SAGA: Safety Analysis in General Aviation
Leveraging Digital Flight Data towards the Development of a Flight Performance Database Tool to improve General Aviation Safety

Josh Schimmel, BS
Department of Integrated Systems Engineering
The Ohio State University
Columbus, OH
Schimmel.5@osu.edu

Shawn Pruchnicki MS, ATP, CFI
Center for Aviation Studies
The Ohio State University
pruchnicki.4@osu.edu

Philip J. Smith, Ph.D
Department of Integrated Systems Engineering
The Ohio State University
Columbus, OH
smith.131@osu.edu

Seth Young, Ph.D., CFI
Center for Aviation Studies
The Ohio State University
Columbus, OH
young.1460@osu.edu

***Abstract* – The emerging prevalence of flight data recording technology in general aviation aircraft presents new analytical potential to enhance safety evaluation. In this paper we describe a prototype tool that not only examines individual flight parameters, but also provides information displays to help the user consider the relationships among related flight parameters. In addition, because the flight data for the broader population of general aviation pilots is stored in a single database, users of the tool can not only review performance data on their own flights, but also can compare their own performance with that of pilots in general.**

Keywords– database; flight safety; exceedances; use cases; analytics

I. Introduction

Part of the excellent safety record experienced by commercial flight operations is due to the ability to examine aggregate flight data in a search for overall trends, as well as to look at parameter exceedances for individual flights. Such analyses offer insight into flight operations that can be used to drive training curricula and to identify needed improvements in procedures or equipment design.

There currently exist two national aviation databases for sharing safety-critical information. Aimed at commercial aviation, the FAA's Aviation Safety Information Analysis and Sharing (ASIAS) system focuses on revealing trends based on aggregations of data drawn from a variety of sources (cf. asias.faa.gov) [1]. However, general aviation operations and data are not represented in the ASIAS data pool. Similarly, NASA maintains the Aviation Safety Reporting System (ASRS), where commercial and general aviation (GA) operators can anonymously submit inputs to a national database [2]. For both of these, the focus is specifically on aviation safety issues where sharing safety-critical information with the aviation community is beneficial.

General aviation flight operations to date have no aggregated data source like ASIAS. As a result, individual pilots are left to decide for themselves, based on anecdotal experience, which areas of their flying abilities need more practice or help from a Certified Flight Instructor (CFI). An analysis tool that provides effective access to aggregated data for GA flights would offer a means for pilots, operators, and regulators to examine GA flights and compare them to either known standards, be that regulatory or community driven, as well as the aggregated performance of the larger community.

This paper will describe recently available sources of general aviation flight performance data, applications or "use cases" for the application of such data to improve flight safety, and the design features of a prototype system, SAGA (Safety Analysis for General Aviation), designed to support analysis of such data.

II. Data

A. Data Collection

The Garmin G1000 integrated avionics unit is among the standard "glass cockpit" avionics packages found on modern general aviation aircraft, including common recreational and training aircraft, such as the Cessna 172. The G1000 contains a data-recording feature that captures 64 flight parameters per second during the time the avionics are powered on, which includes power-on, taxi, all phases of flight, and power down of the aircraft. Standard flight parameters are included (altitude, airspeed, latitude/longitude, etc.) as well as various engine parameters (cylinder head temperature, fuel flow, RPM, etc.) to name a few.

In addition, there are several other parameters, examples of which are listed in Table 1, that are not recorded but that, depending on analytical needs, might contribute significantly to a better understanding of the performance for a given flight. Collectively, these data are the basis for the discussion below regarding the design of a feedback tool for general aviation.

The flight training fleet at The Ohio State University in Columbus, Ohio includes 6 Cessna 172 aircraft with G1000 avionics installed. Data on over 1200 flights were extracted from these aircraft since July 2015 and stored in a prototype *SAGA* database. The process for data extraction was via the removal of a SIMM card from the specific aircraft to transfer the data to the test database.

Table 1: Parameters manually entered into the flight data, not automatically recorded by the G1000

Parameter	Description
Aircraft Code	N-number of the aircraft
Flight Rules	VFR or IFR
Pilot ID	Anonymous identification number assigned to each pilot
Training Flight	Identifies if the flight was used for training with a student
Pilot Skill Level	Private, Commercial, CFI, etc.
Departure/Arrival Airports and Runways	Airport information

B. Data Requirements

Platforms such as the G1000 sample data at sufficient rates (sample rate = 1 Hz) to support data requirements for many (but not all) potential use cases, including hundreds of flight parameters such as airspeed and altitude, but also aircraft centric measures such as radio frequencies and ambient and engine temperatures and pressures. Additional data such as GPS derived location increase the richness of the overall understanding of the context of data sampled. However, not all avionics platforms collect the same number of parameters as there are no agreed standards as to the specific parameters which must be collected at minimum. This lack of standardization will need to be considered in developing further analysis tools.

In addition, designers must determine how many and what types of manual entries should be supported (and determine whether there are automated ways of inferring these data so manual entries are not required). A method of incorporating weather data (temperature, wind vectors, etc.), for example, either manually or automatically from an external source, could add significant analytical benefit. Depending on the user and his/her individual goals (specific flights vs. specific operators vs. industry aggregate), it is important to understand the context of individual flights in order to identify the most meaningful data and patterns.

III. INTERFACE DESIGN & FUNCTIONAL REQUIREMENTS

A. Single Parameter Analysis

With over 1200 flights captured, the development of a prototype database interface raised intriguing questions as to how these data can be made useful for various user populations. In general terms, the database, the supporting algorithms, and the interface need to support use cases involving an individual pilot looking at his/her own flights, as well as use cases focusing on analyses aggregating results from multiple pilots.

In addition, the interface should have a major focus on the value added for individual pilots or flight operators to incentivize them to submit their own data. This added value needs to focus on the benefits of seeing how a particular flight, or a collection of flights managed by a particular operator, compares with other flights in order to help the user better understand whether his/her performance is atypical and therefore perhaps unsafe.

Initial analysis of the entered flight data focused on examining exceedances (already developed by academia, industry and the FAA) [3] that have been flagged by the prototype database interface. This involves examining individual flight parameters and identifying where that parameter crosses a certain threshold. In future designs, this threshold will need to be aircraft specific, based on manufacturer specifications, or may be user defined.

These analyses of exceedances have indicated a need to look closely at the raw data, and the associated calculations for determining exceedances, in order to:

- Ensure that the data collection system is always providing valid data.
- Determine whether richer, more complex metrics need to be used to more precisely and accurately flag exceedances. This is critical for defining requirements, as a system that routinely produces false alarms is not likely to be adopted for use by individual pilots, flight operations, or system analysts.

B. Multi-Dimenstional Analysis

It is important for the designer of a flight repository database and associated interface to understand that a simple flight parameter exceedance may not fully support a pilot in assessing whether there were any safety concerns associated with a particular flight. An exceedance of some set threshold can be useful for alerting the pilot or analyst of a potential concern, but by itself may not provide sufficient information to determine whether this concern is valid. For example, in Figure 1 there appears to be a significant pitch angle that triggered an exceedance. (30 degrees of pitch-up has already been determined to be an exceedance by default.) However, to determine whether this value is really a concern, it is necessary to understand the context in which it occurred, including phase of flight, the concurrent value of other parameters, or what could be considered reasonable performance for an instructional flight, should that be the case.

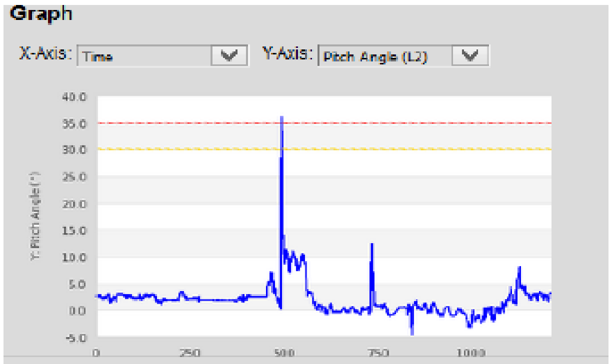

Figure 1: Graph of pitch angle over time for a sample flight

As a further illustration, a low airspeed exceedance at 8,000 ft above the ground (AGL) is probably much less of a concern then one at 500 AGL. Without additional data to help understand the context of the operation in question, a simple numerical exceedance presented to the user does not support a sufficiently clear assessment of safety. In the example above, once we manually examined the location of the flight path for this instructional flight (from GPS coordinate data drawn in Figure 2 below) in conjunction with at a graph of airspeed (shown in Figure 3), we were able to determine that this was a training flight, and the pilot was likely practicing stalls.

Figure 2: GPS-driven flight path for a sample flight

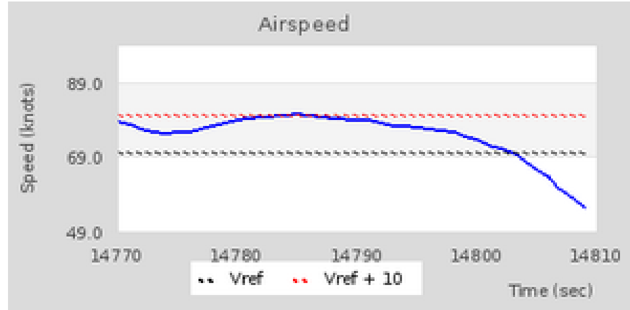

Figure 3: Graph of airspeed over time for a sample flight

This example thus emphasizes the need to consider not only a single value exceedance, but to consider its relationship to other parameters. In some cases, this richer analysis can be supported by more context-sensitive metrics based on combinations of the values of several parameters. In other cases, it may be desirable to develop integrated displays that help the user to visualize these relationships and make a judgment regarding safety implications (for example linking the displays so the user can see the exact point where the flight is on its route when the airspeed is flagged as an exceedance.)

Additionally, in this example, despite being potentially less dangerous at higher altitudes, cases of low airspeeds may still indicate a problem with certain flight operations (such as charter flights) but may be quite normal with others (primary or recurring flight training). Users who are gathering statistics to look at aggregations of flights should be able to filter out or isolate certain types of flight operations in order to make these statistics meaningful. In the example above, a manual entry designating this as a training flight would enable the user to select or deselect training flights. Such capabilities are critical if analyzing diverse and potentially large volumes of data that can be found within the general aviation community.

C. Time-Dependent Analysis

Observing a parameter (or multiple parameters) at a single point in time lacks the contextual support for analyzing flight behavior. An effective examination often must not only show users which flights have exceedances for any given parameter, but also provide the ability to see trends, patterns and the relative stability of a parameter as an indicator of flight performance and risk.

To the extent possible, an effective SAGA interface should be able to automatically consider such context effects and display results appropriately so as to minimize false alarms (or misses). However, there could be important relationships that the user is aware of, but are not captured in an exceedance. Consequently, the user should also be able to select information to display in an integrated fashion (with a matched congruent data collection time scale). Figure 4 illustrates the value of this capability for analyzing approach performance as compared to industry-accepted values for a stabilized approach. In terms of requirements, it is useful to let the user scroll along the horizontal (time) axis of each graph, and view the corresponding values for other specific parameters (and their exceedance ranges) to help the user develop a richer understanding of what is happening at any moment in time.

In this example, the area of interest during this approach is where the vertical speed is in excess of the threshold relative to stabilized approach criteria. However, if only this single parameter's exceedance is flagged and displayed in isolation (without airspeed and altitude and their corresponding exceedance values), it is not possible to determine how egregious this is exceedance actually is. Thus, it is critical for the design to be sensitive to the consideration that exceedances for individual parameters do not necessarily tell the entire story when examining either single flights or aggregates of large numbers of flights.

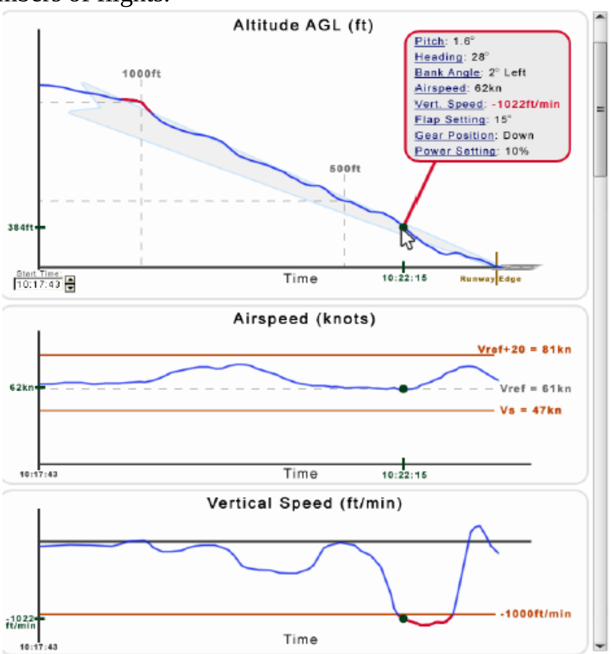

Figure 4: Corresponding displays for multiple flight parameters aligned in time

In addition, the interface and its underlying functionality need to be designed to consider the duration of an exceedance for either a very short duration or for rapid variations in the parameters in order to assess significance.

For example, consider the vertical speed illustration shown in Figure 5 below. At no point does an exceedance of a vertical speed of greater than 1,000 ft/min exist in the data over this time period. If the interface were designed to simply find exceedances, this descent/approach would never be flagged in the data stream. It is not the presence or lack of exceedance that is interesting here, but rather the variation in the rate is potentially more telling. While this flight does not violate the general descent rate exceedance, one can clearly identify visually that this descent is far from stable. By combining this flight parameter with other relative parameters (as determined in the design phase or as selected by the users), a more complete understanding of the context is possible.

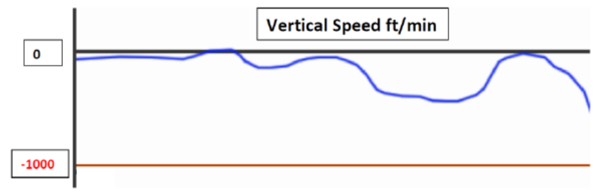

Figure 5: Graph of vertical speed over time during a sample flight's approach

Thus, one challenge in specifying design requirements is to define what constitutes significant variation or instability, and over what period. This is no small task for the designer and should be considered for every parameter, not only in isolation but also in a constellation with others.

D. Benchmarking

Thus, one challenge is designing effective algorithms and displays that help users determine when their performances are significantly different from those of others, and to help them to understand whether such differences warrant serious consideration about how to improve performance in the future. Additional value can be generated by comparing the performance of a single pilot or a particular flight operation with a larger population of flights.

Consider an aircraft operator interested in studying his/her conformance to a certain noise abatement procedure in comparison to other aircraft operating the same procedure. Figure 6 below displays the power setting of four flights, with the user's flight highlighted red. All flights were flown with the same aircraft type.

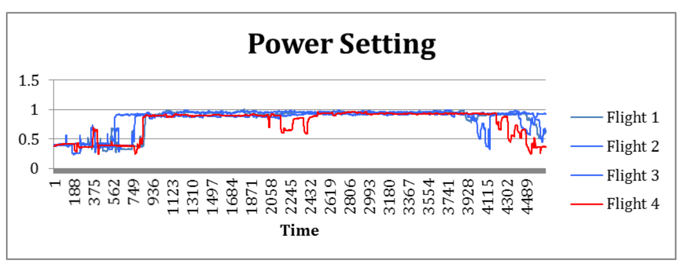

Figure 6: Power settings for 4 flights aligned in congruent time

Any deviations from the procedural power settings would be evident, and the user could modify his/her practices accordingly based on these results. This same type of display could also be beneficial for a systems analyst to assess conformance and recommend necessary changes. This type of comparison ability, over multiple flights, makes it easier for users to identify trends and patterns in flight behavior.

IV. USE CASES

To aid in determining interface design and functional requirements, a variety of use cases were developed. [4]. Considering this diverse user population, the design requirements were developed using an outside-in approach, using storyboarding to initially explore interaction design concepts. Cognitive walkthroughs were then performed to assess the usability of the design, iteratively making changes as needed [5, 6]. As this interface is still in development, not all requirements and features have been implemented to date. However, this section details an example of how a user would use the interface to complete her/her own analysis given the interface in its current state.

A. Users and Goals

Consider a single general aviation pilot who is familiar with general aviation terminology, is using a computer or tablet and wants to review his/her flights over the past few months.

B. Searching for Flights of Interest

The interface is able to generate three separate reports based on the type of search desired. To filter results and create a more defined search, various fields can be specified, such as:

- Departure/Arrival Airports
- Departure/Arrival Date & Time Ranges
- Pilot IDs
- Aircraft Type & N-numbers
- Aircraft Equipage

After the search is run, the user may go back to modify the search options, or save the search query if the user routinely performs the same search. This query is displayed in the top box, and updates automatically as fields are modified below.

For this example, the user enters a date range between 5/12/2015 and 11/11/2015 to collect all flights in that range. His/her anonomyous Pilot ID is entered [3], but all other fields are left blank. See Figure 7 for the initial search page of the interface.

C. Results

Once the search was performed, 87 flights were found, listed in the bottom half of the screen in chronological order (see Figure 8). This order can be changed by clicking one of the header categories at the top of the list.

The top half of the screen is devoted to exceedances. The number of flights listed under “high” or “low” indicates the exceedance threshold for that parameter was surpassed during those flights, while a “caution” indicates a parameter was close to the threshold, but never surpassed it. Selecting the blue link displays a list of those flights for a more detaield examination and from there any individual flight can be examined upon selection (See Figure 9).

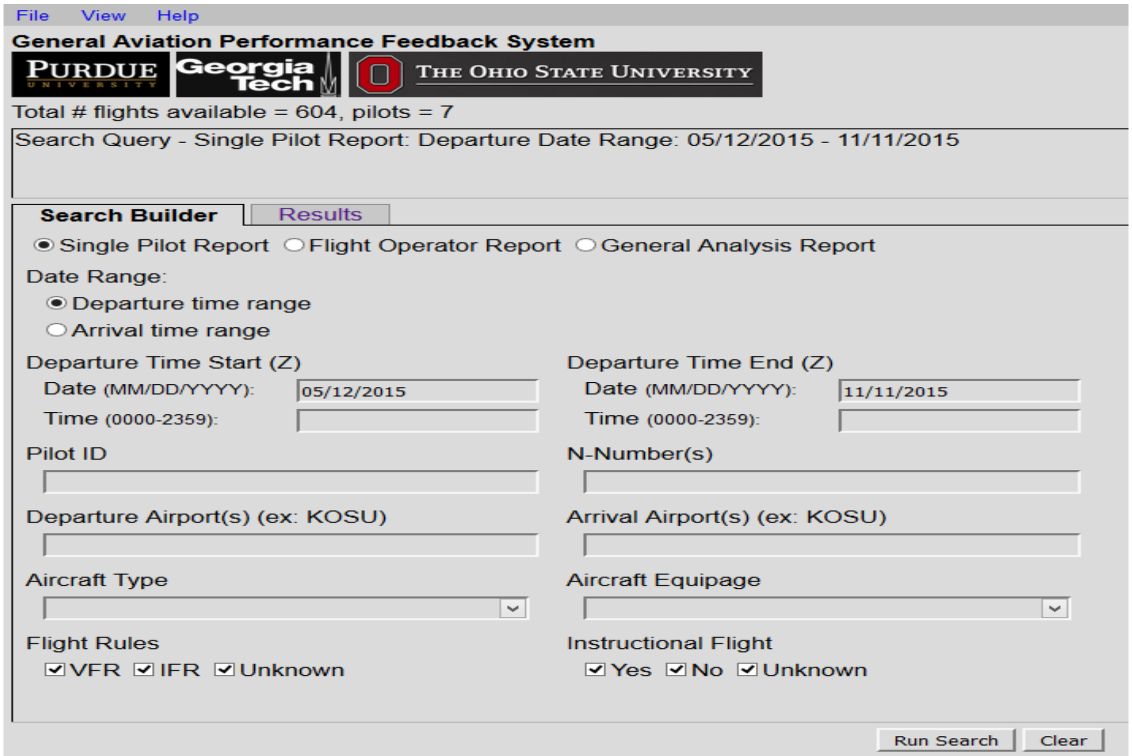

Figure 7: SAGA Flight Data Search Page

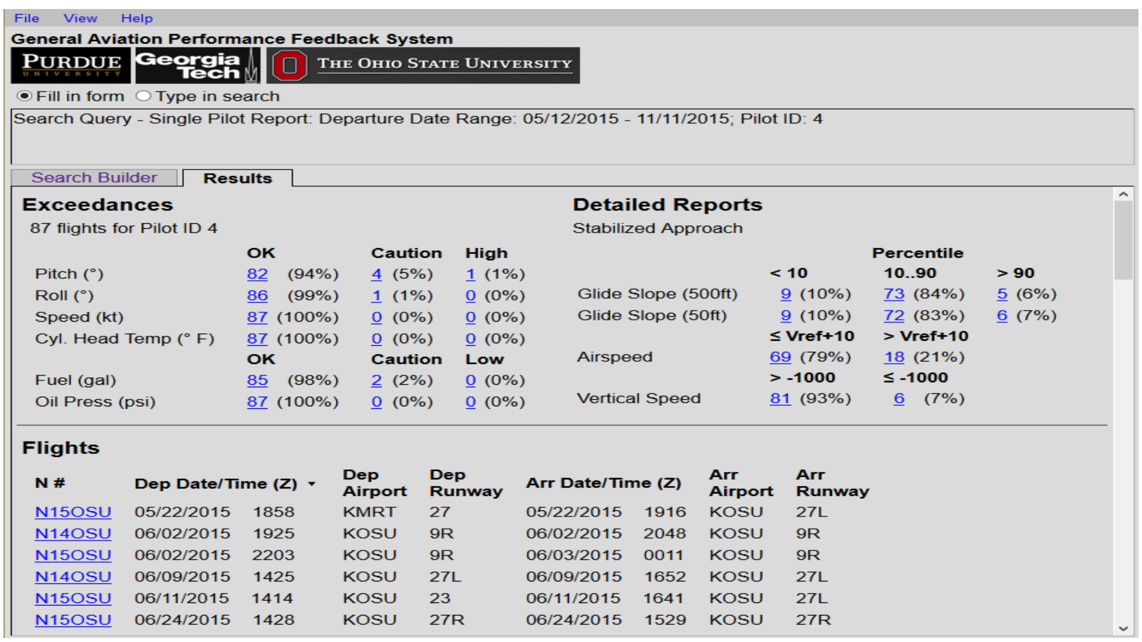

Figure 8: Initial Results and Summary

Pitch = Caution						
N #	Dep Date/Time (Z)	Dep Airport	Dep Runway	Arr Date/Time (Z)	Arr Airport	Arr Runway
N14OSU	08/27/2015 1540	KOSU	9R	08/27/2015 1627	KOSU	9R
N14OSU	09/01/2015 1434	KOSU	27L	09/01/2015 1532	KOSU	27L
N14OSU	09/15/2015 1250	KOSU	27L	09/15/2015 1400	KOSU	9R
N2OSU	10/07/2015 1156	KOSU	27L	10/07/2015 1252	KOSU	27L

Figure 9: “Caution” flight results for pitch

D. Detailed Flight Analysis

Clicking on one of the N-numbers in this list of flights (see Figure 9) displays a window containing detailed infomation about that particular flight, shown in Figure 10. A map display draws the flight path, and graphs display individual flight parameters. While time is the default on the X axis, users are able to modify the parameters on the X and Y axes, depending on the parameters of interest. Up to 6 different graphs can be displayed simultaneously.

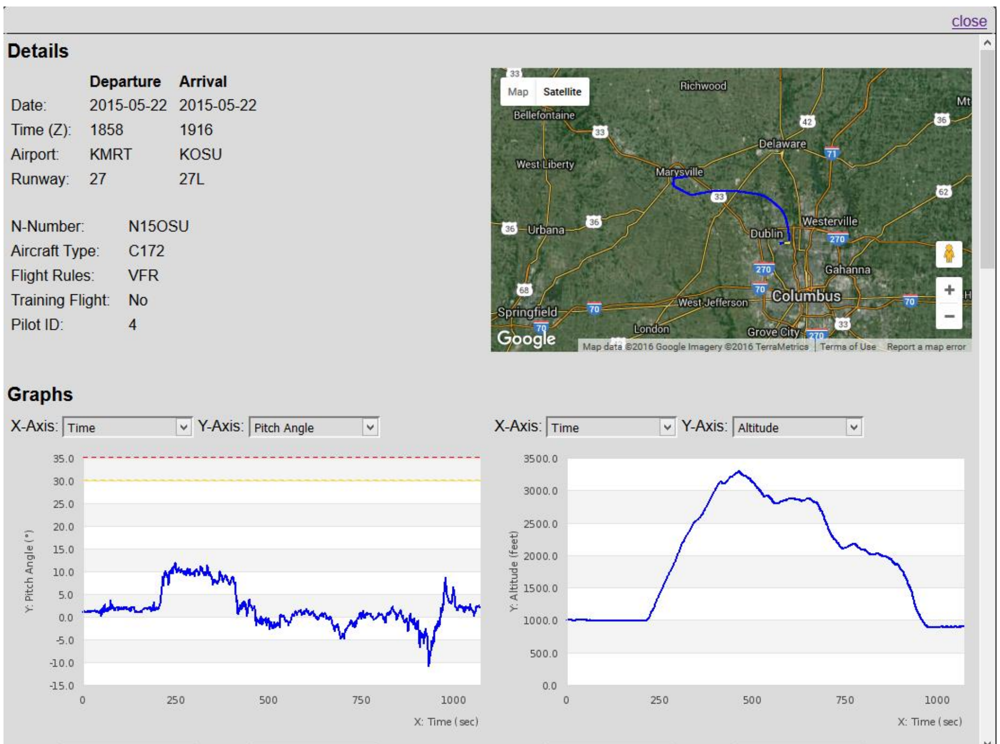

Figure 10: Details for an individual flight

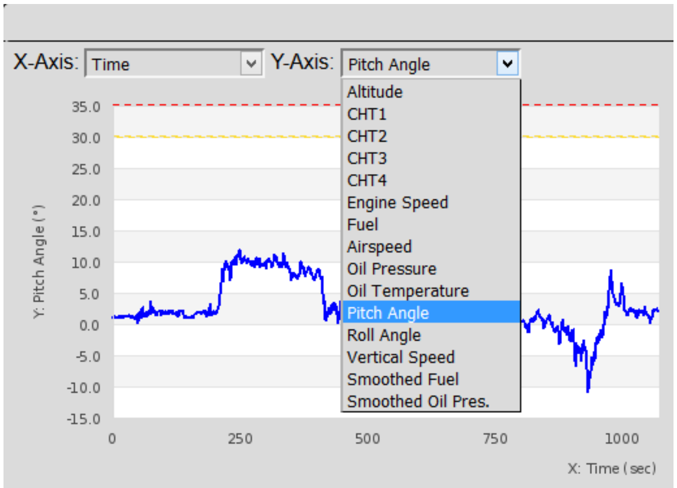

Figure 11: Graph parameter options

Using these interface features, the user is able to examine any flights of interest and decide where additional training or review would be beneficial. This type of self analysis could be very valuable for Flight Instructors to review as they are crafting a required flight review for a pilot.

E. Additional Use Cases

As additional algorithms are built into this database, more advanced statistics and multi-dimensional analysis will be added. Furthermore, additional detailed use cases are being created to further guide the development of the interface. Figure 12 lists some of these use cases.

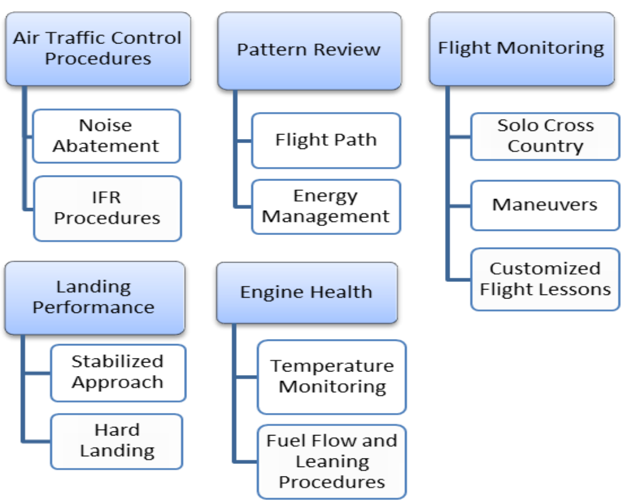

Figure 12: Additional Use Cases

V. Conclusion

By exploiting new data collection capabilities available in general aviation aircraft, an exciting opportunity emerges to conduct a new type of safety analysis. The continued development of a SAGA database, multi-dimensional algorithms, and associated interface will help explore and evaluate the design of a system that allows pilots to easily conduct their own safety investigations; not only analyzing their own flights, but also benchmarking their flights against the larger population. This same opportunity can be expanded to larger GA operators, airport authorities, or aviation regulators, looking to make policy or procedural decisions based on current operational performance..

Acknowledgments

The authors would like to thank researchers at Purdue and Georgia Tech for their contributions to the SAGA project, as well as the FAA General Aviation Center for Excellence (PEGASAS), for their support and funding of this research. we further thank Dustin Johnson for all of his work on this project building the database and interface. Finally, we thank The Ohio State University Flight School for its contribution of flight data to help develop the initial versions of a workable database.

This work was sponsored by the Federal Aviation Administration under Project 5 of the FAA Center of Excellence for General Aviation, PEGASAS (Partnership to Enhance General Aviation Safety Accessibility and Sustainability), via grant No. 12-C-GA-OSU to The Ohio State University. The authors are grateful to FAA project technical monitor Mr. Michael Vu, and to project partners at Purdue University and the Georgia Institute of Technology. The information in this research does not constitute FAA Flight Standards or FAA Aircraft Certification policy.

References

[1] Aviation Safety Information Analysis and Sharing (ASIAS). Retrieved from http://www.asias.faa.gov/pls/apex/f?p=100:1:

[2] Aviation Safety Reporting System. Retrieved from http://asrs.arc.nasa.gov/

[3] Clachar, S. A. (2015). Identifying and analyzing atypical flights by using supervised and unsupervised approaches. *Transportation Research Record: Journal of the Transportation Research Board*, (2471), 10-18.

[4] Wickens, C. D., Lee, J. D., Liu, Y., & Gordon-Becker, S. (1998). Introduction to human factors engineering.

[5] Lewis, C. & Wharton, C. (1997). Cognitive walkthroughs. In M. Helander, T. Landauer and P. Prabhu (Eds.), *Handbook of Human-Computer Interaction, 2nd Edition*. Amsterdam: Elsevier, 717-731.

[6] Smith, P.J., Stone, R.B. & Spencer, A. (2006). Design as a prediction task: Applying cognitive psychology to system development. In W. Marras and W. Karwowski (eds.), *Handbook of Industrial Ergonomics, 2nd Edition*. New York: Marcel Dekker, Inc, 24-1 -24-18.