Flight Anomaly Tracking for Improved Situational Awareness:

Case Study of Germanwings Flight 9525

Murat Yasar[1]

[1] *United Technologies Research Center, East Hartford, CT, 06118, USA*
yasarm@utrc.utc.com

ABSTRACT

Surveillance technologies play an important role in civil aviation safety by providing situational awareness to air traffic controllers and pilots. Tens of thousands of aircrafts operate daily, and each one of them needs to be tracked by combining data from several data sources including ADS-B and radar data for maintaining the safety of airspace. These heterogeneous data sources are aggregated together with schedule and flight status data from airlines and airports. This aggregate data is used to provide appropriate flight tracking of individual aircrafts and helps ensuring that the air traffic operates with maximum safety and minimum delays. This is achieved by a complex system of command centers, control towers, radar ground stations, and automated surveillance equipment. As air travel grows each year, global aviation safety continues to improve thanks to these sophisticated systems. Yet, it is unrealistic to expect that the system would detect, identify and respond to all flight anomalies. As was the case in Germanwings flight 9525, the flight anomalies that are not detected in time may result in catastrophes. This paper analyzes the unfortunate case of Germanwings flight 9525 and proposes an automated flight anomaly detection technology to improve situational awareness for air traffic controllers and pilots, and enhance aviation safety.

1. INTRODUCTION

Commercial flights fly on airways under the direction of air traffic control. For safe air travel, flights are assigned a route that consists of various way points between airports. Air traffic controllers help pilots keep aircrafts safely separated from other aircrafts or obstacles while in flight or on the ground, ensuring safe, orderly and efficient traffic flow. When an aircraft deviates from the planned route, the results are often catastrophic. Recent examples are Malaysia Airlines flight 370 and Germanwings flight 9525. Although a summary flight plan is provided to air traffic control before the flight, it is prohibitive for air traffic control to respond to this type of flight anomalies in a timely manner due to the massive volume of flights that needs to be tracked in real-time.

The current practice for flight tracking relies heavily on monitoring the air traffic by ground based Secondary Surveillance Radars (SSRs) and aircraft based radar transponders. This technology is based on sweeping the airspace by a narrow band antenna that transmits an interrogation signal. The transponders located on the aircraft responds with a coded message that includes the aircraft's identification, altitude and other coded information. The ground station determines the aircraft's direction and distance using the radar data and transmits to the air traffic controller. Although its coverage and capacity is limited by line-of-site, the radar-based tracking is predominantly used across the world.

Flight plans are filed by a pilot or flight dispatcher with the local civil aviation authority (e.g. FAA in the USA) prior to departure to indicate the plane's planned route. The air traffic controllers follow the motion of multitude of aircrafts on their screen to ensure that all aircrafts are keeping to their assigned routes, altitudes and speeds. This creates a high load in the decision making process.

Recently, a new flight tracking technology, called Automatic dependent surveillance – broadcast (ADS–B), is being implemented on board the aircrafts. ADS-B is a cooperative surveillance technology in which an aircraft determines its position via satellite navigation and periodically broadcasts it, enabling it to be tracked. Unlike the radar-based surveillance technique, an ADS-B equipped aircraft actively broadcasts its positions once per second. The information received by air traffic controllers, and other ADS-B equipped aircraft includes the aircraft's identification, altitude, speed, velocity, projected path and

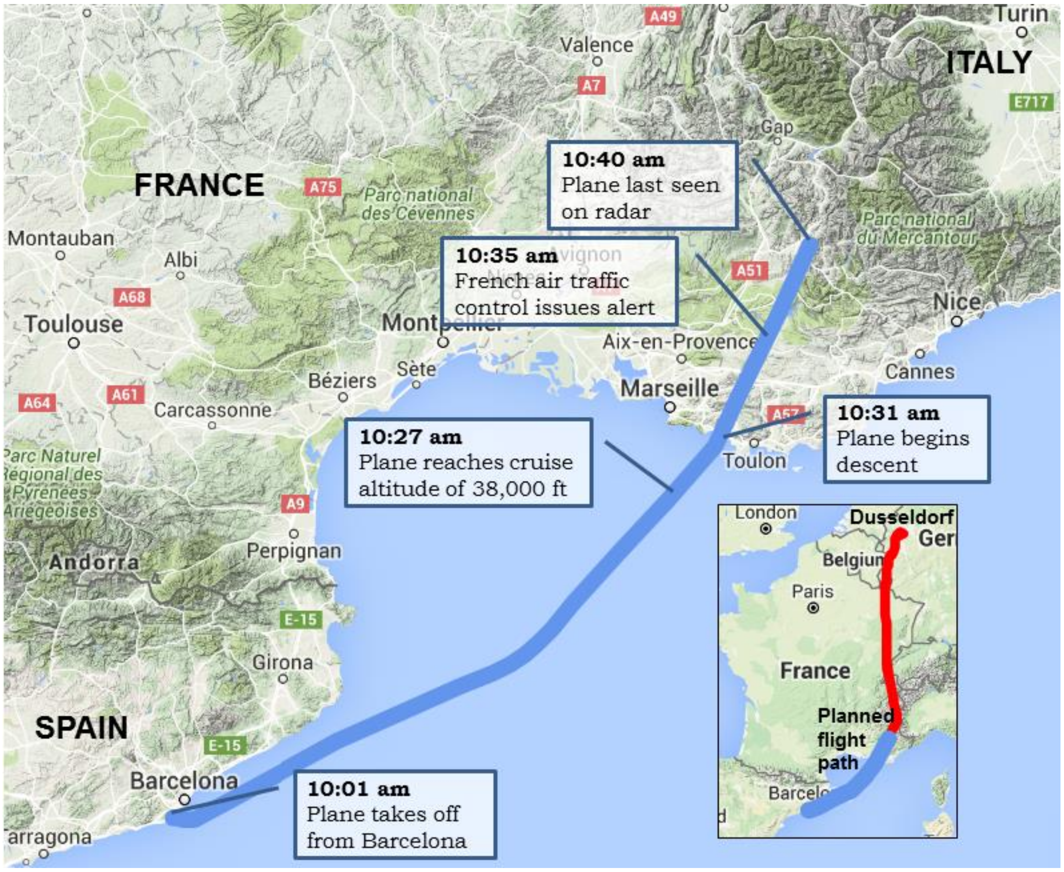

Figure 1. Germanwings flight 4U9525 operating from Barcelona, Spain to Düsseldorf, Germany on March 24, 2015.

other useful information. Since this information is automated and can be received by air traffic control ground stations as well as any other ADS-B receivers, it provides improved situational awareness. Based on the ADS-B technology and data fusion from multiple sources such as flight plan and local weather conditions, it is now possible to automate the flight tracking and to detect, identify and respond to flight anomalies in a timely fashion while reducing the work load of air traffic controllers.

In this paper, a case study of the last flight of Germanwings 9525 will be presented. The flight data will be analyzed and compared to its predecessors from anomaly tracking and prognostics point-of-view. This paper also proposes to use the ADS-B data for flight tacking together with auxiliary information such as weather data in order to determine deviations from flight route and anomalous flight behaviors in real time. The paper is organized in five sections. After the introduction, section 2 will present the specific circumstances around the last flight of Germanwings 9525 and provide information about the accident. Section 3 will elaborate the proposed technique for automated, data-driven flight anomaly tracking. Section 4 will present the results on Germanwings flight case study and the paper will be concluded in section 5. An appendix on Bayesian Networks is included for completeness of the paper.

2. CASE STUDY: LAST FLIGHT OF GERMANWINGS 9525

Germanwings flight 4U9525 was a commercial international passenger flight from Barcelona, Spain to Düsseldorf, Germany on March 24, 2105. The Airbus A320-211 operated by the low-cost carrier was destroyed in an accident over a mountainous area of southern France killing all 144 passengers and six crew members on board.

The Airbus A320-211, registration D-AIPX, was a 24 year old aircraft, which entered into service on 02/05/1991, with more than 58,000 flight hours and 46,000 flights (BEA, 2016). It was equipped with two CFMI CFM56 engines. The last maintenance of the aircraft was performed the day before the accident. According to the METAR data, the weather during takeoff at Barcelona Airport was rainy 12°C, wind from 50 degrees at 16 knots, visibility 10+ km, few clouds at 2000 feet, broken clouds at 4500 feet. Marseille Airport weather at 10:30 am local time was 15°C, wind from 240 degrees at 6 knots.

The flight departed Barcelona, Spain at 10:00 am local time (09:00 UTC) on a regular passenger service to Düsseldorf, Germany. The actual flightpath and the timeline of events can be followed on maps given in Figure 1. The inset of the figure shows the planned flight path. According to the flight data, the plane reached a cruising altitude of 38,000 at 10:27

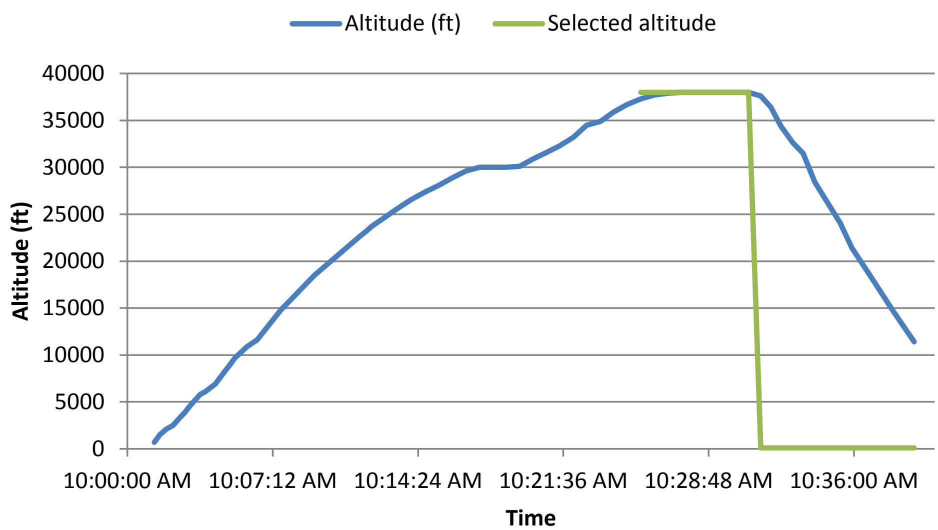

Figure 2. Germanwings flight 4U9525 vertical flight path.

am over the Mediterranean Sea. Vertical flight path of D-AIPX can be seen in Figure 2. According to the accident report (BEA, 2016), the plane was cleared direct to the IRMAR waypoint at 10:30 am by the Marseille control center. An excerpt from the accident report lists the timeline of events afterwards:

"At 10:30:08, the captain told the co-pilot that he was leaving the cockpit and asked him to take over radio communications, which the co-pilot acknowledged."

"At 10:30:53, the selected altitude on the Flight Control Unit (FCU) changed from 38,000 ft to 100 ft. One second later, the autopilot changed to OPEN DES mode and autothrust changed to THR IDLE mode. The aeroplane started to descend and both engines' speed decreased."

"At 10:33:12, the speed management changed from managed mode to selected mode. One second later, the selected target speed became 308 kt while the aeroplane's speed was 273 kt. The aeroplane's speed started to increase along with the descent rate, which subsequently varied between 1,700 ft/min and 5,000 ft/min, then was on average about 3,500 ft/min."

"At 10:33:35, the selected speed decreased to 288 kt. Then, over the following 13 seconds, the value of this target speed changed six times until it reached 302 kt."

"At 10:33:47, the controller asked the flight crew what cruise level they were cleared for. The aeroplane was then at an altitude of 30,000 ft in descent. There was no answer from the co-pilot. Over the following 30 seconds, the controller tried to contact the flight crew again on two occasions, without any answer."

"At 10:34:23, the selected speed increased up to 323 kt. The aeroplane's speed was then 301 kt and started to increase towards the new target."

"At 10:35:03, the selected speed increased again to 350 kt, and until the end of the recording, the selected speed remained at 350 kt and the aeroplane's speed stabilized around 345 kt; the autopilot and autothrust remained engaged."

Between 10:35:07 and 10:39:2, the Marseille control center and French Air Defence system tried to contact the flight crew on multiple occasions and on multiple frequencies without any answer.

According to the flight data recorder, the GPWS started warning: "Terrain, Terrain, Pull Up, Pull Up" at 10:40:41 until the aircraft impacted a sloping rocky ravine in mountainous terrain at an elevation of 1550 meters at 10:41:06.

This record shows that for the first time the air traffic control suspected an anomalous condition 3 minutes after the plane started its descent over France, and after it lost 8000 feet from the cruise altitude it was cleared for. Not until 5 minutes into the descent that the controllers declared an emergency.

3. AUTOMATED FLIGHT ANOMALY TRACKING

Although anomaly detection and loss-of-control analysis is not new, the previous approaches are developed based on aircraft dynamics (Kwatny et al., 2012). Here, we propose a data-driven approach that utilizes probabilistic predictive models (Bishop, 2006 and Sarkar et al., 2014) to compute an estimated flight route and determine the flight anomalies (Rao et al., 2009). We specifically use a model structure called Bayesian Network that learns the statistical flight behavior and predicts the flight route parameters such as latitude, longitude and altitude. The method has four steps, which are shown in Figure 3:

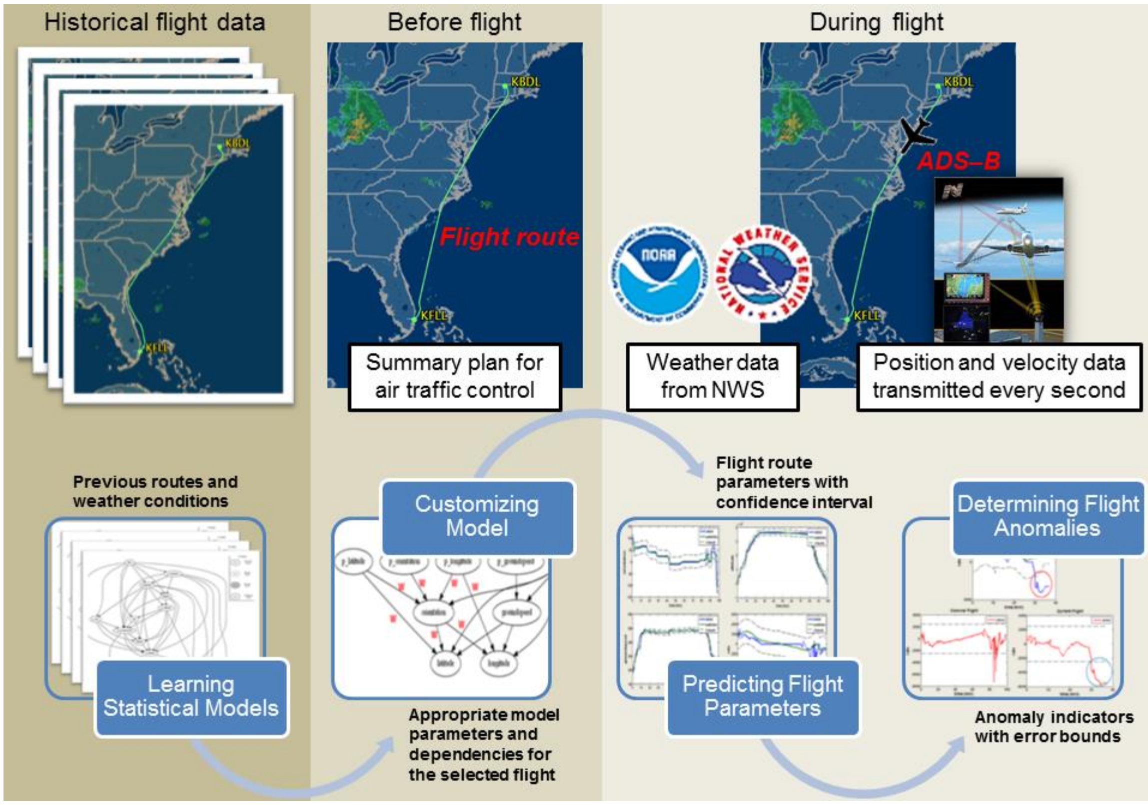

Figure 3. The architecture of automated flight anomaly tracking technique and the step-by-step depiction of preflight and during flight processes.

1. Learning statistical models: Given a flight between two airports, there are a handful of routes (set of way points) that an aircraft can fly. For each possible route, we provide the historical data from previous flights that includes the time series of the aircraft position and weather conditions at the time of flight. Multiple Bayesian Network models are trained using the historical data, each corresponding to a specific route.
2. Customizing the Model: When a summary flight plan is generated for the flight, we pick the appropriate model from the set of models created at Step 1. This is done by simply matching the way points of the summary plan to the way points listed in the possible routes.
3. Predicting Flight Parameters: During the flight, real-time aircraft information (e.g. position and velocity) is received by the ADS-B receiver. At each time-step, the Bayesian Network model, selected at Step 2, predicts the future state by using the current data. The predictions also include a measure of confidence in terms of standard deviation for each flight parameter. This is a desirable feature of Bayesian Networks that differs from other predictive models such as neural networks (Desell, et al., 2014).
4. Determining flight anomalies: While the predicted flight route is computed in real-time at Step 3, the difference between the predicted route and actual flight data is used to determine if the flight is still on track or not. Under normal flight conditions, the actual flight route is expected to be within the confidence interval of the predicted flight route. Similarly, when the prediction errors are persistently out of confidence bounds, the flight behavior is flagged as anomalous.

4. RESULTS OF THE CASE STUDY

With the purpose of detecting the anomalous behavior in Germanwings flight 9525, we analyzed the data from flights on March 13, 14, 16, 17, 19, 20, 23 and 24, 2015 for the same flight from Barcelona, Spain to Düsseldorf, Germany. All these flights followed the same route as seen in Figure 5.

The specific data that was used in this study was obtained from flight tracking website, flightaware.com, and collected immediately after the news of plane crash filled the airwaves. The available flight data consists of time (EDT), latitude, longitude, orientation, ground speed (kts), altitude (ft) and altitude rate of the aircraft. Using the data from flights on March 13, 14, 16, 17, 19 and 20, 2015, we trained a Bayesian Network model shown in Figure 4.

To capture the dynamic nature of the flight, the Bayes Net is structured in a way that the previous states of flight conditions, denoted as "p_" nodes, are connected to the current states of the flight in a directed acyclic graph. The

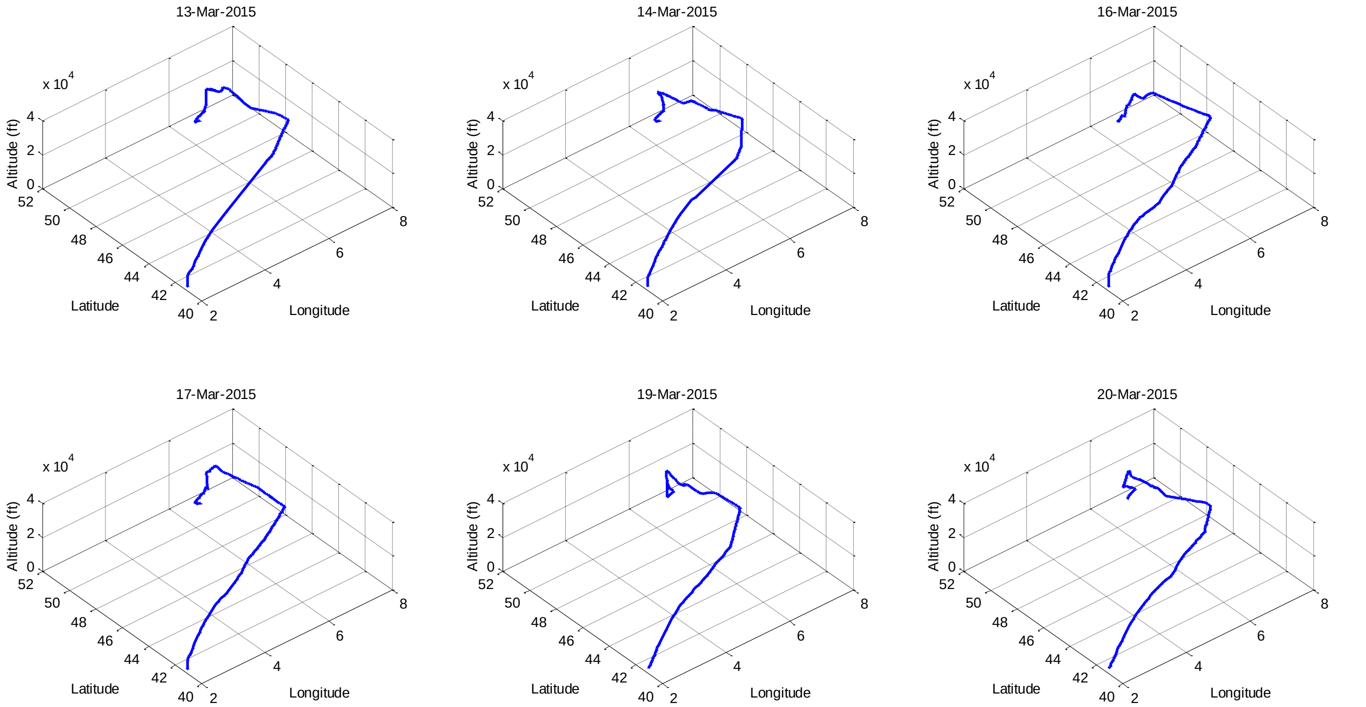

Figure 5. Germanwings flight 4U9525 flight path on March 13, 14, 16, 17, 19 and 20, 2015. Only the data that was used for Bayesian Network training is shown.

"time" node corresponds to the time difference between the start of the flight and the time of the current state. In this way, the dynamic motion of the flight is captured as a Markov process. During the training phase, the Bayes Net learns the joint probability density functions between the previous states, the time difference and the current states.

The flight data from March 23, 2015 (a normal flight) and March 24, 2015 (the doomed flight) was used to predict the route and compare the results of two flights that had two very different outcomes. For predictions, at each time instance the previous flight states are provided as evidence to the Bayes Net and the learned parameters are used to predict the current state of the flight. The predicted current state is then compared to the actual current state to derive the error between the predictions and actuals.

For the March 23, 2015 flight, the flight state predictions, the prediction confidence bounds and the actuals are shown in Figure 6. The prediction errors are only instantaneously out of confidence bounds and the Bayesian Network does a good job of estimating the flight route. We also used the same model to analyze the data of flight on March 24, 2015. This aircraft was destroyed in an accident in a mountainous area in southern France. Upon crossing the French coastline east of Marseille, the airplane began losing altitude and descended from 38,000 feet to 6,800 feet at an average rate of 3710 feet/minute. Our analysis indicated an anomaly in flight altitude rate almost immediately. The Bayes Net flight route prediction results are shown in Figure 7. As seen in the figure, latitude, longitude, orientation and ground speed of the aircraft are within the confidence bounds of the learned model. However, the predicted and the actual altitude and the altitude rate quickly diverge during the accident and move beyond the confidence bounds.

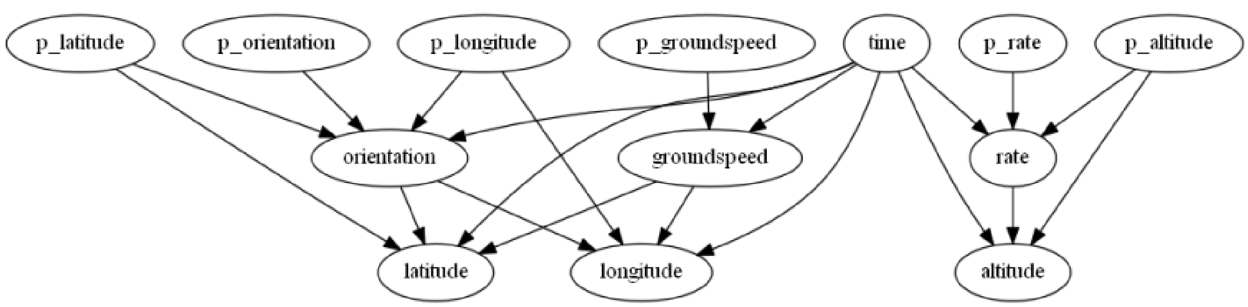

Figure 4. The Bayesian network that was trained to learn the nominal flight behavior.

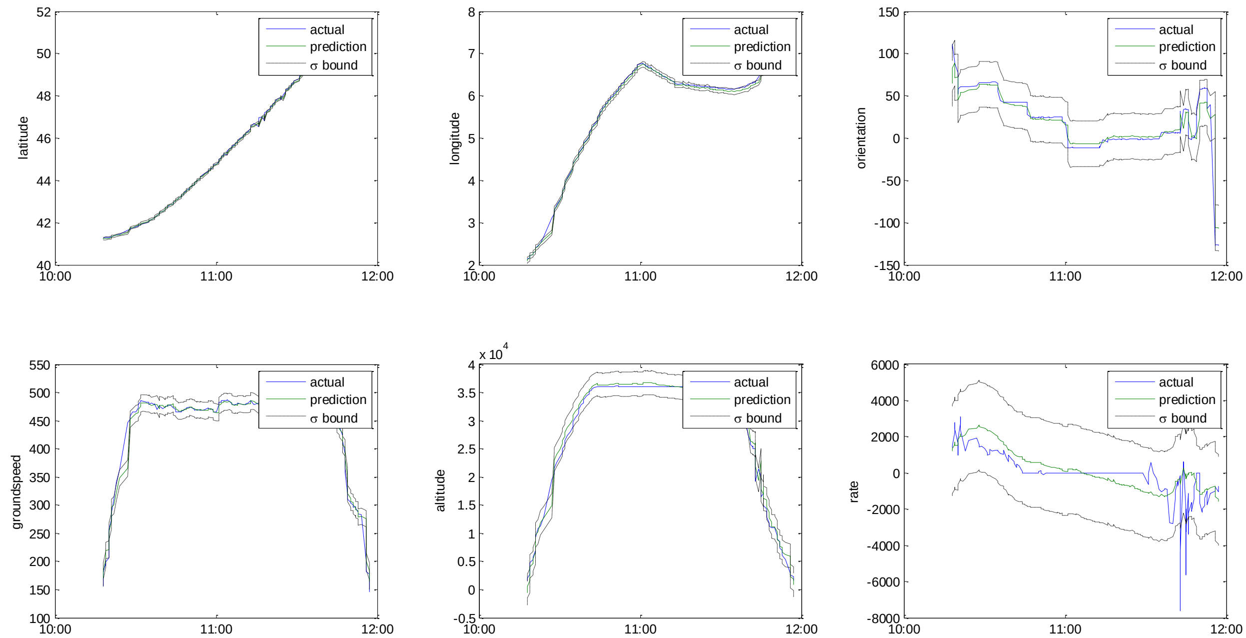

Figure 6. Germanwings flight 4U9525 flight path predictions on March 23, 2015.

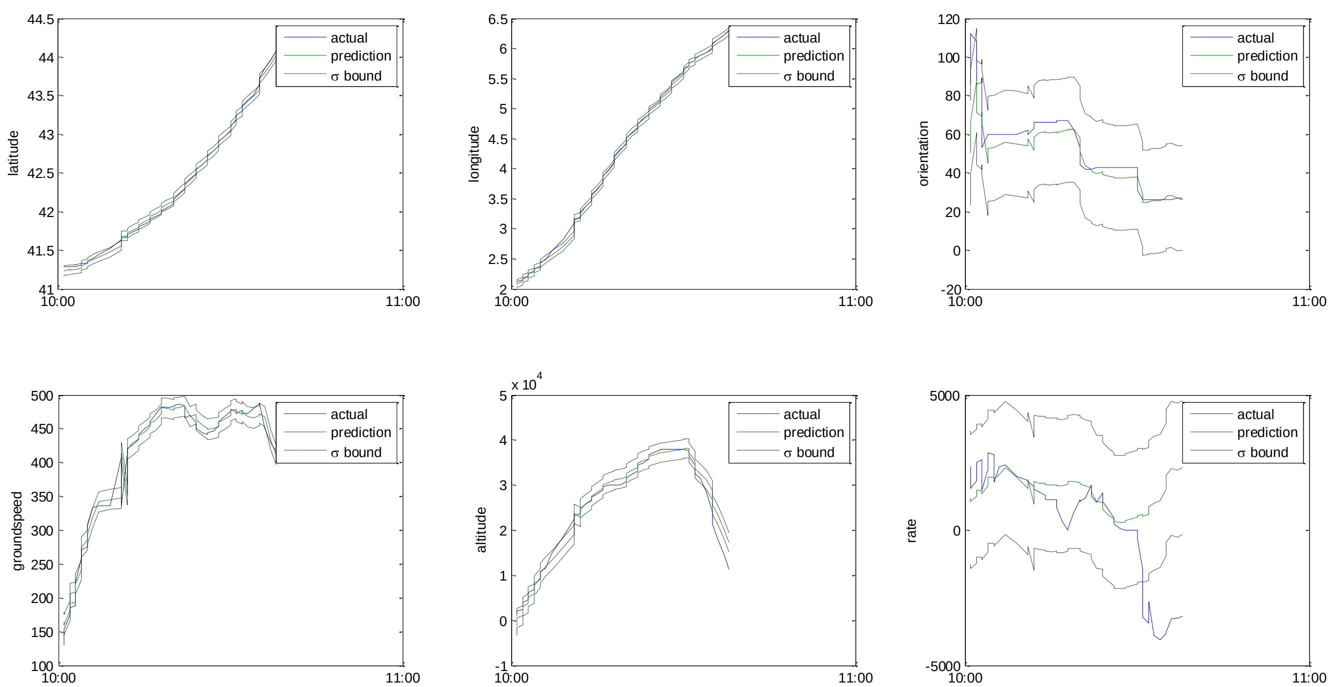

Figure 7. Germanwings flight 4U9525 flight path predictions on March 24, 2015.

The reason for the predicted altitude to also decrease during the incident is the dynamic nature of the predictive model. It essentially makes the predictions based on the previous (actual) states. It could be also be possible to use the previous predictions for calculating the current predictions, but this would result in confidence bounds to get larger at each step. It should be noted that it is further possible to have multiple previous flight states as nodes of the Bayes Net for more accurate predictions, albeit with additional computational cost.

5. CONCLUSION

The study presented in this paper is motivated by an air disaster occurred on March 24, 2015 over the mountainous region of Southern France. On that day, Germanwings flight 4U9525 was perished with all its crew and passengers on board. Although global aviation safety has been steadily improving over the last decades, it is still not possible to respond to all flight anomalies and emergencies in time.

The ever increasing computational capacity and sophistication of machine learning and big data technologies, it might be possible to have automated monitoring of multitudes of flight in real-time. The technologies such as the one described in this paper will help improve the aviation safety to a higher level. It is the opinion of the author that these technologies will eventually make their way into the air transportation system. They will increase the situational awareness and help air traffic control to analyze the emergency situations faster and more effectively. Although there is no commercial technology currently available for flight anomaly detection, intelligent decision making based on historical data and probabilistic models that perform real-time predictive analysis are not too far away.

REFERENCES

Bureau d'Enquêtes et d'Analyses (2016). *Final Report, Accident on 24 March 2015 at Prads-Haute-Bléone (Alpes-de-Haute-Provence, France) to the Airbus A320-211 registered D-AIPX operated by Germanwings*. France: BAE.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

Desell, T., Clachar S., Higgins, J. & Wild, B. (2014). Evolving Neural Network Weights for Time-Series Prediction of General Aviation Flight Data. *Parallel Problem Solving from Nature – PPSN XIII*, pp. 771-781.

Kwatny, H. G., Dongmo, J. T., Chang B. C., Bajpai G., Yasar, M., & Belcastro C. (2012). Nonlinear Analysis of Aircraft Loss of Control. *Journal of Guidance Control and Dynamics*, vol. 36, no. 1, pp. 149-162.

Rao, C., Ray, A., Sarkar S. & Yasar, M. (2009). Review and comparative evaluation of symbolic dynamic filtering for detection of anomaly patterns. *Signal Image and Video Processing*, vol. 3, no. 2, pp. 101-114.

Sarkar, S., Sarkar, S., Virani, N., Ray A. & Yasar, M. (2014). Sensor Fusion for Fault Detection & Classification in Distributed Physical Processes. *Frontiers in Robotics and AI*, vol. 1, article 16.

ACKNOWLEDGEMENT

The author would like to thank Dr. William Veronesi.

BIOGRAPHIES

Murat Yasar, Ph.D. is with the United Technologies Research Center (UTRC) working on diagnostics, prognostics and PHM applications related to elevators, HVAC equipment, gas turbine engines and aircraft components. He holds Master's degrees in Mechanical Engineering and Electrical Engineering and a Doctorate in Mechanical Engineering, all from the Pennsylvania State University, University Park, PA. Before joining UTRC, Dr. Yasar was involved with various programs ranging from aircraft health management to supervisory control design for shipboard systems at Techno-Sciences, Inc. and InnoVital Systems, Inc. for seven years. Dr. Yasar is an expert in control systems design, signal processing and predictive analytics and has worked on many projects as a principal investigator throughout his career. His recent focus has been on development and implementation of machine learning and predictive analytics for PHM applications. Dr. Yasar is a member of ASME and IEEE.

APPENDIX: A BRIEF SUMMARY OF BAYESIAN NETWORKS

Bayesian Networks (also known as Bayes Nets or belief networks) provides a framework to model the joint distribution of random variables. They belong to the family of models called probabilistic graphical models. Essentially, a Bayes Net is a statistical machine learning tool that allows one to build probabilistic models from data. What is powerful about Bayes Nets is their capability of computing reasonable predictions under partially observed data. Another desirable property is built-in uncertainty quantification on the model output, which is particularly useful in the presented application. There have been many applications of Bayes Nets in the literature, specifically for two types of data modelling: the first one includes learning model parameters (i.e. the conditional distributions of random variables) given a structure of a directed acyclic graph (DAG), and the second one includes learning both the model structure (i.e. graph topology) and the parameters. In this application, the model structure was provided for the Bayes Net to be able to incorporate the domain knowledge and specifically define the time dependencies.

As an example, the parameter learning process might involve determining the following joint probability function for random variables x_1, x_2, x_3, x_4, x_5:

$$P(x_1, x_2, x_3, x_4, x_5) = P(x_1)P(x_2 \mid x_1)P(x_3 \mid x_1, x_2)P(x_4 \mid x_3)P(x_5 \mid x_3, x_4)$$

It involves an iterative maximum likelihood estimation step starting from (usually random) priors. A standard approach is the expectation-maximization algorithm which alternates computing expected values of the unobserved variables with maximizing the likelihood.

Once the parameters of the joint probability function are learnt there are multiple inference algorithms available for Bayes Nets. For this application a well-known junction tree algorithm was used for exact inference. It applies the explicit Bayes Rule to determine the expected values and the variances of the unobserved variables based on evidence (i.e. observed data). What is further possible is to use soft evidence where a distribution of values could be provided as evidence.