

INTENT INFERENCE FOR FREE FLIGHT AIRCRAFT

Jimmy Krozel, Ph.D.
Metron Aviation, Inc.
131 Elden Street, Suite 200.
Herndon, VA 20170

Abstract

In today's air traffic control system, a level of intent of an aircraft is revealed in its flight plan and a pilot's adherence to published navigation routes. In the future Free Flight environment, aircraft may be allowed to fly any route they choose. Intent will likely be broadcast in the form of Trajectory Change Points (TCPs), for instance, the next two TCPs called TCP and TCP+1 in an Automatic Dependent Surveillance – Broadcast (ADS-B) message. However, if TCPs do not accurately represent intent, do not get received, or do not exist, then an aircraft has a need to infer the intent of another aircraft. In this paper, a method of inferring intent is investigated which is based on artificial intelligence models of actions and a process for best fitting an intent model to observed actions of an aircraft. With such an intent inference algorithm, the TCPs for an aircraft can be inferred and used when the TCP data in a datalink message cannot be validated. While intent inference may not be needed in a normal situation, its use may enrich the presentation of data to the flight crew in a Free Flight environment.

Introduction

As new Free Flight procedures remove jetway routing, positive control, and other constraints [14], an added emphasis will be placed on distributed control techniques. In Free Flight, a datalink technology, for instance, ADS-B, will be used to communicate state and intent data between aircraft. State and intent information (e.g., position, speed, and TCPs) will be

exchanged between aircraft. In this paper, we focus on a datalink which provides a TCP and TCP+1, although, in general, the approach also applies to the case of n TCPs or TCP, TCP+1, ..., and TCP+ n .

There is a need for intent inference in Free Flight since there is no guarantee, even if flight plans or intent are broadcast to nearby aircraft, that these plans will be followed. Even with acceptable Required Navigation Performance (RNP) levels which stipulate how closely the flight crew will locate the aircraft relative to the intent, there still is the possibility that the crew might have changed their plans without inputs into the FMS or ADS-B broadcast, the equipment for ADS-B broadcast might fail, or deviations from a flight plan might be excessive due to poor wind predictions, weather avoidance, or emergency diversions.

Figure 1 illustrates the intent inference problem. In this example, an aircraft has lateral deviations due to Flight Technical Errors (FTE) while flying to a way point, severe weather is just ahead, and there is a good reason why the aircraft might be (1) flying North of the storm (because the current heading is that way), (2) flying directly to the next way point and ignoring the weather (because the flight plan indicates this), or (3) flying to the South of the storm and skipping a way point (because airlines have policies to avoid storms and arrive on schedule). In general, the problem of intent inference is to determine: What is the intent of the aircraft?




Figure 1. What is the intent of the aircraft?

While intent data is broadcast from one aircraft to another, the following cases can occur:

1. The broadcast intent message is present, but needs to be verified true;
2. The broadcast intent message is present, is verified false (or erroneous), and an intent must be inferred from all available information; or
3. The broadcast intent message is missing (due to data dropout or an unequipped aircraft), and the intent needs to be inferred.

In each case, the intent inference algorithm designed in this paper will result in the intent of the aircraft and a level of confidence.

In the air traffic management domain, there are certain rules and procedures which help explain aircraft actions. Examples of data that indicate useful intent information include: the filed flight plan, Standard Instrument Departures (SIDs), Standard Arrival Routes (STARs), and others (listed later) which often stipulate a set of well defined and ordered actions. In the literature, Zhao, et al [19] specify several classifications of such pilot intent models, including: motive intent, objective intent, trajectory intent, and cost intent. These actions are modeled in an intent inference knowledge database using several artificial intelligence models, and the knowledge base can be searched for an inferred intent with one of several correlation functions. We present an approach to solving this problem which combines a tracking filter (Kalman filter) to estimate the state information and to predict short term trajectories, and intent models to predict long term paths.

Our approach to intent inference has unique benefits and does not parallel any other approach in the literature. Other approaches include:

- Operator Plan Analysis Logic (OPAL) [3, 4, 12], where plan-goal graphs are used to explain the causal relationship between the observed discrete actions and the goals or intent of the operator.
- Operator Function Model expert system (OFMspert) [1, 13], where an expert system uses a blackboard architecture and operator function models, a heterarchic-hierarchic network of nodes that represent operator control functions, to replicate and infer an explanation of the actions of the human operator.

- Generalized plan recognition [5], where a deductive inference system performs intent inference by recognizing the plan an operator is implementing, based on observations of actions, and an action taxonomy (an exhaustive set of sequences of discrete admissible actions).
- Event tracking using Soar [18], where the Soar knowledge based system reasons about the decisions made in the task of the operator, triggering rules that create problem space, goals, and subgoals that lead to an explanation.

These methods are primarily used for discrete intent inference problems and are only in special cases applicable to the continuous intent inference problem.

As shown in Table 1, continuous intent inference fills in the gap between trajectory prediction and discrete intent inference. Tracking filter-based prediction, for instance with a Kalman filter, while useful for short time periods ahead of an aircraft, does not address the intent inference problem when it comes to explaining factors like avoiding weather, Special Use Airspace (SUA), turbulence, etc., which are factors not easily incorporated into the aircraft equations of motion modeled in most tracking filters. Using a tracking filter is most applicable when *coupled* with the intent inference method as we discuss later in this paper.

The next sections of this paper detail the theory behind intent inference, present some demonstration scenarios, and state our conclusions.

Theory

The intelligent inference algorithm is based on the theoretical foundation *Intelligent Path Prediction for Vehicular Travel* [6,7,8,9]. Our solution incorporates any broadcast or datalinked information from nearby aircraft being tracked. Furthermore, when domain knowledge data are available, e.g., traffic data, weather data, terrain data, SIDs, STARs, etc., these data are taken into consideration. The intent inference module is able to verify that a tracked aircraft is following the intent that is broadcast, and, if the aircraft is seemingly following some other intent, then the module will identify the most plausible intent(s) for the aircraft being tracked. Finally, the intent inference module uses the most plausible intent to predict the future path of the vehicle being tracked.

Table 1. Prediction theories, applications, and characteristic time horizons.

PREDICTION THEORY	Tracking Filter Prediction	Continuous Intent Inference	Discrete Intent Inference
APPLICATION	Trajectory Estimation and Prediction	Path Prediction	Discrete Event Changes
TIME HORIZON	Seconds – Tens of Seconds	Tens of Seconds – Minutes	Asynchronous / Event Driven

Basis of Theory. We analyze the process of human flight control through two classical theories: control theory and psychology. From control theory, flying an aircraft incorporates: stability, control, guidance, and navigation. From psychology, reasons for making control decisions while flying an aircraft are explained by understanding human decision making processes.

Based on the analysis of Stengel [17], a model for intelligent flight control is postulated. This model includes sensing, regulation, and decision making, as in Figure 2. The sensory inputs to the human are associated with seeing (states from displays, other aircraft out the window, runways, etc.), feeling (forces felt by the body), and hearing (sounds correlated with speed, engine noise, and inner ear effects from gravity) – taste and smell are assumed negligible in the model. These senses are connected to the brain to influence decision making. Decision making governs regulation, with neuro-muscular responses driven by learned associations between stimuli and desirable actions. Furthermore, through the decision making process the pilot plans and sets goals that effect navigation and guidance command inputs. Such planning and decision making is dependent on the pilots knowledge base of both rules of flight and past experience.

For the aircraft model, navigation, guidance, and control functionality might be supplied by the pilot or by computer-based control systems (e.g., FMS). For the modestly equipped aircraft, a pilot will perform all guidance and navigation logic and will combine these with regulation logic to affect control inputs (minimally, the stick, throttle, and rudder). In the most advanced cockpits, navigation, guidance, and

control systems are automated with potentially minimal input by the pilot. For instance, autopilots allow the pilot to command a constant heading angle or airspeed. The maneuvering precision will differ based on whether the tasks of navigation, guidance, and control logic are performed by a human pilot or by automatic control.

The human decision making model has a hierarchical structure. In the inner loops, stability augmentation and reflexive control are performed with fast update rates, perhaps 1-10 Hz. In the outer loops, guidance and navigation changes may be made every 1-10 minutes. The central nervous system supports such a hierarchy for decision making by providing a structure that includes declarative actions, procedural actions, and reflexive actions – the basis of the intelligent flight control system proposed by Stengel [17]. In the research of intent inference, we will next show that this hierarchical structure and separation between declarative actions from reflexive actions is also useful to model.

Consider the relationship between the human decision making process of flying an aircraft and the problem of predicting the intent and path of an aircraft being tracked by an outside observer. These concepts are similar but different:

- **Intent Inference** is the process of determining *what* the tracked vehicle is most likely attempting to do.
- **Path Prediction** is the process of determining *how* the tracked vehicle will most likely accomplish what it is intending to do.

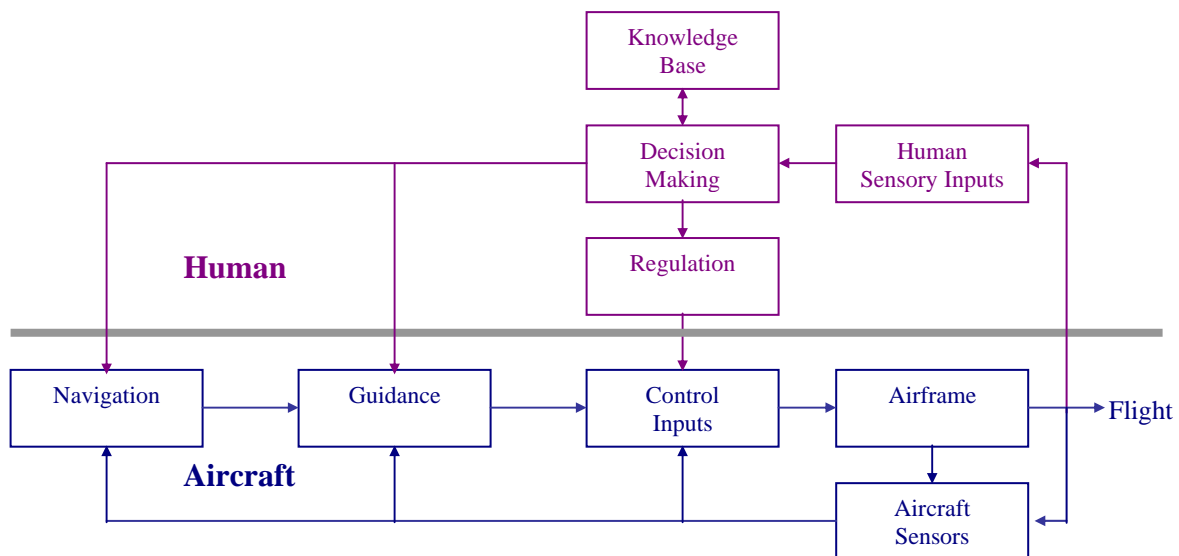


Figure 2. Process for human decision making while flying an aircraft.

The path of an aircraft is directly a result of the control inputs determined by regulation and/or guidance decisions. The pilot intent is to use guidance and control inputs to follow a path plan, e.g., a flight plan. In accordance with the aircraft flight control model set forth so far, intent inference is related to inferring the declarative and procedural decisions of the pilot, and path prediction is related to inferring the path that the pilot attains from regulatory and reflexive control inputs.

Intent inference and path prediction are solved simultaneously in our approach. We first discuss these individually, and then show how the two are combined. Intent inference may be considered to be at a higher level of abstraction than path prediction, since intent inference specifies *what* and path prediction specifies *how*. Predicting the intent of a vehicle can be abstracted into the process presented in Table 2. Likewise, predicting the motion of a vehicle can be abstracted into the process presented in Table 3.

Since the processes for intent inference and path prediction relate in a hierarchical nature, we combine these processes into a single algorithm that both predicts intent and the future path of the aircraft, as in Table 4. Note that because of the hierarchical nature of decision making, intent is likely to remain constant for a long time as inner loop commands change more frequently. If path prediction is pursued with very short look-ahead times, then the outer loop guidance and navigation commands can be assumed constant. However, as look-ahead times become longer, one should expect that the guidance and navigation commands may change, depending on how far along the process the pilot is in terms of following the outer loop decision logic. Thus, the model used for the path prediction must adapt to the context set forth by the outer loop decisions. From this discussion, the hierarchical structure explains why intent inference is necessary to do accurate path prediction, especially when the prediction time is long enough to include declarative or procedural decisions.

Table 2. Process for intent inference.

Step	Process
1	Build a knowledge database of plausible intent models for the vehicle being tracked and domain knowledge about the environment the vehicle is traveling through,
2	Observe and record state information about the vehicle being tracked and nearby traffic,
3	Combine the plausible intent models with the observed state information and determine the “correlation” of fit, and
4	Rank the intent models by “correlation” of fit to determine the most likely intent of the vehicle.

Table 3. Process for path prediction.

Step	Process
1	Build a dynamics model for the vehicle being tracked, based on the inferred intent,
2	Observe and record state information about the vehicle being tracked,
3	Estimate states and/or parameters in the dynamics model whether observed directly or not, and
4	Project the dynamics model into the future in order to predict future motion.

Table 4. The combined process for intent inference and path prediction.

Step	Process
1	Build a knowledge database of plausible intent models for the vehicle being tracked and domain knowledge about the environment the vehicle is traveling through,
2	Observe and record state information about the vehicle being tracked and nearby traffic,
3	Build a dynamics model for the vehicle being tracked, based on the plausible intent,
4	Estimate states and/or parameters in the dynamics model whether observed directly or not,
5	Combine the plausible intent models with the observed state information and determine the “correlation” of fit,
6	Rank the intent models by “correlation” of fit to determine the most likely intent of the observed vehicle, and → RETURN INTENT
7	Project the dynamics model into the future in order to predict future motion. → RETURN PREDICTED PATH

Intent Inference Knowledge Base and Domain Knowledge.

The intent inference knowledge base consists of plausible cost criteria for travel and domain knowledge. The first step in building an intent inference module is to use knowledge engineering and operations analysis to build this knowledge base.

Intent inference is based on identifying a particular intent model from a set of cost criteria that explain intent. For this purpose, the aircraft flight domain has a rich set of cost criteria that will reside in the intent inference knowledge base, including but not limited to the cost criteria listed in Table 5. A database of intent models holds all the associated constants and parameters needed to provide the direction that an aircraft should proceed given any arbitrary set of initial conditions. While some intent models are relatively simple, for instance the model to hold heading, others are very complex, for instance the model to avoid weather.

The intent inference knowledge base requires domain knowledge. Domain knowledge is used to build the situation assessment model, or a map of the situation, which includes the factors that might influence the pilot’s decision making. The domain knowledge database may include the information in Table 6. Our particular application includes the following:

1. Inputs via a data fusion – this module takes inputs from ADS-B [16], Flight Information Services (FIS), and Controller Pilot Data Link Communications (CPDLC).
2. Other onboard systems – these modules include the GPS receiver, the Conflict Detection and Resolution (CD&R) module, and hazard avoidance module, which provide way point data to the intent inference module.
3. Onboard databases – these databases provide static data, including airport, navigation, SUA, and terrain elevation data.

Method of Intent Correlation. The actions of a decision maker may be observed and analyzed both locally (instantaneously) and globally (over a time window). For instance, a local decision is a change in state in an instantaneous or very short time horizon, such as a single decision to hold course, turn, climb, descend, etc. A global decision is a series of local decisions which occur over a longer time horizon, perhaps a sequence of decisions to solve a navigation task to capture a way point, follow a STAR, avoid weather, etc.). For instance, while manually flying a flight plan, if a pilot turns to capture a way point after drifting away from the nominal flight plan, the local decision might be to turn left or right towards the way point. If one looks at the state variables at any one instant, it might be difficult to infer that the pilot is flying to a way point,

Table 5. Plausible cost criteria for intent inference in the aircraft Free Flight domain.

Intent Description	Required Data and Domain Knowledge Needed
Hold Heading	Velocity Data and Wind Data
Hold Altitude	Altitude Data
Hold Speed	Velocity Data or Airspeed Data
Hold Course	Position, Velocity, and Wind Data
Hold Coordinated Turn	Velocity Data and Bank Angle Data or Turn Rate Data
Go To Way point (fly over)	Position and Velocity Data and TCP or Flight Plan Data
Skip Way point; Fly to Next Way point (fly by)	Position and Velocity Data and TCP+1 or Flight Plan Data
Fly Direct To (way point, final approach fix, metering fix, or airport location)	Position and Velocity Data and TCP or Flight Plan Data or Airport, Metering Fix, or Final Approach Fix Data
Return to Flight Plan from Heading Deviation	Position and Velocity Data and Flight Plan Data
Return to Flight Plan from Speed Deviation	Position and Velocity Data and Flight Plan Data
Return to Flight Plan from Altitude Deviation	Altitude and Velocity or Altitude Rate and Flight Plan Data
Return to Flight Plan from Lateral Offset	Position and Velocity Data and Flight Plan Data
Fly a list of Waypoints (Flight Plan, SID, or STAR)	Position and Velocity Data and SIDs and STARs Database or Flight Plan Data
Avoid Special Use Airspace (SUA)	Position and Velocity Data and SUA Data
Avoid Turbulence	Position and Velocity Data and Turbulence Report Data
Avoid Hazardous Weather	Position and Velocity Data and Weather Precipitation Data
Avoid Aircraft	Position and Velocity Data for Ownship and Intruder(s)
Avoid Terrain	Position and Velocity Data and Terrain Database
Avoid Hazard Region (General)	Position and Velocity Data and Hazard Region Data
Fly Holding Pattern	Position and Velocity Data and Holding Pattern Data

Table 6. Data categories and domain knowledge databases needed for intent inference.

Data Description	Database or Source of Information
Position (Latitude, Longitude, Altitude)	State Data Estimates from Onboard Data Fusion Module
Velocity (North, East, Vertical)	State Data Estimates from Data Fusion Module
Acceleration (North, East, Vertical)	State Data Estimates from Data Fusion Module
Bank Angle	State Data Estimates Module
Airspeed	State Data Estimates Module
TCP (Way point)	ADS-B Module; Navigation Aid Database
TCP+1 (Way point)	ADS-B Module; Navigation Aid Database
Navigation Equipment Performance	ADS-B Module
Flight Plan	CPDLC Module
SIDS and STARS	Navigation Aid Database
Airport Locations	Airport Database or Flight Plan
Alternate Airport Locations	Airport Database or Flight Plan
Final Approach Fix Locations	Airport Database
Metering Fix Locations	Airport Database; Navigation Aid Database
Conflict Resolution Way points	Conflict Detection and Resolution CD&R Module
Precipitation Hazard	FIS Module
Precipitation Hazard Avoidance Route	Hazard Avoidance Module
Special Use Airspace (SUA)	FIS Module
Special Use Airspace Avoidance Route	Hazard Avoidance Module
Turbulence	Turbulence Data from FIS PIREP
Turbulence Airspace Avoidance Route	Hazard Avoidance Module
Terrain Data	Terrain Database
Terrain Avoidance Route	Hazard Avoidance Module
Hazard Region (General)	Hazard Region Data
Hazard Region Avoidance Route	Hazard Avoidance Module

way point, as depicted in Figure 3. For this paper, we assume that the heading direction and the course are the same directions. Mathematically, by defining a unit vector $\vec{\psi}$ in the direction of the heading ψ and a unit vector $\vec{\phi}$ in the direction ϕ to the way point, then the dot product:

$$\text{Local Correlation} = \vec{\psi} \cdot \vec{\phi} \quad (1)$$

will indicate (locally) if the aircraft is heading to the way point (dot product of 1) or away from the way point (dot product -1). This dot product acts as a good measure of correlation (providing a number between -1 and 1) between the intent to fly directly to the way point and the current aircraft state as indicated by the current heading.




Figure 3. A pilot is not heading directly to a way point, but is generally heading that way.


Next, this local analysis of the intent of the aircraft is extended to a global analysis by considering the correlation between a history of state variables with a series of decisions that would support a consistent intent. For the fly-to-next-way-point example, one would continuously integrate the local correlation measure (the dot product) over the observed flight path:

$$\text{Global Correlation} = \frac{1}{k} \int_{\text{flight path}} \vec{\psi} \cdot \vec{\phi} ds \quad (2)$$

where ds is a differential element along the flight path and k is a non-dimensionalization constant. In this research, a good value for k is simply the arc length between the start location of the aircraft and some characteristic domain location (e.g., current location, a way point, airport, or other location):

$$k = \int_{\text{flight path}} ds. \quad (3)$$

In previous work [6], several global correlation functions were investigated, including fading window and moving window functions, as shown in Figure 4.



Previous Optimal Directions to Way Point

Figure 4. A global analysis investigates the previous headings and previous optimal directions to the way point in order to collect data to support or reject the conclusion that the aircraft is flying to the way point.

The global correlation measure can be used to validate the broadcast intent message. The global correlation measure has a range:

$$-1 \leq \text{Global Correlation} \leq 1. \quad (4)$$

If the broadcast intent is used to model the unit vectors $\vec{\phi}$, as explained in the next section, and if the correlation value is high and close to the value 1, then the broadcast intent can be considered valid. The global correlation can also be used to set a level of confidence for the intent. If the correlation value is zero or negative, then the broadcast intent is either invalid or is simply a poor model of intent.

Engineering judgement must be exercised to determine when to decide if the global correlation value is high enough to consider an intent validated. Experimentation is used to investigate the global correlation values (defined with a moving window) for typical circumstances of aircraft attempting to follow a constant intent but having a typical guidance system and FTE. After experimentation, a cutoff for identifying valid vs invalid broadcast intent messages is determined. In any event, if the intent inference algorithm is capable of presenting an intent that better explains the motion of the aircraft in comparison to the intent that is broadcast by the aircraft, then this additional information should probably not be dismissed. More than one explanation of intent might apply at one time, for instance, flying a particular flight plan, flying to the TCP, and holding altitude all have very high global correlation values and are all quite valid intents that could be occurring simultaneously.

Decision Making Maps for Plausible Intents. A cost of travel or cost criterion C determines the optimal (or near-optimal) decisions that would be made to solve vehicle navigation tasks given any initial condition. For instance, the minimum distance to a way point has a local cost $c=1$ and cost model:

$$C = \int_{\text{flight path}} 1 ds. \quad (5)$$

The generalized cost of travel is modeled as:

$$C = \int_{\text{flight path}} c(x, y, z, t) ds \quad (6)$$

where the local cost function $c(x, y, z, t)$ may be a function of many domain variables (e.g., the presence of weather, turbulence, another aircraft, SUA, etc., as recorded in the intent inference knowledge base), which determines a positive value for c . In a specific intent model, there is only one value for c defined for every point in space (x, y, z) and for any time t . This is the most general form of the cost of travel C which forms the basis for our work.

Mathematically, the motion of the aircraft over time determines the set of heading unit vectors $\vec{\psi}$ which point in the heading direction. The gradient ∇C determines the direction for the unit vector $\vec{\phi}$. A map of unit vectors $\vec{\phi}$ can be derived from the gradient ∇C at multiple points as in Figure 5. In mechanization, the entire map is not needed, only the gradient information at the current and past locations of the aircraft are needed for the correlation function, eq. (2). In our work, we specify how the direction for $\vec{\psi}$ is established within the intent inference module or by analyzing the results from a function calls to other aircraft system modules. Note that the ‘‘correlation’’ of fit is determined by correlating unit vectors $\vec{\phi}$, determined by the gradient ∇C , with the observed state data which determines $\vec{\psi}$. If future motion is to be predicted, the cost criterion C that best explains the motion of the aircraft is used to generate the predicted future motion by exploiting dynamic programming.

Hazardous airspace constraints may be modeled for SUA, turbulence regions, or hazardous weather avoidance regions. A *constraint region* $C_{NO GO}$ is simply defined by a list of vertices that specify a polygon region, as in Figure 6. Figure 7 illustrates how the minimum distance route to a way point can be determined using the SUA as a constraint, that is, all the points (x, y) within the constraint region $C_{NO GO}$ are weighted a very high weight and the dynamic programming solution avoids passing through these regions. In general, a hazard avoidance algorithm may have to be consulted to arrive at the set of way points that avoid a hazard. In such cases, a set of hazard avoidance routes is more useful than simply the optimal hazard avoidance route, since the pilot may fly any of the hazard avoidance routes. As discussed in [10] for hazardous weather avoidance, Optimal Path Maps (OPMs) can serve this purpose.

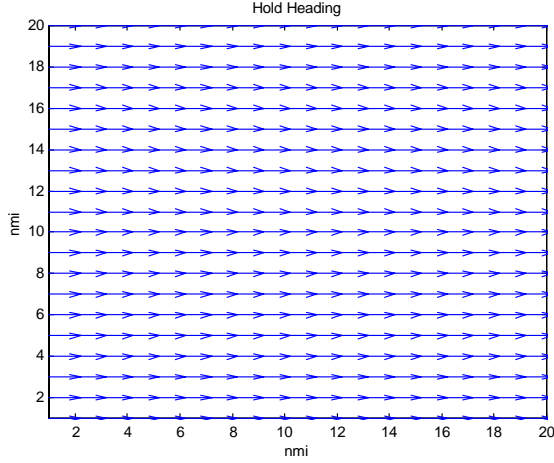


Figure 5. Example of the optimal decisions (headings) that should be used to fly constant heading at any initial condition.




Figure 6. Modeling of a Special Use Airspace (SUA) region North of Las Vegas, NV.

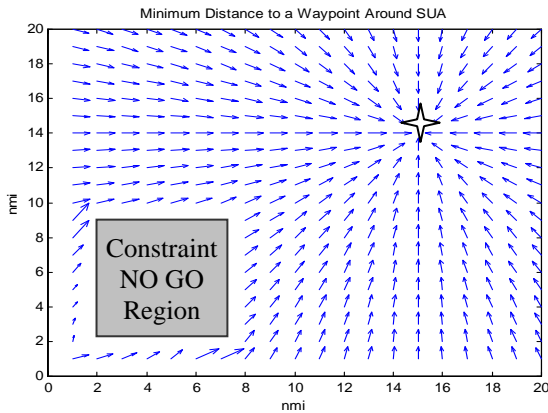


Figure 7. Example of the optimal decisions (headings) that should be used to arrive at the way point using minimum distance to travel to the way point subject to avoiding SUA airspace from any initial condition.

Scenarios

Some scenarios are presented next. First, we present a theoretical scenario to understand analytical variations in intent model correlation values with respect to some common variations in flight conditions. A second scenario demonstrates the merits of several plausible intent models explaining the motion of an aircraft.

Theoretical Scenario. The geometry of this scenario is described within Figure 8, where an aircraft is attempting to fly along the x -axis from the origin to the way point located at position $x=a$. The location of the aircraft is described by (x,y) . In this analytical study, we consider flight paths that are modeled through analytical mathematical expressions.




Figure 8. A pilot with aircraft located at (x,y) is heading to a way point located at $(a,0)$.

In this scenario, we consider intent model for the cost to travel to have a uniform unitary weight over the horizontal plane $c(x,y,t) = I$, to model a pilot trying to fly a straight line from the origin to the way point, thus minimizing the distance to the way point. To define the constant k , we use the arc length of the path flown. We also consider a pilot who drifts off of the x -axis in a motion that is described by the analytic function:

$$\begin{aligned} x &= v_x t \\ y &= A \sin\left(\frac{n\pi t}{a}\right) \end{aligned} \quad (7)$$

where $v_x=1$ is the ground speed, t is time parameter, n is a variable that describes the sinusoidal drift with amplitude A of a fictitious flight path described by (x,y) , and a describes the location of the way point, as shown in Figure 8. The functional evaluation of $\vec{\psi}$ (not normalized here) is defined by the heading direction (the derivative of eq. (7)), that is,

$$\vec{\psi} = \left[v_x \quad \frac{n\pi A}{a} \cos\left(\frac{n\pi t}{a}\right) \right], \quad (8)$$

and the direction of $\vec{\phi}$ is described by the direction from the point (x,y) to the point $(a,0)$, that is,

$$\vec{\phi} = \left[\frac{a-x}{\sqrt{(a-x)^2 + y^2}} \quad \frac{-y}{\sqrt{(a-x)^2 + y^2}} \right]. \quad (9)$$

Figure 9 illustrates the comparison between the baseline and multiple variations near the baseline path. In this situation, as the deviation away from the baseline increases, the correlation function decreases.

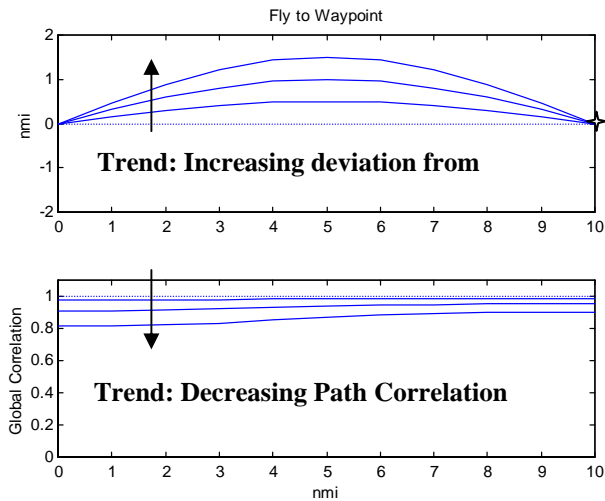


Figure 9. If a pilot is heading to a way point located at $(10,0)$ and the flight path veers away from the optimal path (the x -axis), then an increase in the local deviation from the baseline (top plot) causes a decrease in the correlation value (bottom plot).

Figure 10 illustrates the comparison between the baseline and variations in the number of turns. In this case, as the amount of turning increases from the baseline increases, the correlation function decreases.

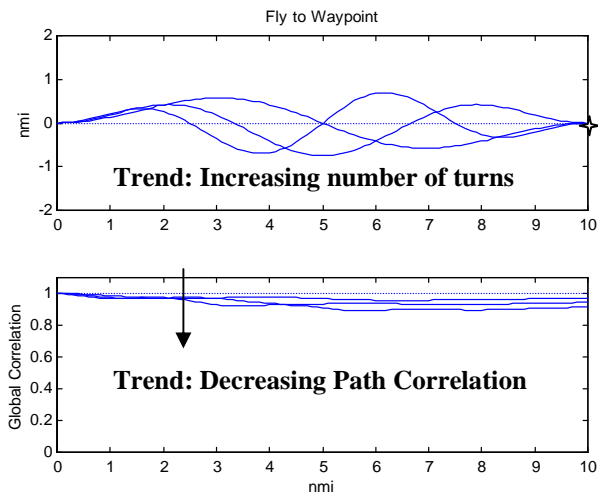


Figure 10. If a pilot is heading to a way point located at $(10,0)$ and the flight path turns from the optimal path (the x -axis), then an increase in the local deviation from the baseline (top plot) causes a decrease in the correlation value (bottom plot).

In general, local path variations around a baseline path cause a decrease in the correlation function. As expected, all of the resultant correlation values are

relatively close to, yet less than, the baseline value. When multiple cost criteria are considered, local variations in aircraft motion will act to reduce correlation values in a similar manner for all cost criteria – an observation based on our experience.

Competing Intent Models. In this scenario, four models of intent compete to explain an aircraft's motion:

- A. Fly-to-Way point TCP (Figure 11)
- B. Fly-to-Way point TCP+1 (Figure 12)
- C. Hold Heading (Figure 13)
- D. Avoid Hazardous Weather (Figure 14)

Given the directions that a pilot might fly to optimally follow the above travel criteria, the local and global path correlation measures are computed. Figure 15 illustrates two example flights being tracked. Figure 16 illustrates a time history of the competing intent models and the output from the intent inference algorithm.




Figure 11. Fly-to-Way point TCP optimal decisions.




Figure 12. Fly-to-Way point TCP+1 optimal decisions.




Figure 13. Hold heading optimal decisions.

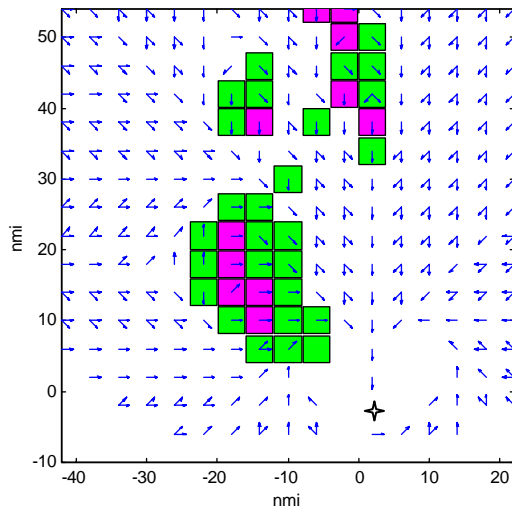


Figure 14. Avoid hazardous weather optimal decisions.

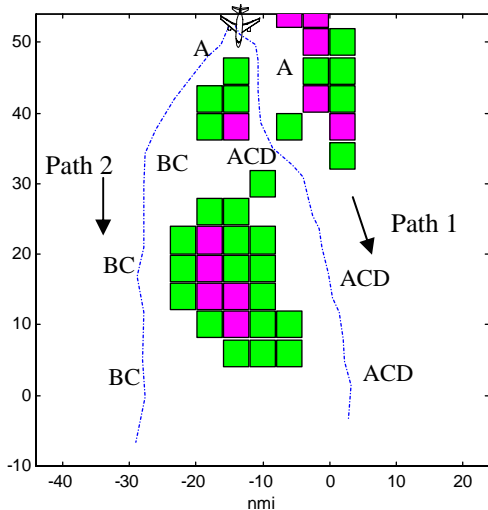


Figure 15. Two example flight paths and intent inference results.

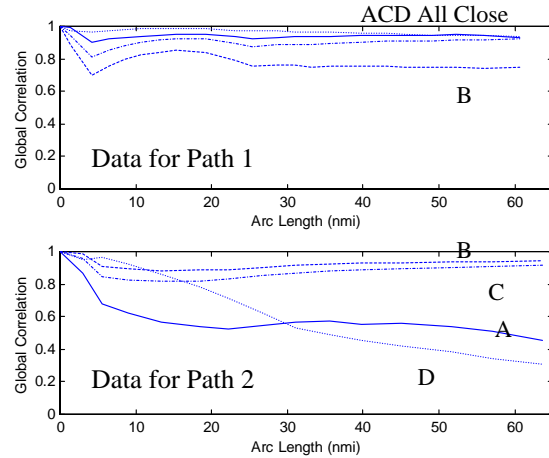


Figure 16. Time history of global correlation measures and the intent inference output.

Discussion

There are a few limitations to our work in this paper. While the results that we have generated to date look promising, we have not tested this approach with real flight data. We have also not modeled all the plausible cost criteria mentioned in Table 5, and anticipate that some of them may be difficult modeling tasks. Some models are appropriate for Free Flight, and some are more appropriate for today's operating conditions. Finally, for each intent model, there may be parameters which identify plausible intent models that increase the complexity of the solution approach. For instance, hold heading could be investigated with several headings (the heading that maximizes global correlation returned as the most likely intent).

Adding a Kalman Filter. The proposed intent inference method can be coupled with Kalman filter to produce state estimates and predict motion spanning near term (e.g., seconds to minutes) to far term (e.g., minutes to hours) prediction times. A Kalman filter by itself could not within reasonable accuracy predict intent beyond near-term time horizons, since the filter in its general form is not capable of modeling a knowledge base as diverse as the one described in this paper. However, the intent inference method described so far, while applicable to both near-term and far-term predictions, can be enhanced by a tracking filter for producing nearby aircraft state estimates, better (more accurate) near-term trajectory predictions for nearby aircraft, and a measure of uncertainty of these estimates.

The Kalman filter provides a convenient sequential fading window filter. It can provide the best estimate of the state variables needed for the correlation function, eq. (2), of the intent inference method even if, for instance, the state variables do not arrive at a continuous uninterrupted rate (e.g., during

lost transmissions of data from ADS-B or other input sources). Gelb [2] goes into detail as to how the tracking filter is used for estimating state variables. In [11], we present how our Kalman filter assists the intent inference algorithm with stable inputs, conversely, the intent inference method is used to adjust tracking filter parameters, such as the process noise. With these adjustments, a Kalman filter can produce better near-term trajectory predictions for nearby aircraft.

Conclusions

In this paper we present a method for inferring the intent of an aircraft given state information, domain knowledge (e.g., SIDs, STARs, weather data, SUA, etc.), and ADS-B broadcast intent information (e.g., TCP and TCP+1). Domain knowledge is used to build plausible models of intent, and through a correlation measure, the models of intent are ranked based on data that supports the plausible models of intent. The highest ranking models of intent from the set of plausible models constitute the best estimate of the intent of the aircraft. This technique can be used to verify the broadcast intent in the ADS-B message, or it can be used to infer an intent when there is no intent information in the ADS-B message. Additionally, the output of the intent inference algorithm can adjust the process noise parameters of a Kalman filter to result in better trajectory predictions of nearby aircraft.

Acknowledgments

This research was funded by NASA Langley Research Center under contract NAS2-98001 for the Advanced Air Transportation Technologies (AATT) program. The authors appreciate the program support of Mark Ballin and Sheila Conway at NASA Langley.

References

- [1] Bushman, J.B., Mitchell, C.M., Jones, P.M., and Rubin, K.S., "ALLY: An Operator's Associate for Cooperative Supervisory Control Systems," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. 23, No. 1, pp. 111-128, Jan./Feb., 1993.
- [2] Gelb, A., *Applied Optimal Estimation*, MIT Press, 1974.
- [3] Geddes, N.D., "A Model for Intent Interpretation for Multiple Agents with Conflicts," *IEEE Intern. Conf. on Systems, Man, and Cybernetics*, San Antonio, TX, Oct., 1994.
- [4] Hoshstrasser, B. and Geddes, N., "OPAL: Operator Intent Inferencing for Intelligent Operator Support Systems," *Proc. of the IJCAI-89 Workshop on Integrated Human-Machine Intelligence in Aerospace Systems*, Detroit, MI, Aug., 1989.
- [5] Kautz, H. and Allen, J., "Generalized Plan Recognition", *National Conf. on Artificial Intelligence*, Philadelphia, PA, Aug., 1986.
- [6] Krozel, J., *Intelligent Path Prediction for Vehicular Travel*, Ph.D. Thesis, Dept. of Aeronautics and Astronautics, Purdue University, W. Lafayette, IN, May, 1992.
- [7] Krozel, J., and Andrisani, D., "Predicting Mission Plans for an Observed Land Vehicle", *DARPA Symposium on Associate Technology*, Fairfax, VA, June, 1991.
- [8] Krozel, J., and Andrisani, D., "Intelligent Path Prediction for Vehicular Travel", *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. 23, No. 2, March, 1993.
- [9] Krozel, J., and Andrisani, D., "Intelligent ϵ -Optimal Path Prediction for Vehicular Travel", *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. 25, No. 2, Feb., 1995.
- [10] Krozel, J., Lee, C., and Mitchell, J.S.B., "Estimating Time of Arrival in Heavy Weather Conditions," *AIAA Guidance, Navigation, and Control Conf.*, Portland, OR, Aug., 1999.
- [11] Mueller, T. and Krozel, J., "ADS-B Intent Verification by Kalman filter and Intent Inferencing," *AIAA Guidance, Navigation, and Control Conf.*, Denver, CO, Aug., 2000.
- [12] Rouse, W., Geddes, N., and Curry, R., "An Architecture for Intelligent Inferences: Outline of an Approach to Supporting Operators of Complex Systems," *Human-Computer Interaction*, Vol. 3, 1988.
- [13] Rubin, K. S., Jones, P. M., and Mitchell, C. M., "OFMspert: Inference of Operator Intentions in Supervisory Control Using a Blackboard Architecture," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 18, No. 4, pp. 618-637, 1988.
- [14] RTCA, *Report of the RTCA Board of Directors' Select Committee on Free Flight*, RTCA, Inc., Washington, DC, Jan., 1995.
- [15] RTCA Special Committee-169, *Operations Concepts for Data Link Applications of Flight Information Services*, RTCA/DO-232, March, 1996.
- [16] RTCA Special Committee-186, *Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast (ADS-B)*, RTCA/DO-242, Feb., 1998.
- [17] Stengel, R.F., "Toward Intelligent Flight Control", *IEEE Systems, Man, and Cybernetics*, Vol. 23, No. 6, Nov./Dec., 1993.
- [18] Tambe, M. and Rosenbloom, P., "Event Tracking in Complex Multi-agent Environments", *4th Conf. on Computer Generated Forces and Behavioral Representations*, Orlando, FL, May, 1994.
- [19] Zhao, Y., C. Haissig, and M. Hoffman, "Analysis of Pilot Intent in Air Traffic Management," *Proc. of the American Control Conf*, Philadelphia, PA, June, 1998.