The 2016 AAAI Fall Symposium Series:
Cognitive Assistance in Government and Public Sector Applications
Technical Report FS-16-02

Digital Copilot: Cognitive Assistance for Pilots

Steven L. Estes, Kevin J. Burns, John R. Helleberg,
Kevin M. Long, Matthew E. Pollack, Jeffrey L. Stein
The MITRE Corporation
7515 Colshire Dr.
McLean, VA 22012

Abstract

Over the past several years, pilot-oriented mobile applications have seen widespread adoption among recreational pilots. Pilots have reported they provide significant workload savings by eliminating the need to manage paper charts, manuals, and checklists in the cockpit. The pilot, nonetheless, still must go looking for the information when it is required, increasing accident risk by diverting attention away from control of the aircraft. In this paper, we provide an overview of a cognitive assistant that determines when information is required based on flight context and automatically provides it to the pilot at the appropriate time. In addition to an overview of the concept, a recent evaluation is discussed alongside future plans to evaluate the safety of the Digital Copilot.

Introduction

In an old aviation anecdote, a commercial pilot, accustomed to flying as part of a two-person crew in a multi-engine jet, decides to fly a solo General Aviation (GA) flight in a small, single-engine piston aircraft for fun. Prior to the flight, he declares an emergency with the Flight Service Station (FSS). The briefer, somewhat incredulous, asks why the pilot is declaring an emergency before the plane has even left the ground. The pilot says "I am down to one engine, one pilot, and one radio; in my line of work that is an emergency."

It is often said that flying is the safest form of transportation, and for commercial air transport that is absolutely true. Recreational flight is, however, a wholly different animal. Failure Analysis Associates (1993) estimates that, while the risk of dying in a commercial aviation accident was 0.2 per million flight hours, the recreational flight rate was 15.6 fatalities per million hours (nearly double the motorcycle rate).

Equipment that is less reliable and lacks redundancy plays a large role in recreational flight risk, but it is not the only contributor; flying solo (i.e., without a trained copilot) is another important factor. According to Robert E. Breiling and Associates (2010), the accident rate for turbine aircraft certified for single-pilot operations is 3.4 times higher than the rate for aircraft requiring two pilots. The fatal accident rate was 13 times higher.

There's no mystery as to why this is. Flying an aircraft by yourself is demanding. Among other things it requires frequent mental calculations, communications with air traffic control, accessing paper or electronic charts and manuals, and scanning for traffic. On top of all that, the pilot must fly the airplane, controlling for all four dimensions. This last and most critical task can be deceptively difficult as an input on one dimension can impact all other dimensions. Consider that when banking an aircraft into a turn, the pilot must raise the nose to avoid losing altitude, which in turn causes a loss of speed unless the pilot also increases the power. So, the pilot must adjust the aircraft power while using two different control inputs (stick and rudder) to effect the bank and, at the same time, monitor the airspeed, bank angle, and altitude.

Those in commercial aviation have long recognized the safety benefits of having two pilots in the cockpit, and have formalized each pilot's role under the concept of Crew Resource Management (CRM). CRM works under the principles that each pilot has a unique role to play during the flight and each should crosscheck the other (Helmreich, Merritt, & Wilhelm 1999). In other words, the pilots act as living, breathing cognitive assistants to one another. CRM has been a hugely successful safety innovation in commercial aviation, and its success invites the question: can we use existing technologies to bring some components of CRM to pilots flying without a copilot? More to the point, can we build a Digital Copilot?

Design

Before deciding whether it was possible to build a Digital Copilot, it was necessary to define what the system must do. In CRM, the second pilot (often referred to as the Pilot Monitoring) has a wide variety of responsibilities, some of them more amenable to a digital cognitive assistant than others. For instance, the second pilot is typically responsible for communications with Air Traffic Control. While allocating that responsibility to a Digital Copilot is a technical possibility, the implementation would require a near zero failure rate (and, would surely be met with stiff regulatory resistance). Other constraints existed as well. Most recreational aircraft, for example, do not have an internet connection nor do they have readily available connections to the avionics or other aircraft systems.

With tradeoffs between feasibility, level of effort, and effectiveness in mind, a set of design guides and guardrails were constructed to direct development. The first clear guide coming from design team discussions was the principal of providing the right information at the right time. By using context to anticipate a pilot's need for specific information, the Digital Copilot could perform one of the essential roles of the Pilot Monitoring: retrieving and managing situational information based on operational context.

Any functionality developed was also to rest safely between the guardrails of reliability and recoverability. If a feature could not be developed so that it worked reliably, it was not to be considered for inclusion. Similarly, if loss of a feature would put the pilot in a situation from which he or she might not recover, it was eliminated. Reliability and recoverability are both principles that are widely applied when evaluating automation in aviation (Sheridan and Parasuraman 2006) and, as cognitive assistance is an intelligent form of automation, those principles should be applied to the design of a Digital Copilot as well.

Feature Types

Features for consideration were collected by reviewing the tasks performed by the Pilot Monitoring, taking flights in which we observed a recreational pilot flying solo, and hosting a workshop in which we asked pilots with a wide range of experience to suggest tasks that a Digital Copilot may usefully perform. Features were then downselected by applying the guides and guardrails as well as assessing technical feasibility and cost. What was left was a set of features that largely focused on reduction of workload through information management. These information management features can be grouped into categories of on-demand information and contextual notifications.

On-Demand Assistance

Using a speech recognition based interface, the pilot can request information directly. For example, the pilot can ask "Will the tower be open" to which the Digital Copilot will reply either yes or no. The reply is based on the Digital Copilot's estimate of time to destination and the air traffic control tower schedule (stored on the device). As a result, the pilot does not need to pull the appropriate book from the flight bag to get the tower closing time nor is it necessary to calculate the time to destination and then determine whether the arrival will occur before the tower closing time. Workload is offloaded with simple cognitive assistance. Examples of other information that the pilot can request include: information about the approach altitude and direction, basic position information relative to the destination, and checklists, which the Digital Copilot can also read to the pilot.

Assistance via Contextual Notifications

The heart of a Digital Copilot concept is its ability to infer pilot intent. Accurate and reliable inferences allow the Digital Copilot to automatically deliver relevant information to the pilot based on context. Currently, the Digital Copilot is able to infer the destination airport, the phase of flight, and phase of the approach. Using these inferences, the Digital Copilot determines when to provide some types of information. For example, prior to contacting the air traffic control tower on arrival, the pilot must have the current weather information. That information is broadcast on a unique frequency at each airport. So, the pilot must remember to look up the frequency, tune it in, and listen to the current weather broadcast at an appropriate point prior to contacting the tower. The Digital Copilot helps the pilot complete this task by automatically presenting the correct frequency to the pilot at the onset of a particular inferred phase of flight. This assists the pilot both by eliminating the workload of having to lookup the frequency as well as by reminding the pilot to complete the task.

Evaluation

The current set of features were implemented in Swift for iOS and deployed on an Apple iPad for testing. In all, about 25 different cognitive assistance features were developed, driven by a set of 10 algorithms. A speech recognition system running locally on the device (most recreational aircraft do not have internet access) was deployed for on-demand information requests. Visual and aural interfaces were developed for contextual notifications. To determine whether the cognitive assistance features reduced pilot workload, a longitudinal evaluation was conducted. The evaluation goal was to assess the usability

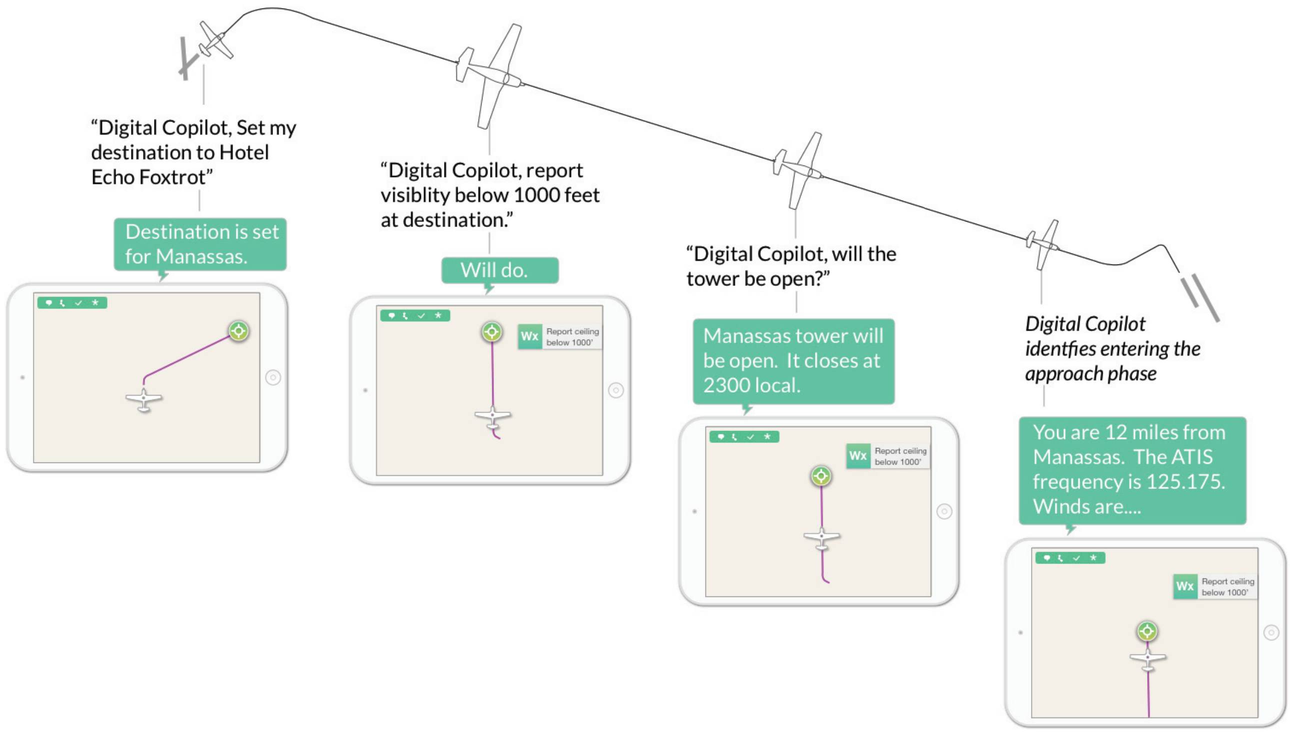

Figure 1: Typical Interaction with Digital Copilot During Evaluation.

and utility of the assistance provided by the Digital Copilot.

Fourteen recreational pilots participated in the study during which they flew a series of flights in a modified FRASCA flight simulator while using the Digital Copilot. An example of some of the interactions between the pilot and Digital Copilot can be seen in the Figure 1. Pilots returned a month after the first simulation activity to complete a second. In the second simulated set of flights they used a version of the Digital Copilot that had been updated based on comments received during the first simulation activity.

In exit surveys all fourteen pilots agreed that the Digital Copilot reduced workload and had the potential to improve safety. When asked whether the Digital Copilot could potentially be distracting, one of the fourteen pilots replied in the affirmative and two were unsure.

Overall, pilots found that the Digital Copilot - using a combination of if-then, fuzzy, and Bayesian logic to enable cognitive assistance for pilots - reduces workload and improves safety. While these findings suggest the human factors methods used to identify and select features were successful, there remains a need for quantifying the safety benefits that the system provides. Ultimately, utility must be measured by a user's ability to understand flight situations and take appropriate action.

Measuring Safety

In parallel with current system design efforts, we are developing methods for formally evaluating the system's safety. Our approach combines quantitative methods at three levels of analysis, each addressing a different level of utility in the hierarchy illustrated in Figure 2.

At the topmost level we are using methods of Probabilistic Risk Assessment (PRA), such as event trees and fault trees, to model critical safety functions and potential failure modes (Modarres 2008). These methods are used to compute fatal and non-fatal accident frequencies, which represent overall measures of operational safety.

At the middle level we are using methods of Human Reliability Analysis (HRA), such as time-reliability correlations, to model pilot errors that contribute to accidents (Bell and Holroyd 2009). The methods compute error probabilities as a function of the time windows for action along with other performance shaping factors such as stress and workload. These error probabilities are needed as input to the PRA models that quantify accident sequences as combinations of aircraft system failures and human pilot errors.

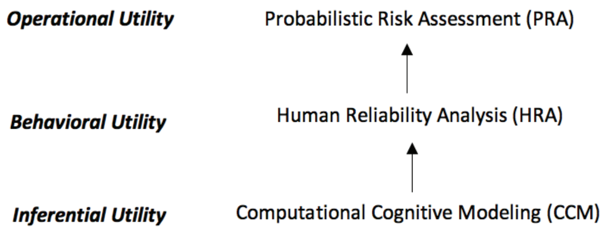

Figure 2: Process for Assessing Digital Copilot Utility.

Finally, at the bottom level we are using Computational Cognitive Modeling (CCM) to model errors in pilot situational awareness and decision making, which are the functions for which Digital Copilot is intended to provide assistance. Failure probabilities for these cognitive functions are needed as input to the HRA models of pilot errors, which in turn are needed as input to the PRA models of accident sequences.

Starting from the bottom of the hierarchy, our approach starts by quantifying the inferential utility of the system to the pilot; as a basis for quantifying the behavioral utility of the pilot to the mission; and as a basis for quantifying the operational utility of the human-system teaming needed to achieve flight safety. Clearly the bottom level of inferential utility is crucial, and yet methods for analyzing utility at this level are currently much less mature than existing methods for PRA and HRA used to quantify risk in aviation and other hazardous industries.

To advance the state of analysis at the bottom level, our research is leveraging a Bayesian-probabilistic framework presented at the 2015 AAAI Symposium on Cognitive Assistance in Government and Public Sector Applications. The framework, dubbed HELP (Hypotheses, Evidence, Likelihoods, and Priors and Posteriors), offers a principled structure for analyzing how, and how well, a human reasons to the hypothesis that most likely explains available evidence (Burns 2015). This framework can be used to quantify the probability that a pilot will correctly diagnose a given situation, as a function of the information that the pilot obtains manually or automatically with or without the Digital Copilot.

Our research in this area has three objectives: develop formal methods that can be used to quantify the operational safety of the Digital Copilot; quantify the safety benefits of potential changes to the system, thereby informing ongoing development efforts; and to identify how methods of CCM, HRA, and PRA can be extended beyond evaluation of Digital Copilot to quantify the inferential, behavioral, and operational utilities of cognitive assistance in domains other than single-pilot aviation.

Summary

General Aviation single-pilot operations offer an excellent applied problem for cognitive assistants, but success in this area requires a thorough knowledge of the task at hand. The flight deck, whether a small recreational aircraft or large commercial jet, already demands the pilot's attentional resources in the form of aural alerts, communications, displays, paper, and the like. A cognitive assistant must lessen those demands, not increase them. As such, applying basic human factors principles to the system design and, just as importantly, evaluating the overall utility is critical. The system should do no harm. Developing methods to ensure this is the case is an important component of not just the Digital Copilot, but any cognitive assistant.

References

Bell, J. and Holroyd, J. 2009. *Review of human reliability assessment methods*. Derbyshire, UK: HSE Books.

Burns, K. 2015. Bayesian HELP: Assisting Inferences in All-Source Intelligence. In Proceedings of the 2015 AAAI Fall Symposium Series. Arlington, VA: 2015 AAAI Fall Symposium Series.

Failure Analysis Associates Inc. 1993. Comparative Risk of Different Activities. *Design News* October 4, 1993.

Helmreich, R., Merritt, A. and Wilhelm, J. 1999. The Evolution of Crew Resource Management Training in Commercial Aviation. *The International Journal of Aviation Psychology*. 9(1):19-32.

Modarres, M. 2008. *Probabilistic Risk Assessment*. In ed. K. B. Misra, 669-718. Handbook of Performability Engineering. London: Springer.

Robert E. Breiling and Associates. 2010. *Business Turbine Aircraft Accident Review*. Robert E. Breiling and Associates, Boca Raton, FL.

Sheridan, T., and Parasuraman, R. 2006. *Human-automation interaction*. In ed. R. S. Nickerson, 89-129. Reviews of human factors and ergonomics. Santa Monica, CA: Human Factors and Ergonomics Society.