Visualization of Pilot-Automation Interaction

Todd J. Callantine and Barry W. Crane
San Jose State University/NASA Ames Research Center
MS 262-4, Moffett Field, CA 94035-1000
[tcallantine, bcrane]@mail.arc.nasa.gov

ABSTRACT

This paper presents a methodology and interface for visualizing pilot-automation interaction. Extensions to the Crew Activity Tracking System (CATS) enable it to associate prompting events with pilot activities and subsequent automated aircraft behaviors, and to display the associations graphically. The depictions provide qualitative and quantitative information in a form that supports both rapid detection of high-level interaction patterns and detailed analysis of salient events, pilot activities, and aircraft behaviors. The paper provides examples of data visualizations from a recent NASA simulator study and discusses their value for analysis.

Keywords

Visualization, pilot-automation interaction, aviation

INTRODUCTION

Understanding pilot-automation interaction is central to the development of advanced air traffic management concepts. New operational procedures rely on the use of aircraft automation as a complement to ground-based Air Traffic Control (ATC) automation (e.g., [5]). Of particular interest is whether and how pilots properly implement procedures that require the use of automation. This requires assessment of the modes used, and the timing of their setup and engagement with respect to the flight context. The timing and content of ATC clearances and the location of the aircraft in the airspace must also be considered. Using conventional techniques, it can be tedious to extract data from multiple sources and quickly assimilate it in a form that supports such analyses.

To address this problem, this research applies information visualization concepts to the analysis of pilot-automation interaction. Information visualization concepts are attractive because they support rapid exploration of complex data, as well as detailed examination of areas of interest (e.g., [9]). For example, for a new procedure that requires the use of high-level aircraft automation to accurately fly a projected trajectory, researchers could quickly identify the undesirable trend for crews to revert to lower levels of automation. They could then examine the characteristics of the flight context (e.g., ATC clearance amendments, information requirements, workload, etc.) surrounding each instance and address problems.

The visualization implementation presented here utilizes the Crew Activity Tracking System (CATS) [2]. CATS has been applied to the analysis of procedural deviations [3] and, more recently, has incorporated visualization techniques to help understand the operational context in which activities take place [1]. This foundation made CATS the natural choice for the framework upon which to base the visualization system. Both CATS and the visualization system are implemented in the Java™ language.

The remainder of the paper is organized into four parts. It first describes the visualization methodology, selection of information to include, and how CATS constructs the visualization. Second, it describes the features of the visualization interface. Third, it demonstrates by example how the visualization can be used to detect key trends and examine important factors in detail. Finally, the paper discusses the results of the research and future directions.

VISUALIZATION METHODOLOGY

The methodology for visualizing pilot-automation interaction records the flow of salient events, pilot activities, and aircraft behaviors, together with the automation modes used. The methodology is unique in that it also constructs a graphical depiction of the relationship between a pilot activity and the aircraft behavior it enables. This helps elucidate important characteristics of automation, viz., dependencies between modes and behaviors and latencies between the time when a mode is configured for use and the time when the automation actually performs the behavior. Visualizations are selectable as either sequential, where events and activities are simply depicted in the order in which they occur, or temporal, where they are plotted according to their actual time of occurrence. Temporal visualizations highlight clusters of activities and again emphasize workload shifts between control and monitoring activities. In addition, the visualizations relate pilot-automation interactions to the state and location of the aircraft in the airspace.

Implementation Issues

Implementing the methodology to allow inspection of pilot-automation interaction in a particular operational environment requires the researcher(s) to identify the prompting events of interest, along with the pilot activities and aircraft behaviors they wish to visualize, and ensure that all are detectable from data. The general visualization methodology allows considerable latitude in this process.

The visualizations described below, for example, extend the general methodology to highlight distinctions between salient classes of pilot activities or aircraft behaviors. One such distinction is between the dimension of flight (e.g., lateral, vertical, or speed control). Another is between strategic and tactical pilot activities, which are often interspersed. For example, by programming information into the Flight Management Computer (FMC) and engaging high-level automated modes, pilots control the lateral and vertical flight of the aircraft strategically. At the same time, however, they may engage in tactical speed control by manually overriding FMC speed targets. Aircraft behaviors may similarly be assigned to distinguishing classes, such as latent and immediate. Latent behaviors may be enabled by a strategic pilot activity, but may not occur until the aircraft attains a particular state. By contrast, other behaviors may occur immediately as a result of a tactical pilot activity.

Specifying the information to be visualized generally entails constructing a table that represents the pilot activities and aircraft behaviors at appropriate levels of abstraction; if appropriate, the information might further be divided according to the dimensions of control and classes of activities and behaviors (Table 1). As Table 1 suggests, there may be instances in which a given activity may result in a latent aircraft behavior or an immediate aircraft behavior. For example, suppose the researcher is interested in understanding the activity ‘engage Vertical Navigation (VNAV) mode.’ The activity belongs to the strategic and vertical classes, because it entails first programming the FMC with the vertical profile to be flown. Further suppose that the researcher wishes to develop a visualization that shows how the ‘engage VNAV’ activity links to the aircraft behavior ‘begin VNAV descent.’ Two instantiations of the ‘begin VNAV descent’ activity may be relevant: a latent one when the aircraft does not begin the descent until the VNAV computed top-of-descent point, and an immediate one when the aircraft starts descending as soon as VNAV mode is engaged. Creating such a visualization entails including context-specific rules for linking pilot activities to aircraft behaviors. The cells in Table 1 therefore contain such rules.

Once the relationships to be visualized have been specified, a computer-based implementation of the visualization may be constructed. This research extends CATS for this purpose. The CATS data server merges computer-readable data on the state of the aircraft and its automation, ATC clearances, and pilot activities. It also contains event filters for extracting salient events from the data. CATS can further ‘detect’ aggregate (multi-action) activities (e.g., load Arrival procedure), so that such activities can be visualized.

CATS extensions are, first, a facility to record the sequences of the events (including mode changes), activities, and behaviors required to construct the visualization. A second extension is a ‘linking’ capability to assess the context-specific rules required to associate a pilot activity with a specific aircraft behavior and record the links. In addition, the record is time-stamped to enable either temporal or sequential viewing. The following section describes the interface format used to implement the visualization methodology.

Table 1. Example form of a table of links between pilot activities and aircraft behaviors decomposed into salient classes, with context-specific linking rules in cells.

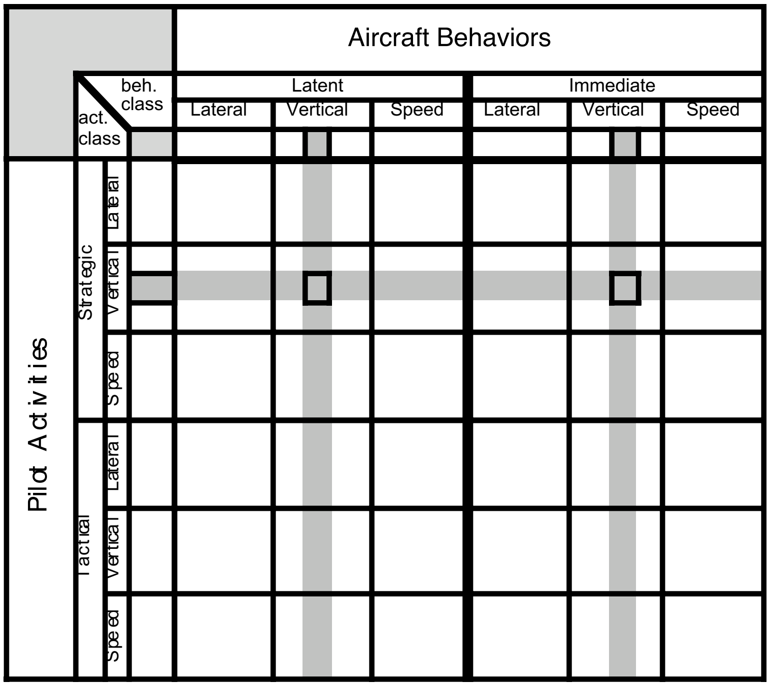

VISUALIZATION INTERFACE

The visualization interface, as implemented, is organized in a combined graphical and tabular format (see Figure 1). It is divided into five rows that represent the following data categories: prompting events (①), crew activities (②), automation behaviors (③), engaged automation modes (④), and ATC facility (⑤). These data categories are labeled in a column along the left side of the display.

Data of each category are presented graphically in the body of the interface. The top four rows represent the flow, from top to bottom, of events, activities, behaviors, and state changes. Prompting events at the top serve as triggers for crew activities or latent automation behaviors. Crew activities on the second row link to automation behaviors on the third row. Automation behaviors are often manifest as mode changes, which are depicted on the fourth row. The fifth row displays the ATC facility controlling the aircraft during a given period, which provides position context for the other data categories.

Prompting events, crew activities, and automation behaviors are discrete events and are represented by small boxes on the display; connecting lines represent links between activities and behaviors (⑥). Continuous data on currently active automation modes are depicted in a linear ribbon-like format on the display. The mode ‘profiles’ also depict qualitative information about when the aircraft has leveled off (⑦). Finally, vertical lines demarcate controlling ATC facilities (⑧).

The present implementation further decomposes the data categories to distinguish different aspects of a particular data set. The prompting events are decomposed into ATC clearances and aircraft position. The crew activities, automation behaviors, and automation states are each decomposed according to whether they affect the speed, lateral, or vertical aspects of the aircraft trajectory. Furthermore, to differentiate between strategic and tactical crew activities and resulting mode engagements, as well as latent and immediate aircraft behaviors, the present implementation color-codes the boxes and mode profiles. In the grayscale versions shown here, light shades indicate strategic activities and modes, and latent aircraft behaviors; dark shades indicate tactical activities and modes, and immediate aircraft behaviors. The present implementation also displays each component with a numeric code assigned during the specification process, which can be suppressed from the display if desired. Clicking the mouse on an event, activity, or behavior box displays specific information (⑨) to support detailed analyses.

Figure 1 also illustrates how the visualization interface can be dynamically reconfigured to visualize information sequentially or temporally. Information is easier to distinguish and inspect in the sequential view, while the temporal view highlights activities and aircraft behaviors clustered in time. A scale adjustment facility is provided for the temporal view (⑩). The present implementation also supports printing for offline analyses and collects a variety of statistics.

EXAMPLE APPLICATION

The visualization methodology was implemented and

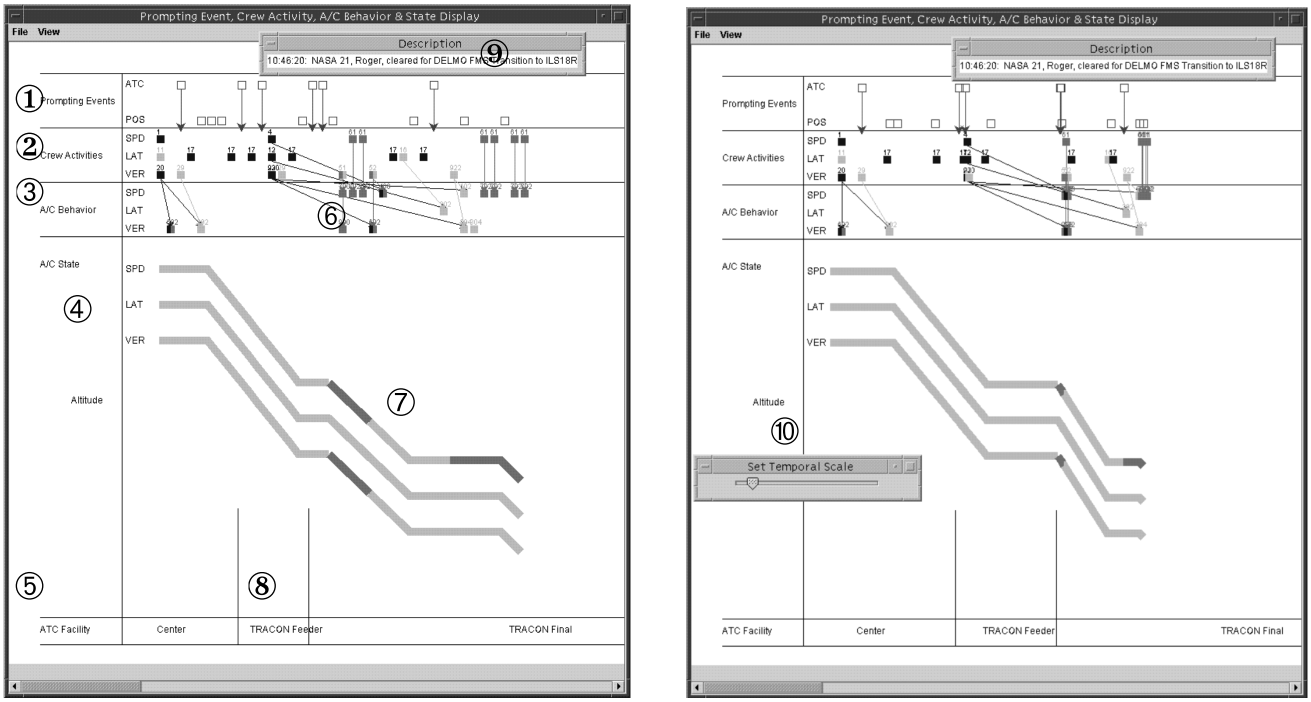

10:46:20: NASA 21, Roger, cleared for DELMO FMS Transition to ILS18R

Figure 1. The visualization interface in sequential (left) and temporal (right) modes.

used to support human factors analyses of data from a NASA study of a new arrival procedure [5]. The study sought to address, among other issues, whether and how pilots could use high levels of automation during the arrival phase of flight, so that actual aircraft trajectories could match trajectories predicted by ATC automation.

In the first study condition, airline pilot subjects flying the NASA Advanced Concepts Flight Simulator (a full-motion glass cockpit simulator) were allowed to comply with ATC clearances just as they would during current line operation. The second study condition required the use of the new 'FMS Arrival' procedure, designed to allow high levels of automation to be used comfortably in the terminal airspace. A third study condition added an additional challenge: the FMS Arrival procedure was interrupted and later resumed to investigate how contingencies might affect pilot workload and automation usage.

To supplement other data analyses [4], CATS processed data from each experimental trial and constructed visualizations using the approach and implementation scheme described above. For illustration, Figure 2 shows visualizations from each of the three study conditions. The numbers in Figure 2 refer to various pilot-automation interaction issues that can be rapidly identified in the samples. For example, interactions and resultant behaviors for the current-day operations samples are reasonably consistent (❶). The dark-colored vertical and speed profiles in the visualizations indicate the subject pilots opted for lower levels of automation when arriving in the terminal airspace.

Visualizations of the FMS Arrival condition show that it is possible to fly the FMS Arrival procedure and remain solely at the highest levels of automation (❷); however, there are also examples in which pilots had to intervene to various degrees (❸ and ❹). Closer inspection of these cases confirms speed control can at times be difficult, and that VNAV mode may be not suitable. Whether VNAV was not used for a specific reason is not immediately obvious from the visualization, but may be further elucidated by querying a computer-based version of the visualization.

The areas marked by ellipses in ❸ and ❹ demonstrate how the visualization can support in-depth examination of particular interactions. Each reveals different reasons for opting for a lower level of automation. In ❹, the pilots did indeed feel VNAV was unsuitable, and chose a different mode ('Vertical Speed'). Additional visualization tools described in Callantine (1999) confirmed speed control was indeed the problem. In ❸, however, the pilots tried to reengage VNAV but could not, because the correct target altitude was not set. This points to a different sort of difficulty.

Finally, returning to rapid identification of interactions, there is a vast increase in the number of pilot activities, salient aircraft behaviors, and associated mode transitions for the samples in which the FMS procedure is interrupted and resumed (❺). Moreover, resumption of the procedure is not smooth, if it occurs at all.

RESULTS AND FUTURE DIRECTIONS

The visualization methodology presented here is unique in that it provides a picture of the overall system and the relationships between its elements: the operational environment, operator activities, and machine behavior. As implemented, it proved useful as a supplementary tool for investigating pilot-automation interactions in a proof-of-concept application. The visualizations allowed rapid identification of differences between the study conditions and individual trials, and favorably supported detailed investigations of particular interactions.

Applying visualization techniques is a natural way to further our understanding of pilot-automation interaction. Issues regarding aircraft compliance with ATC clearances using automation may hinge on clearance timing and content, which in turn reflects ATC automation issues. Besides aiding visualization of interactions as they relate to ATC operations, visualization could help efforts to understand how to better support mode usage (e.g., [6], [7], [8]). Visualizing the links between a mode usage activity and the resultant aircraft behavior is useful for understanding shifts between control and monitoring activities and their effects on cockpit resource management and workload, as well as when and how pilots shift between levels of automation.

Future work may include refinement of the visualization tool as a part of a CATS-based design and analysis suite. The refined tool could integrate data-mining capabilities that enable the researcher to, for example, highlight activities of a particular type, or replay the portion of the data that contain an interesting interaction pattern. The visualization methodology could also be applied to visualizing air traffic controller-automation interactions.

REFERENCES

1. Callantine, T. (1999). Visualization and analysis of flight crew procedure performance. Proceedings of the 8th International Conference on Human-Computer Interaction, Munich, Germany, 1281-1285.
2. Callantine, T., Mitchell, C., and Palmer, E. (1999). GT-CATS: Tracking operator activities in complex systems. NASA Technical Memorandum 208788. Moffett Field, CA: NASA Ames Research Center.
3. Callantine, T., Palmer, E., and Smith, N. (1997). Activity tracking for procedure refinement. Proceedings of the 1997 Conference on Systems, Man, and Cybernetics, Orlando, FL, 3126-3131.
4. Crane, B. (1999). Flight crew factors for CTAS/FMS integration in the terminal airspace. (Unpublished Masters Thesis) San Jose, CA: San Jose State University.
5. Crane, B., Prevôt, T., and Palmer, E. (1999). Flight crew factors for CTAS/FMS integration in the terminal airspace. Proceedings of the 8th International Conference on Human-Computer Interaction, Munich, Germany, 1276-1280.

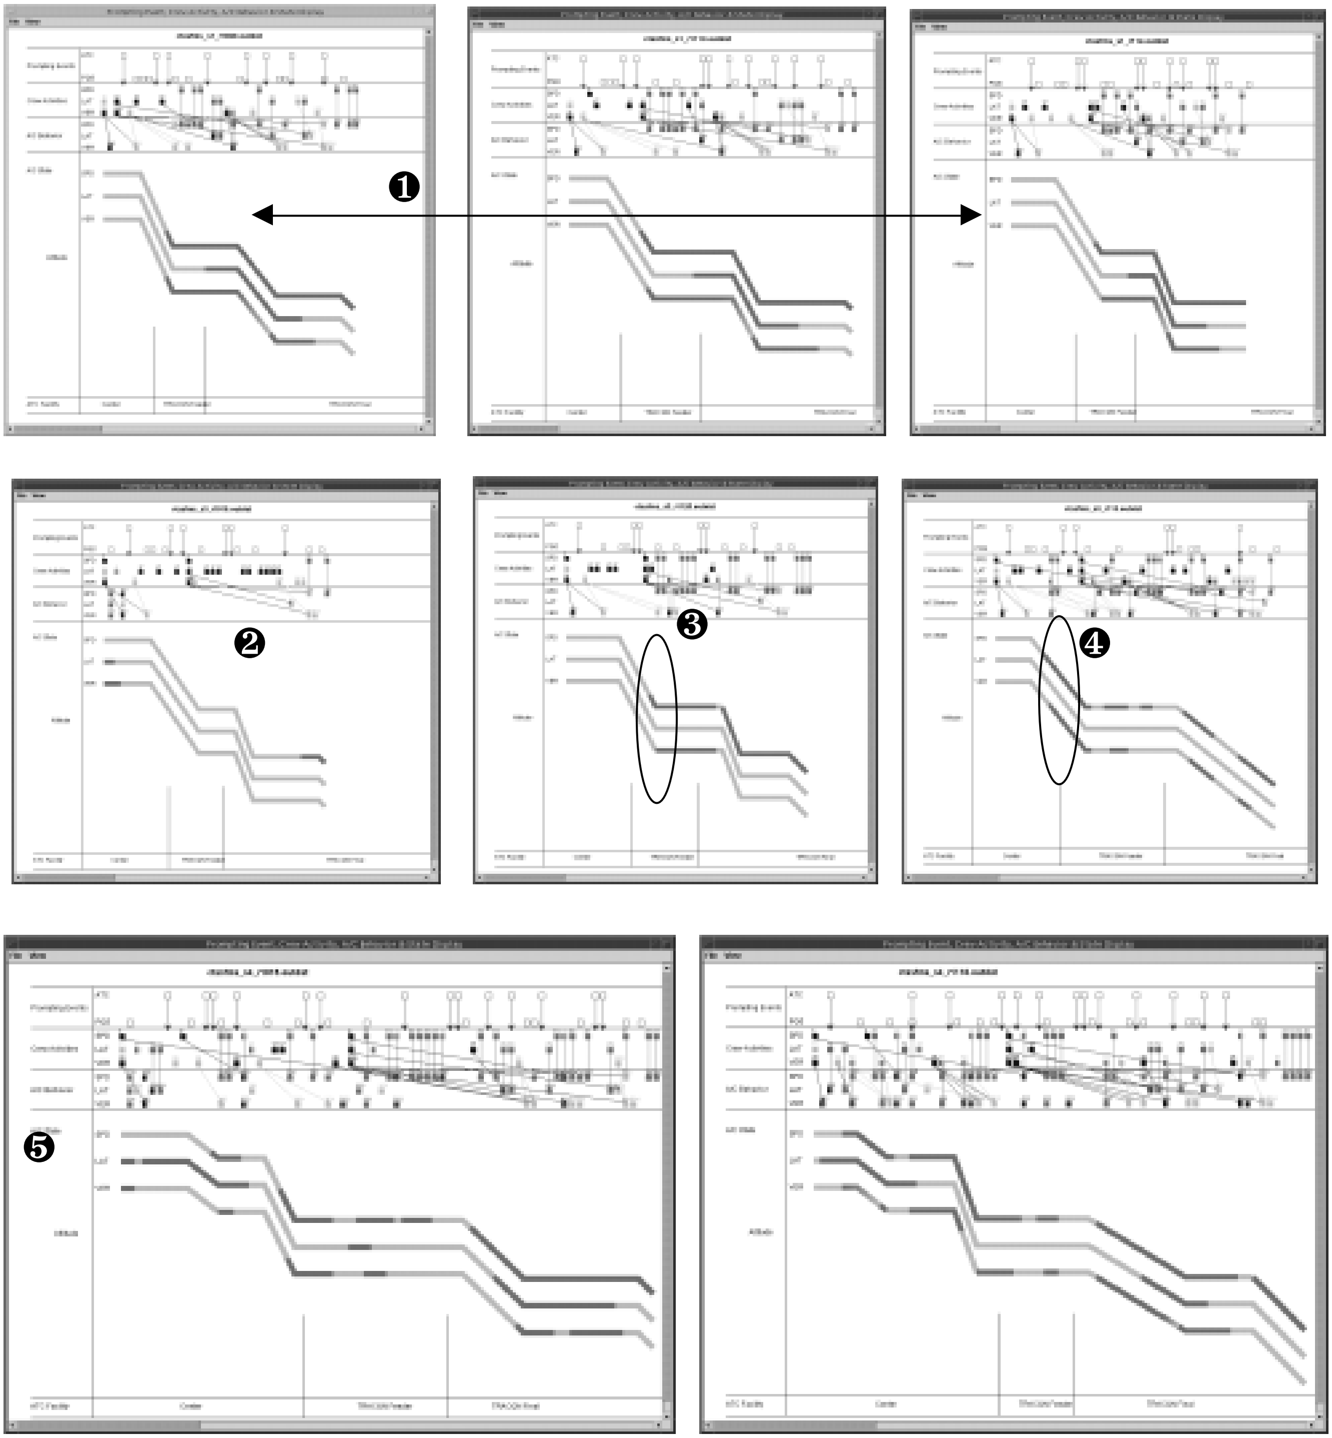

Figure 2. Example visualizations of experimental data from Crane, Prevôt, and Palmer (1999). The top row depicts current-day operations, the middle row operations using the FMS arrival procedure, and the bottom operations when the FMS procedure must be interrupted and resumed. All views are sequential.

6. Degani, A., Shafto, M., and Kirlik, A. (1999). Modes in human-machine systems: Review, classification, and application. International Journal of Aviation Psychology, 9(2), 125-138.
7. Javaux, D., and De Keyser, V. (1998). The cognitive complexity of pilot-mode interaction: A possible explanation of Sarter and Woods' classical result. In G. Boy, C. Graeber and J. Robert (Eds.), Proceedings of the International Conference on Human-Computer Interaction in Aeronautics Conference. Montreal, Quebec: Ecole Polytechnique de Montreal, 49-54.
8. Sherry, L., and Polson, P. (1999). Shared models of flight management system vertical guidance. International Journal of Aviation Psychology, 9(2), 139-154.
9. Schneiderman, B. (1998). Designing the user interface: Strategies for effective human-computer interaction, Third Edition. Reading, MA: Addison-Wesley.