address or flight ID to identify the historical information for the aircraft. Accordingly, the next step is to select the most probable route(s) (step S404)  
30 which the aircraft may be following based on the time and day and position information for the aircraft. The selection may be based on a learnt model which

has been trained on historic data from aircraft flight patterns. The model may include expected tracks together with common deviations therefrom such as deviations from a great circle path or spread in cruising altitude. At step S410, there is a check to see if the aircraft is on the probable route. There is also a check shown at step S412 in which the current position and time of the aircraft  
5 against historical data for the probable route to determine whether the aircraft is where it would be expected if following this route.

If none of the checks described above prompts an alert (S114), the track which has been analysed can be classified as neutral and the method can simply  
10 wait until the next data is received (step S116). The method will be then loop back to the first step (S100) shown in Figure 2a. Otherwise, an alert is output as described above (step S114).

Figure 3 is a schematic block diagram of the classifying system 10 showing the components which may be used to implement the method shown in Figures  
15 2a to 2d. The classifying system may comprise a civil air classifier module 30 which communicates with a central module 32. The central module 32 also communicates with each of a combat management system (“CMS”) 34, a prediction module (such as a “ForeSee” software module) 36, a condition monitoring module 38, and a simulation tool module 40. It will be appreciated that  
20 the functionality of the various modules could be combined in to fewer modules or separated in to more modules. Similarly, the depicted flow of information and requests for information is merely indicative.

The civil air classifier module 30 may be activated by an API request from the central module 32. As explained in detail above, the civil air classifier module  
25 30 may implement a series of checks, including one or more of a duplicity check, a data pairing check, an air lane check and a current position check or other tests such as testing whether the aircraft is on a probable route. If any of these checks or tests fails, an alert is issued and sent to the central module 32. In addition to issuing an alert, as explained in more detail below, the civil air classifier module  
30 30 may be used to identify the aircraft and this identification is also sent to the central module 32.

The civil air classifier module 30 requires information to be able to identify the aircraft and/or issue an alert. For example, the civil air classifier module 30 may request system track and radar data. In this example, the request is routed through the central module 32. The requested radar data may include positional data in a Cartesian format; i.e. in the form of X, Y, and Z positions and/or radar cross-section of the aircraft. The requested radar data may comprise predicted radar performance data which may be requested from a prediction module 36 (such as a ForeSee module). The requested radar data may also or alternatively comprise actual radar performance data which, in this example, is requested from the condition monitoring module 38. The radar performance data may indicate any maximum range for the radar in terms of its mode and the conditions. The classifier module 30 may itself be located on a moving vessel, e.g. a ship, and there may also be data, e.g. positional data such as latitude, longitude and heading, to capture the information about the vessel. The system track information may be requested from the combat management system (CMS) 34.

The civil air classifier module 30 may send the track of interest to a cross-section classifier module 33, via the central module 32, to compare a cross-section derived from the radar data with a cross-section within a database. The cross-section classifier module 33 requests a stored cross-section value for the same type of aircraft from a target database 43. The target database 43 may be a lookup database comprising cross-sectional values stored against general names for types of aircraft, e.g. jumbo jet, small passenger jet, private jet, military aircraft, or against the specific names for types of aircraft, e.g. Boeing 747. The cross-section classifier module 33 may adjust the received stored cross-section value to calculate an expected cross-section value based on several parameters, e.g. the direction of the illuminating radar, the radar transmitters frequency, the material of the aircraft, the geometry and exterior features of the aircraft. The cross-section classifier may comprise a machine learning algorithm module which would be trained using historical data to build a model to adjust the stored cross-section value to match the physical conditions for the real world measurements. Once the expected value of the cross-section has been calculated, this is compared with the cross-section value from the radar data. A

confidence value which is indicative of the likelihood that the two values are matched may then be determined. If there is no match, an alert can be issued.

The civil air classifier module 30 may also be used to create synthetic tracks in conjunction with a simulation tool 40. In this arrangement, the simulation  
5 tool 40 activates the creation of the synthetic tracks by issuing an API request to the central module 32, which is routed to the civil air classifier module 30. The synthetic tracks may be created using historic data which has been received by the civil air classifier module 30. The simulation tool may be used to visualise the  
10 created synthetic tracks which are representative of the time and region under consideration by a user. For example, the triggering API request may be for track(s) at a particular point in time or over a period in time. The request may specify a single track or a plurality of tracks.

The central module 32 brings all the data from the various modules together to create fused data which is sent to the combat management system  
15 (CMS) 34. The alerts are also forwarded from the central module 32 to the combat management system (CMS) 34 (via ALMAS - Alert Management System).

The central module 32 may also comprise an internal performance monitoring module which can be used to log or record every time a classification of a track is made (and also when an alert is issued or a track is classified). A  
20 threshold for the correct identification of aircraft tracks can be set at any suitable value between 0 to 100% and a typical confidence threshold may be 90% or 95%. The internal performance monitoring module could also be used to check whether the input or output data is corrupt and warn a user accordingly. The internal performance monitoring module could also be used to check that if the system is  
25 shown the same or similar data, the output from the system is the same or similar, i.e. the system assigns the same identity to the same aircraft if the data is repeated. The internal performance monitoring module could also be used to check how quickly the system is identifying tracks. For example, the acceptable threshold for identifying a new track may be set at a few seconds from initial  
30 detection. In other words, the speed of identification may be at least as fast as the new tracks are received, subject to there being sufficient computing capability.

If a track is not assigned within this system, an alert may be issued for that track or the internal performance monitoring module may flag that the system is not functioning properly.

At least some of the example embodiments described herein may be  
5 constructed, partially or wholly, using dedicated special-purpose hardware. Terms such as 'component', 'module' or 'unit' used herein may include, but are not limited to, a hardware device, such as circuitry in the form of discrete or integrated components, a Field Programmable Gate Array (FPGA) or Application Specific Integrated Circuit (ASIC), which performs certain tasks or provides the  
10 associated functionality. In some embodiments, the described elements may be configured to reside on a tangible, persistent, addressable storage medium and may be configured to execute on one or more processors. These functional elements may in some embodiments include, by way of example, components, such as software components, object-oriented software components, class  
15 components and task components, processes, functions, attributes, procedures, subroutines, segments of program code, drivers, firmware, microcode, circuitry, data, databases, data structures, tables, arrays, and variables. Although the example embodiments have been described with reference to the components, modules and units discussed herein, such functional elements may be combined  
20 into fewer elements or separated into additional elements.

Although a few preferred embodiments have been shown and described, it will be appreciated by those skilled in the art that various changes and modifications might be made without departing from the scope of the invention, as defined in the appended claims.

25 Attention is directed to all papers and documents which are filed concurrently with or previous to this specification in connection with this application and which are open to public inspection with this specification, and the contents of all such papers and documents are incorporated herein by reference.

30 All of the features disclosed in this specification (including any accompanying claims, abstract and drawings), and/or all of the steps of any

method or process so disclosed, may be combined in any combination, except combinations where at least some of such features and/or steps are mutually exclusive.

Each feature disclosed in this specification (including any accompanying  
5 claims, abstract and drawings) may be replaced by alternative features serving the same, equivalent or similar purpose, unless expressly stated otherwise. Thus, unless expressly stated otherwise, each feature disclosed is one example only of a generic series of equivalent or similar features.

The invention is not restricted to the details of the foregoing  
10 embodiment(s). The invention extends to any novel one, or any novel combination, of the features disclosed in this specification (including any accompanying claims, abstract and drawings), or to any novel one, or any novel combination, of the steps of any method or process so disclosed.

- 17 -

## CLAIMS

1. A method for classifying aircraft behaviour, the method comprising:  
receiving aircraft data from an aircraft which is to be classified;  
5 determining whether the received aircraft data comprises identification  
information for the aircraft;  
in response to a determination that the received aircraft data comprises  
identification information, using said identification information to classify the  
behaviour of the aircraft; and  
10 in response to a determination that the received aircraft data does not  
comprises identification information, obtaining the position of the aircraft and  
comparing the obtained position to an expected route for the aircraft to classify  
the behaviour of the aircraft.
- 15 2. The method of claim 1, further comprising:  
obtaining the position of the aircraft which is to be classified;  
comparing the position of the aircraft to known routes to determine  
whether the aircraft is within any known route; and  
in response to determining that the aircraft is not on any known route,  
20 classifying the behaviour as suspicious.
3. The method of any one of claims 1 to 2, further comprising:  
receiving identification information from a plurality of aircraft at the same  
time as receiving the aircraft data from the aircraft to be classified; and  
25 wherein using said identification information to classify the behaviour of  
the aircraft comprises:  
determining whether at least part of the identification information  
duplicates identification information received from one of the plurality of aircraft.
- 30
4. The method of any one of claims 1 to 3, further comprising:  
storing historical data for a plurality of aircraft; and

wherein using said identification information to classify the behaviour of the aircraft comprises determining whether at least part of the identification information matches identification information in the historical data.

- 5 5. The method of claim 4, further comprising, in response to determining that at least part of the identification information matches identification information in the historical data:

obtaining a route associated with the matching identification information in the historical data;

- 10 obtaining the position of the aircraft;

comparing the obtained position to the associated route to determine whether the aircraft is travelling along the associated route; and

classifying the behaviour as suspicious when it is determined that the aircraft is not travelling along the associated route.

15

6. The method of claim 5, further comprising, when it is determined that the aircraft is travelling along the associated route:

determining an expected position along the associated route;

- 20 comparing the obtained position to the expected position along the associated route; and

classifying the behaviour as suspicious when it is determined that the aircraft is not at the expected position.

7. The method of any one of claims 1 to 6, comprising selecting the  
25 expected route by comparing the position data for the aircraft to be classified with a plurality of routes which have been generated by a model which has been trained using historical data received from a plurality of aircraft.

- 30 8. The method of any one of the preceding claims, wherein the identification information comprises at least one of an aircraft identification which uniquely

identifies the aircraft and a flight identification which identifies a flight path along which the aircraft is travelling.

9. The method of claim 8, further comprising:

5 determining that the identification information comprises both the aircraft identification and the flight identification and

using said identification information to classify the behaviour of the aircraft comprises:

10 comparing the aircraft identification and the flight identification to historical data to determining whether they are a matching pair, and

classifying the behaviour as suspicious when it is determined that the aircraft identification and the flight identification are not a matching pair.

10. The method of any one of the preceding claims, wherein the identification  
15 information comprises an identification of the type of aircraft.

11. The method of claim 10, wherein using said identification information to classify the behaviour of the aircraft comprises:

20 receiving radar data about the aircraft,  
determining a cross-section value for the aircraft from the radar data,  
comparing the determined cross-section value with an expected cross-section value for the type of aircraft, and

25 classifying the behaviour as suspicious when it is determined that the determined cross-section value and the expected cross-section value do not match.


12. The method of claim 11, further comprising determining the cross-section value from the signal to noise ratio and observed radar cross-section for the radar data.

30


13. The method of claim 11 or claim 12, further comprising calculating the expected cross-section value using a stored value for the type of aircraft and

adjusting based on at least a distance to the aircraft and an orientation of the aircraft.

14. A system for classifying aircraft behaviour, the system comprising:
- 5 at least one sensor for receiving aircraft data from an aircraft which is to be classified; and
- a processor which is configured to:
- determine whether the received aircraft data comprises identification information for the aircraft;
- 10 in response to a determination that the received aircraft data comprises identification information, use said identification information to classify the behaviour of the aircraft; and
- in response to a determination that the received aircraft data does not comprises identification information, obtain the position of the aircraft and
- 15 comparing the obtained position to an expected route for the aircraft to classify the behaviour of the aircraft.
15. A computer-readable medium comprising processor control code which when running on a system causes the system to carry out the method of any
- 20 one of claims 1 to 13.



*Fig. 1a*



*Fig. 1b*




Fig. 2a




Fig. 2b

4/6




Fig. 2c


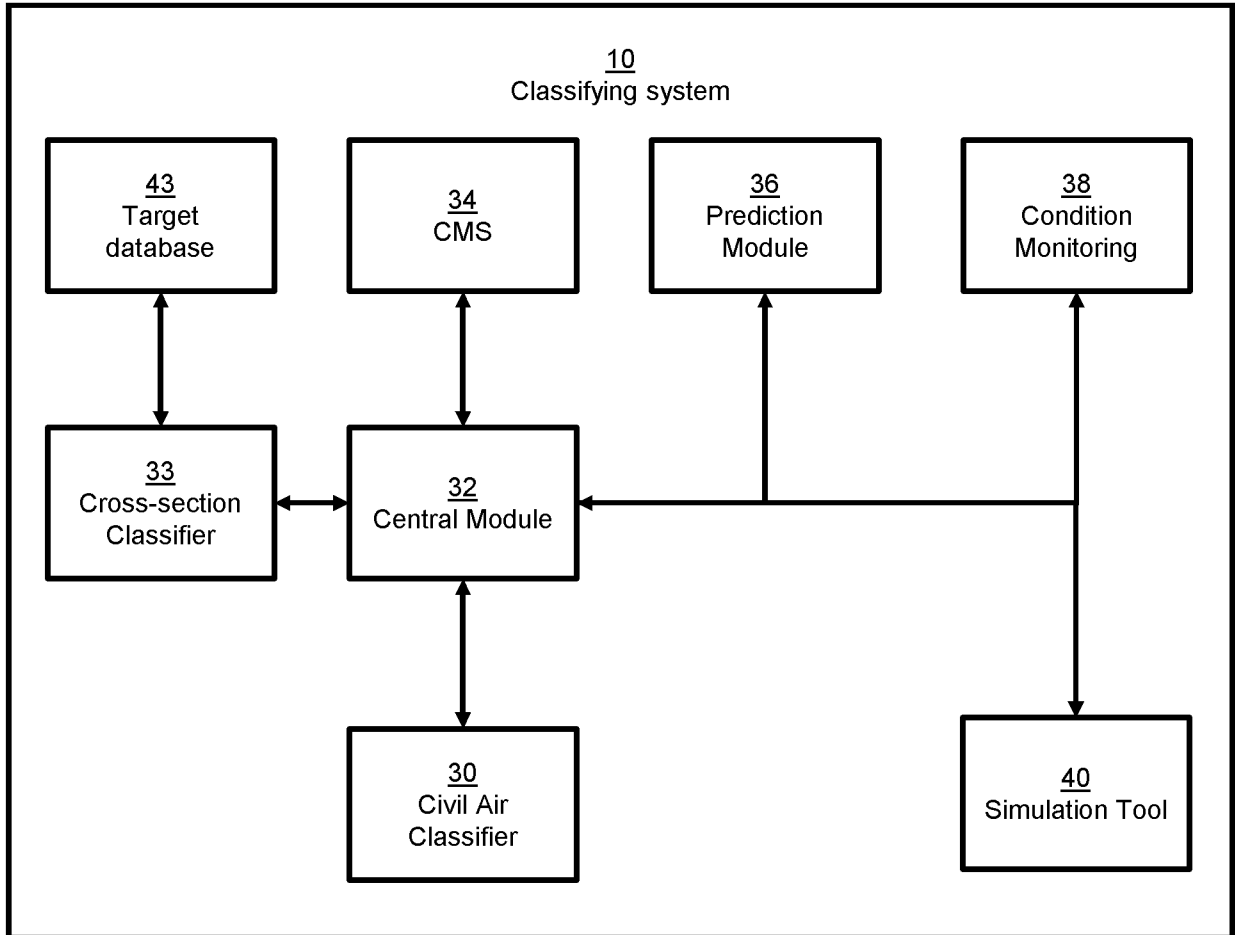



Fig. 2d



*Fig. 3*



*Fig. 1a*