Pilot Activity Profile Aggregation based on Flight Data Analysis

Laurent Chaudron, David Guéron and Nicolas Maille
ONERA - The French Aerospace Lab
F-13661, Salon de Provence, France
{laurent.chaudron, david.gueron, nicolas.maille}@onera.fr

Abstract

Various methodologies directed towards improving flight safety proceed via systematic flight data analysis. Some rely on the search for prescribed deviations, but the work presented here focuses on methodologies that signify the activity of the pilot via an *activity profile* whose definition usually requires a huge amount of expert knowledge and experimental setups. We here describe a semi-automatic construction of these profiles through a three-step process of *supervised aggregation*. Using this methodology, we could single out a sequence of pivotal rules giving account for singular events of the observed activity and combine them in a computational profile.

1. Introduction

1.1 Operational challenge

Aiming for increased safety-management during air operations, ICAO (International Civil Aviation Organization) prompts airline companies to set up Flight Operations Quality Assurance (FOQA) programs based on a systematic analysis of in-flight recorded digital data.

The studies carried out through these programs chiefly rely on overcame thresholds and on observed discrepancies; in particular, theses studies are not relevant for examining odd contexts and can not account for unexpected differences in a relevant way ([9]). To make up for this inconvenience, several approaches relying on the identification of atypical flights by statistical tools have been recommended ([1, 4]). Nonetheless, these approaches require a significant amount of observations (at least 1000 for each kind of activity) as well as a sort of regularity in the conducts of the flights: these restrictions make them hardly adaptable to less restrictive tasks, such as helicopter rescue assignments or military operations.

1.2 Flight Profile

ONERA[1] has been concerned for some time about this problematic of systematic analysis, favoring an approach based on the search for deviations in the *activity* of the crew members.

This approach relies on a model able to account for human activity ([8]), using an activity language stemming from N-MDA (a French acronym for Kernel of a Model for Activity Description, [5]), the LDA (the French acronym for Language for Activity Description, [3]). Using this language, basic actions of the crew (for example, the retraction of the landing gear) can be described and grouped within a shared *Flight Profile*.

Flight Profiles roughly describe usual implementations of flight procedures. Each Flight Profile contains sequences of elementary actions, rules allowing to identify these actions, and constraints about the possible sequences of actions. In that way, each profile stands for a specific way to carry one phase of flight, thus appearing as an abstraction, a *meta-object*, that gathers together specific combinations: a Flight Profile is a *way of doing*, a way of piloting, that consists in the set of attended procedures, as well as in several indicators that can distinguish between situations.

[1] *The French Aerospace Lab*. Public Establishment in charge of Aerospace Resarch. http://www.onera.fr/english.php .

Given the digital flight parameters recorded during a flight, one must be able to connect the portrayed activity with an adequate Flight Profile, and to compare it to the one captured in the profile.

A major issue in this new approach lies in the time and the expert knowledge that are required for the determination of these *operational* Flight Profiles. The here discussed work, that intends to carry this step in a semi-automatic manner, consists in finalizing a set of tools intended to help the expert in the construction and the layout of the rules, starting from a set of exemplary flights. It relies on the second step of the supervised aggregation mechanism described in [7].

1.3 Outline of the article

The aggregation mechanism used is detailed in section 2, Supervised Aggregation. Section 3, Application..., pictures the context in which our demonstration of the second step of supervised aggregation, the maieutic step, takes place. Finally, section 4, Results..., examines the results obtained using the methodology on a set of commercial flights.

The last step of the aggregation procedure consists in finalizing methods and tools for combining these basic activities into a Flight Profile.

2. Supervised aggregation

2.1 An abstraction mechanism

Flight Profiles are currently individually handmade by experts, using a costly process based on:

- standard procedures described in flying guides (which we call *theoretical expertise*),
- interviews conducted with operators, as well as with instructors (which we refer to as *operational expertise*),
- regular updates of expertise and regular fitting of profiles, first of which the feedback of security officers on observed activities.

In some way, a Flight Profile can be looked at as capturing the *mean activity* of the studied flights. Notwithstanding, the word "mean" too strongly reflects a) characters that can be best signified by formulas (of which each numerical mean, *i.e.*: geometric, arithmetic, quadratic, harmonic...) and, more important, b) operations whose outcome is similar to the incoming material. Yet, inasmuch as we are concerned with human activities, which are complex objects (at least more complex than straight figures), the outcome at this step can not be another activity. As a matter of fact, a Flight Profile is a set of rules (*association rules*) associating the various observed patterns with one another: it thus appears as an object of a higher abstraction level which signifies the mechanisms displayed during the studied activity.

2.2 A 3-step process

This document introduces the first steps in the setup of the supervised aggregation process which, starting from a set of exemplary flights and an expertise on them, will lead to the discovery of a Flight Profile that reports for the activity of the crew.

Figure 1 on the facing page pictures the examined aggregation process which, starting from flight tracks (in the above diagram, these appear to be the paths of the aircraft, but other kinds of tracks exist), identifies their significant features (meaning: those to which we can attach a semantic value), the *patterns*, before synthesizing these objects by establishing a Flight Profile.

The supervised aggregation process that we devised, composed of the three steps of 1. decomposition, 2. maieutics, 3. reconstitution, is based on the generic aggregation portrayed in [2]. It nevertheless differs from the process there described in the nature of its second step which, in Barberà's process, is but an averaging step that ends up in a quantity similar to the income one, obtained from the decomposition step. And, as we have already mentioned, concerning objects as complex as human activities (which include subjectivity and which are liable to be tackled from various angles), it seems hopeless to expect getting a coherent object merely by gathering the individual means of the different characteristics used to signify the activities.

This process has been portrayed in [6], that sums up our; later enrichments have been described in [7].

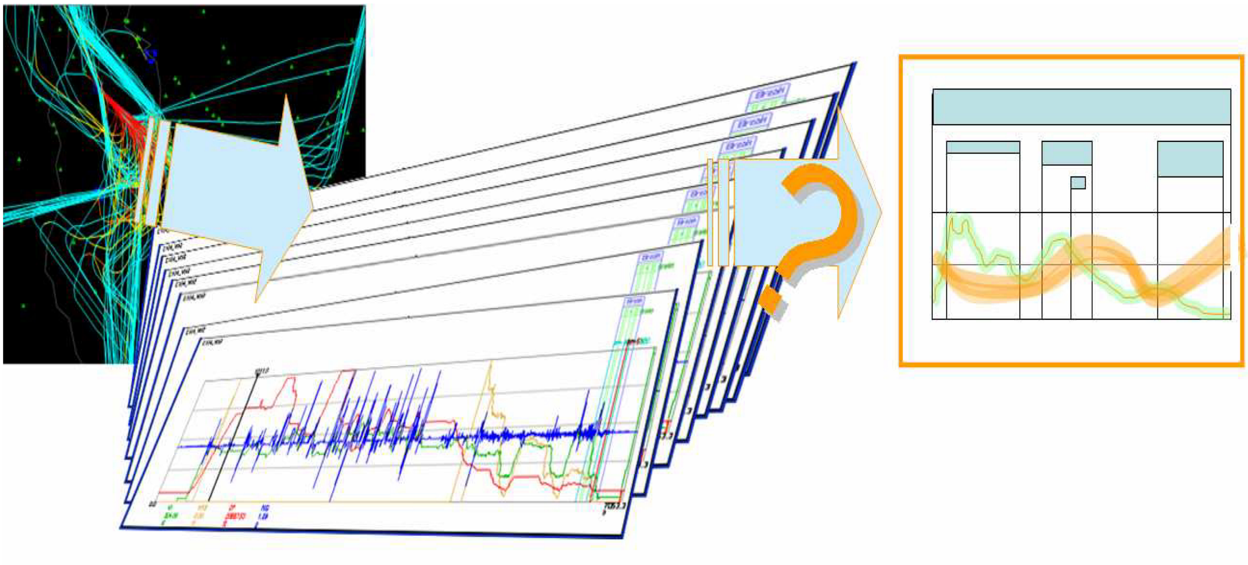

Figure 1: From flights to Flight Profiles

2.3 The 3 steps of supervised aggregation

We thus fitted the aggregation process to our needs in the following manner:

2.3.1 Decomposition

This first step consists in choosing a set of simple characters in order to signify the properties of interest in the description of the objects. This set will be submitted to our aggregation procedure. In actual practice, not much choice is left and this step is technically easy.

In the work here described, each flight is decomposed in a set of parameters of various natures (numeric or symbolic) that can stand for 3 different kinds of indications:

- actions of the pilot (via the operated controls),
- measures of surrounding environment (for example, temperature or luminosity),
- measures of inside environment, that describe the state of the aircraft and the one of the pilot.

These parameters are objective witnesses of what occurred during the flight, in terms of crew actions as well as in terms of aircraft reactions.

2.3.2 Maieutics

This step is the core of the supervised aggregation process, on which depends the following step, that uses its results.

It consists in providing, for several elementary activities, a set of indicators (such as ”raising the nose wheel” or ”extending the main landing gear”) in the form of *patterns* providing access to characters (at present, dates of beginning and end of an activity) related to these activities.

This phase is carried out by way of a supervised aggregation process that uses a set of exemplary flights so as to learn the convenient patterns.

2.3.3 Reconstitution

This step combines the patterns sprung from the previous one in order to make up a new object — the outcome of the aggregation of the original objects — whose kind is nonetheless different (this new object is the Flight Profile, made up of the relations between the observed patterns). Such a profile can then signify a specific way to carry a takeoff, depending on the ordered set of the elementary actions that constitute it, together with their conditions of realization.

In other words, the patterns obtained in the last step are aggregated into the more general *ways of doing* which, as for themselves, are specific executions of existing *procedures*.

We will now describe how we conducted the first two steps of this approach (being mainly concerned with the second), for use on a set of commercial flights.

3. Application to A330 commercial flights

The first two steps of the here-described supervised aggregation process have been used on a set of 20 real flights (meaning that they were no simulation flights), for which a few elementary activities of the takeoff stage were given by the experts. We now briefly describe the representation choices that have been made; in the following section 4, we will indicate some of the obtained results.

3.1 Decomposition step: using flight data

As noticed in paragraph 2.2, a flight leaves a wide variety of tracks: aside the visual, acoustic or atmospheric tracks, several indicators also exist that describe the human and mechanical activities during the flight.

We here restrained ourselves to 83 (including time) recorded parameters which, despite being of various natures, could all be handled like numerical figures, and which appeared to us as constituting a thorough-enough set to allow the first intended tests of methodology. Especially, the selected parameters allowed us to follow the states of flight controls (control column, flaps, vertical fin…), the modes of available automated systems (autopilot), as well as the main aerodynamic characters of the aircraft (speeds, altitudes…).

Figure 2 on the current page shows the evolution of some of the parameters for one of the examined flights.

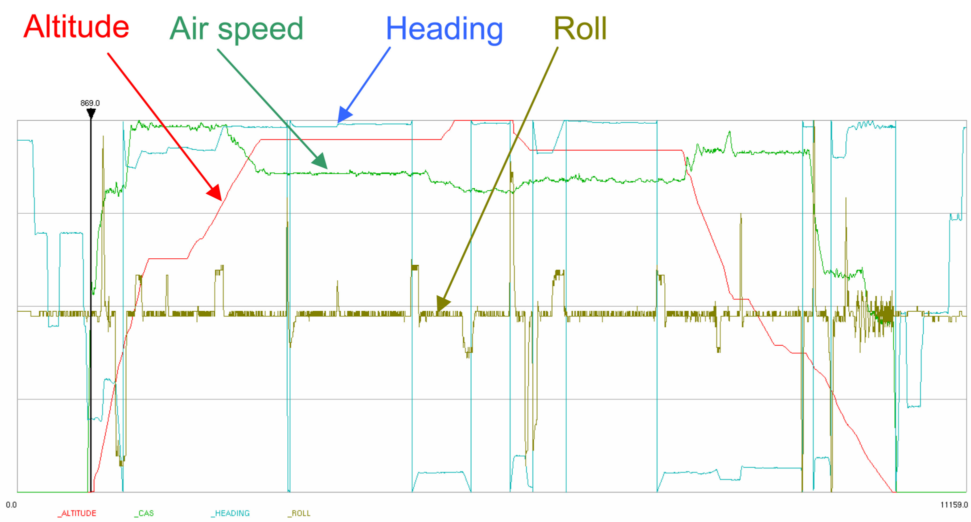

Figure 2: A subset of the measured parameters, throughout a flight

Notwithstanding an insignificant approximation, we could consider that the parameters were measured at regular time intervals of 1 second. In what follows, $\mathcal{R}_T$ stands for the set of times when some measure was made, and $\underline{\mathrm{X}}(t)$ refers to the value taken by parameter X at the closest time of measure prior to time t.

3.2 The maieutic step: creating patterns

3.2.1 The training set

The maieutic step, dedicated to the build-up of patterns which identify elementary activities, consists in a supervised learning process that uses a set of flights for which is available some kind of *expert knowledge*.

In the work related here, this expert knowledge consists in the dates of 6 events, common to each of the 10 flights of our "learning set". On table 1 on this page, these "Key Points" are listed KP1 through KP6.

	KP1	KP2	KP3	KP4	KP5	KP6
Flight1	869	877	944	988	1031	2877
Flight2	1178	1185	1264	1303	1330	2907
Flight3	1355	1361	1444	1481	1511	3228
Flight4	1027	1033	1090	1146	1193	2697
Flight5	1459	1465	1542	1588	1616	3190
Flight6	955	961	1046	1082	1114	3044
Flight7	1045	1053	1138	1166	1223	2872
Flight8	533	541	626	660	704	2478
Flight9	989	995	1080	1111	1138	2575
Flight10	843	849	934	969	999	2664

Table 1: Dates (expressed in seconds after the start of the flight recorder) of the 6 events, in each of the 10 flights.

During the learning stage, the meaning of the 6 elements to be described was unknown, but to the expert by whom they had been identified, and we will discuss in the following section the relevance of the obtained patterns, according to this meaning.

3.2.2 The set of candidate patterns

We previously stated that, for each of the searched-for activity element, the outcome of the maieutic stage was a pattern identifying (the date of) this element from the set of available data on the considered flight.

A basis of *candidate patterns* will be considered, and be made operational through a routine for instantiating its thresholds. In this test, conducted for evaluating the selected methodology, we considered a set of 4 kinds of patterns, each requiring the specification of 2 thresholds, expected to reveal the variations of a numerical figure. This basic set may and shall later be extended by addition of numerical stabilization patterns, as well as of patterns that concern symbolic parameters.

The supervised learning process first consists in estimating the thresholds for each of the 4 kinds of patterns, applied to each of the measured parameters, assuming that the thus-instantiated pattern would result in the available expert knowledge. The final step is then to select the best amongst all the built-up patterns.

The shapes of the increase and decrease patterns can be found on the next page. A parameter is observed to increase when its value gets over a specified threshold, called "confirmation threshold" (`cTh`), the increase having been instantiated at the latest anterior date for which this parameter has crossed another (smaller) threshold, called "start threshold" (`sTh`) (see figure 3).

3.2.3 Pattern construction and selection

The basis of candidate patterns we used for the learning process is made of 4 kinds of generic patterns, that can be used with each of the 82 measured parameters. Following the stage in which thresholds are established, 328 instantiated patterns are now available to describe each Key Point, and it now remains to decide which ones, amongst those, identify this Key Point the best. This done, will have been selected: 1) a kind of pattern, 2) an application parameter, and 3) two thresholds.

For each elementary activity to be characterized, the learning process includes two steps: the instantiation of the 328 candidate patterns, followed by the selection of the most relevant among them.

PATTERN 1 *Parameter increase (resp. decrease)*
Given two thresholds $\mathtt{sTh}$ *and* $\mathtt{cTh} \in \mathbb{R}_+$ *such that* $\mathtt{sTh} < \mathtt{cTh}$ *(resp. >), parameter* X *will be said to increase (resp. decrease) significantly if and only if*

$$\exists\, \mathtt{sT},\ \mathtt{cT}\ \in \mathcal{R}_T :$$
$$\begin{cases} |X(\mathtt{sT})| \leq \mathtt{sTh} \quad \textit{and} \quad |X(\mathtt{cT})| \geq \mathtt{cTh}\ , & (\textit{resp.}\ \ \geq, \leq) \\ \forall\ \mathtt{t} \in]\mathtt{sT}; \mathtt{cT}[\, \cap \mathcal{R}_T\ , \quad \mathtt{sTh} < |X(\mathtt{t})| < \mathtt{cTh}\ . & (\textit{resp.}\ \ >, >) \end{cases}$$

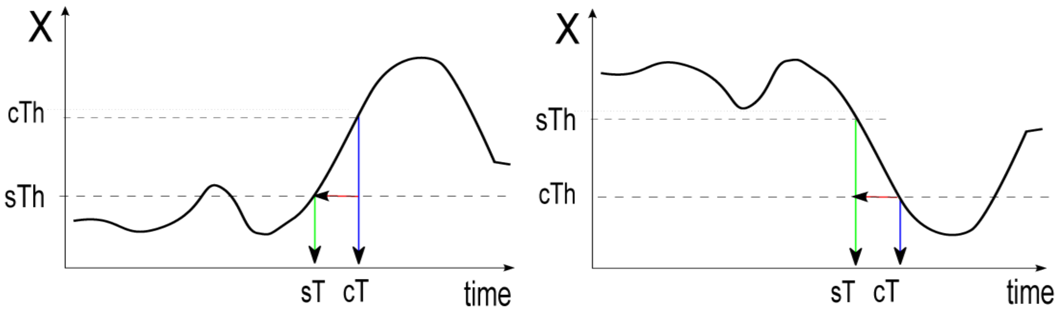

(a) Increase (b) Decrease

Figure 3: Evolution of a parameter

Threshold instantiation The algorithm used for computing the thresholds for the increase and decrease rules is the following:

Pattern assessment The process for selecting a pattern is made of two successive steps, during which are evaluated 1) the consistency of the thresholds, 2) the relevance of the results it provides.

The coherence criteria permits pruning a large number of instantiated patterns. For example, as the increase patterns are concerned, those whose confirmation threshold, $\mathtt{cTh}$, is below the start threshold, $\mathtt{sTh}$. This criteria thus appears specific to the generic pattern on which it is used (but unspecific of the considered parameter).

The relevance criteria is obtained by using the pattern for the very flights that were used to instantiate its thresholds. Each pattern is then evaluated according to the discrepancy between the dates specified by the experts and the dates obtained with the pattern. In the described work, the standard deviation σ has been chosen, amongst others that have different characteristics: it is here possible, as well as desirable, to dismiss deviant data and to identify some groups of expert maneuvers. The pattern minimizing this standard deviation is then considered to be the most suitable to reflect the event.

4. Results and discussion

4.1 First Key-point: combination of simple patterns, incidental patterns

Regarding the first date to analyze (KP1), the first patterns obtained are, in relevance order:

1. the increase of the longitudinal trim ($\sigma = 0,707$) ;
2. the increase in altitude ($\sigma = 0,837$) ;
3. the increase in radiosonde altitude [2] ($\sigma = 0,949$) ;

Algorithm 1 *Threshold computation for the increasing (resp. decreasing) parameter pattern.*

1. *Initializing thresholds :* $\texttt{sTh} = 0$, $\texttt{cTh} = +\infty$ *(resp.* $\texttt{sTh} = +\infty$, $\texttt{cTh} = 0$*).*
2. *Updating thresholds :*
choosing to indicate the set of flights constituting the expert basis as $n = |\mathcal{F}_{exp}|$*, the thresholds are updated as indicated below:*

$$\texttt{for i from 1 to n, do :}$$

$$\texttt{sTh} = \max\left(|\underline{\texttt{X}}(\texttt{Kp}^i)|,\ \texttt{sTh}\right) \qquad \textit{(resp. } \min\textit{)}\ , \tag{1}$$

$$\texttt{cTh} = \min\left(\max_{\substack{\texttt{t}\geq\texttt{Kp}^i\\ \texttt{t}\in\mathcal{R}_T}}(|\texttt{X}(\texttt{t})|,\ \texttt{cTh})\right) \qquad \textit{(resp. } \max(\min)\textit{)}\ . \tag{2}$$

4. the increase in longitudinal wind ($\sigma = 1,000$).

This first event corresponds to the takeoff of the aircraft (the time when the main landing gear leaves the runway). In this respect, the first three patterns appear plausible as to the physical phenomenon that is being described. The fourth is, as far as plausibility it is concerned, more surprising: wind doesn't change at takeoff! Its good positioning is due to the peculiar functioning of the data recording system on the plane, the longitudinal wind being kept at 0 as long as the main gear remains on the ground.

It is good to notice that other parameters could have been used in order to identify this first element, particularly boolean parameters related to the state of the main gear (on ground or in the air). Thus, the already evoked utility for patterns concerning other kinds of parameters than continuous (discrete or categorical).

On the other hand, having several candidate patterns (low standard deviation) may allow to build up more robust patterns by combination.

4.2 Second Key-point: unexpected patterns

As for the second Key Point (KP2), the best two patterns are the increase in radiosonde altitude and a decrease in the aircraft's incidence.

This Key Point identifies the landing gear retraction. The pattern applied to radiosonde altitude shows that this action was undertaken at more or less constant heights, independent of the conditions of the various experiments: if this pattern is irrelevant as for the identification of the landing gear retraction in a new flight, it nevertheless allows, by putting into relation two measures *a priori* independent, for the discovery of a specific kind of piloting. As for the pattern concerning incidence, it is related to the physical phenomena caused by the retraction of the landing gear: the shape of the aircraft is modified and its orientation in the airflow changes.

These results show the importance of expert-supervision for the selection of the pattern. Eventually, it should be remarked that this Key Point 2 could as much have been depicted by patterns applied to boolean parameters.

4.3 Subsequent Key points

Relevant patterns have also been found to depict the following Key Points (KP3, 4, 5 and 6). For example, Key Points 4 and 5, that describe the beginning and end of a turn for disengagement from the runway axis after takeoff (and that isn't part of some takeoff procedures) are identified by patterns concerning the aircraft's lateral inclination (roll).

However, if it is still possible to calculate the thresholds of the various candidate patterns for these last four Key Points, we can not anymore assess in a pertinent way the relevance criteria by testing the pattern on the expert flights.

[2] Radiosonde altitude supplies the distance of the aircraft to the ground, as given by a probe. The given figure is restricted by the probe's range. Altitude supplies the distance regarding a specified reference (the takeoff runway, the sea level…), which can be modified by the crew during the flight.

Indeed, in so doing, the pattern is captured by earlier variations of the parameter, probably related to other maneuvers and actions performed by the pilot.

The solution to this inconvenience, which seems easy to put into practice, is to allow the application of the patterns between specific dates, specifying when the pattern is to be looked for, as well as to allow the search engine not to stop at the first case of the pattern that is observed, in order to be able to select the one which suits us the best.

5. Conclusion

In the present paper, we demonstrated the interest and feasibility of a supervised aggregation process in order to come up with Flight Profiles.

In particular, the results demonstrate that using generic patterns allows the characterization of meaningful flight events. It should also be remarked that, in the presently available tools for experience feedback that are based on the recognition of elements of activity, these patterns are produced by experts for whom adjusting internal thresholds is a tricky step. The learning process that was set up proves to be a precious help for ensuring of the sturdiness of the obtained patterns. Between other things, it provides for threshold evolution, in order to account for new flights in a more controlled way, adding these flights to the learning basis. Besides, in order to better characterize the studied activity, this process makes available a few patterns which might not have been considered by the analyst.

The obtained results also show that the pattern basis currently used must be enriched, on the one hand to take into account discrete parameters and, on the other hand, to allow for using more complex patterns, that could be applied to several different parameters. These points are currently under study.

The maieutic step of the aggregation process being for its main part under control, the problematics of the reconstitution step now unlock. This last stage achieves the assembling of elementary activity elements in a Flight Profile. Integrating temporal constraints will thus be at the core of the work to come.

References

[1] Brett G. Amidan and Thomas A. Ferryman. 2005. Atypical Event and Typical Pattern Detection within Complex Systems. In: *IEEE Aerospace Conference, Manhattan Beach*.

[2] Salvador Barberà. 2003. A Theorem on Preference Aggregation. Working Paper n°166. Barcelona Economics Working Paper Series.

[3] Pascale Bonnet. 2008. A description language for piloting activity, intended at the systematical analysis of flights, taking human activity variability into account. In French. French Title: Un langage de description de l'activité de pilotage, pour l'analyse systématique des vols et l'étude de la variabilité de l'activité. ONERA, Mémoire de Formation 1/13534.

[4] Suratna Budalakoti, Suratna Srivastava and Matthew E. Otey. 2009. Anomaly Detection and Diagnosis Algorithms for Discrete Symbol Sequences with Applications to Airline Safety. In: *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*. 39.

[5] Sybille Caffiau, Patrick Girard, Laurent Guittet and Dominique Scapin. 2009. Hierarchical Structure: A Step for Jointly Designing Interactive Software Dialog and Task Model. In: *HCI International 2009*. 1:667–673.

[6] David Guéron, Nicolas Maille, Laurent Chaudron and Jean Caussanel. 2010. Aggregation of Human Activities, Building Flight Profiles. In: *Decade Of the Mind VI* (DOM VI).

[7] David Guéron, Laurent Chaudron, Jean Caussanel and Nicolas Maille. 2010. Aggregation of human activities, discovering flight profiles using recorded flight parameters. In French. French Title: Agrégation d'activités, découverte de profils de vol à partir de traces de paramètres. In: *INFORSID, atelier ICT*.

[8] Patrick Le Blaye, Laurent Chaudron and Jean-Yves Grau. 2000. REX: a human factor flight safety research program. In: *ICAS'2000, 28th Aug.- 1st Sept*. Harrogate, United Kingdom.

[9] Irving C. Statler. 2007. The Aviation System Monitoring and Modeling (ASMM) Project: A Documentation of Its History and Accomplishments: 1999-2005. Available at http://www.nasa.gov/pdf/225024main_TP-2007-214556%20ASMM_Project.pdf.