Compact Preference Representation for Pilot Decision Recommendation

Denys Bernard[1][a], Jean-Claude Méré[1] and William Nicolas[2]
[1]*Airbus S.A.S, Toulouse, France*
[2]*ALTEN SO, Toulouse, France*

Keywords: Decision Making, Flight Diversion, Compact Representation of Preferences, Qualitative Reasoning.

Abstract: Our goal is to apply compact representations of preferences (Kaci et Al., 2020) in pilot recommendation functions for future commercial aircraft. The support of a decision assistant would be helpful in a variety of flight situations, and we focus here on the case of a diversion decision. CRP are based on simple, modular and intuitive representation of preferences among a set of candidate solutions. Solutions are represented as vectors of qualitative variables. CRPs help to define a logical language for specifying «preference statements". Those preference statements are used to sort a set of candidate solutions (or "outcomes"). Each outcome is represented as a vector of qualitative or propositional variables. The conceptual simplicity of CRP facilitates knowledge elicitation and explanation processes. We developed a variant of an existing framework named CP-theories (Wilson, 2011) which fulfils our expressivity and operational constraints. The language and algorithms of our framework have been applied to support pilots to make the best decision about flight diversions.

1 INTRODUCTION

The purpose of this study is to adapt and apply a family of approaches known as "compact representations of preferences" (abbrev. CRP) (Kaci et Al., 2020) to the design of pilot recommendation functions for future commercial aircraft. The support of a decision assistant would be helpful in a variety of flight situations, and we focus here on the case of a decision for diversion.

CRPs are based on simple, modular and intuitive representations of preference statements. Reasoning on those statements is used to sort a set of candidate solutions (or "outcomes"), each of those solutions being represented as a vector of qualitative variables. CRPs theories define logical languages to specify "preference statements". The conceptual simplicity of CRP facilitates knowledge elicitation and explanation processes. We introduce a variant of an existing framework named CP-theories (Wilson, 2011) which fulfils our expressivity and operational constraints. The paper first reminds the state-of-the-art of CRPs, then shows how we adapted the selected pre-existing framework. Finally we detail some diversion examples where different sets of preference statements are invoked to face different types of operational issues.

2 REPRESENTING AND REASONING ABOUT PREFERENCES

Compact representations of preferences (Kaci et Al., 2020) - abbreviated here as CRP - define logic-based formalisms to specify a knowledge base of unitary "preference statements", as well as reasoning mechanisms for: checking the consistency of the knowledge base; determining the dominance relation between two options ("outcomes"); ordering a set of possible options. Among the different frameworks, graphical approaches use oriented graphs where nodes represent the variables of the possible outcomes, and branches represent priorities or dependencies among variables. Conceptual foundations of graphical approaches of preferences

[a] https://orcid.org/0000-0002-0131-4795

Bernard, D., Méré, J. and Nicolas, W.
Compact Preference Representation for Pilot Decision Recommendation.
DOI: 10.5220/0011925000003622
In *Proceedings of the 1st International Conference on Cognitive Aircraft Systems (ICCAS 2022)*, pages 5-9
ISBN: 978-989-758-657-6

are detailed in (Shoham, 1997), who draws a parallel between probabilities and utilities. In particular he notes that utility independence among variables plays the same role as events independence in bayesian networks: by taking account of dependence relationships, one can drastically reduce the effort needed to compute and compare utilities of alternative choices. Shoham develops this parallel between probability and utility under the form of "utility networks". Graphical approaches for compact representation of preferences are built on the concept of variable dependencies, but without requiring the definition of a quantitative utility function. For example CP-nets (conditional preference networks) (Boutilier et Al., 2004), define a preference theory as a pair (G, CT) where G is a dependency graph over the variables, and CT is a function which assigns a conditional preference table to each variable; The conditional preference table for variable X defines the preferences over possible values of X, for each possible value assignment of the parents of X in G, all other things being equal. In CP-nets, the rule that all the variables other than the parent variables in the net must be equal for the preference statement to trigger ("ceteris paribus" assumption), is too constraining in cases where some variables are always less important than others, or when some variables are not relevant for a particular use case. For this purpose, (Brafman et Al., 2006) enriched the graphical language of CP-nets as TCP-nets (trade-offs-enhanced CP nets) by representing variable priorities. TCP-nets not only display dependencies (like CP-nets), but also show priority relations. It is then possible in a TCP-net to state that under certain circumstances, a variable X is much more important than variable Y. In such a case, one can ignore Y to evaluate solution dominance. Example in diversion scenario: if the safety margin is degraded, flight time is a more important criterion than cost of maintenance at diversion airport, which can then be ignored when comparing two options which are different regarding safety margins.

CP-theories (Wilson, 2011) further generalise preferences networks. For a set of variables V, a cp-theory is a set of statements of form:

$$u: x > x' \ [W] \tag{1}$$

Where: u is a value assignment to a subset $U \subset V$, x and x' are possible values of a variable X, s.t. $X \notin U$, and $W \subseteq V-U-\{X\}$. Such a statements says that an outcome t u x w is preferred to any outcome t u x' w' where: t is a value assignment on $V-(U \cup \{X\} \cup W)$ ("ceteris paribus" variables); u is a specified value assignment on U (set of preconditions); w and w' are any value assignments on W (indifferent variables).

Wilson proposes efficient tree-based algorithms for evaluating consistency and dominance in cp-theories.

The compact representations of preferences have also been considered from the viewpoint of logic. There, a possible outcome is a possible world, where a set of formulas in a primary logical language holds. A preference statement in such a logic says that if certain formulas hold in world W1, and other formulas hold in world W2, then W1 is preferred to W2. For example (Bienvenu et Al., 2010) provides a general logical theory of preferences ("prototypical preference logic") by extending a propositional language L with preference statements formed as:

$$\alpha \rhd \beta \parallel F \tag{2}$$

Where α and β are formulas of L, F is a set of formulas of L. It expresses that we prefer an outcome O1 over outcome O2 if $O1 \vDash \alpha$, $O2 \vDash \beta$, and O1 and O2 agree on the formulas in F. The prototypical preference logic generalises most of the graphical representations and CP-theories.

In the literature, computing preference relations often uses graph algorithms: To decide whether an outcome dominates another one, the algorithm tries to find a path of elementary preference relations through the possible outcomes. For example in CP-nets, a preference relation between two outcomes is obtained by generating a "flipping sequence" i.e. a sequence outcomes, where two consecutive outcomes differ only by one variable. Determining that an outcome is preferred over another one consists in finding a flipping sequence from the first one to the other one. This principle is extended in (Wilson, 2011) whose algorithm generates a "cs-tree" (complete search tree) where the terminal leaves are the possible outcomes, and the intermediate nodes are partially instantiated variable assignments.

3 A VARIANT OF CP-THEORIES FOR PILOT DECISION ASSISTANCE

Our approach is inspired from cp-theories (Wilson, 2011). We had to make a few adaptations to the initial theory in order to better fit some requirements of our application:

- Like in classical expert systems, preference statements will be elaborated with human pilots. We need to improve the expressivity of preference statements in cp-theories to facilitate knowledge elicitation.

- The computation time at system utilisation is critical, but we can mitigate this risk through off-line pre-processing.
- The set of preference statements to be used depends on the particular operational situation treated (use case).

The following notations are borrowed from (Wilson, 2011):

- Domain(X) is the domain of feature X
- If $U \subseteq F$, IU denotes the set of possible value assignments to features in U. So an outcome is an element of IF
- If $u \in IU$ and $X \in U$, u(X) denotes the value that u assigns to X
- If $V \subseteq U$, u(V) denotes the projection of u on I_V

The driving idea of our approach is to represent a preference statement as a vector of relations: A criterion on feature X is defined by a binary relation R on Domain(X) (e.g. X="flight time", R=shorter). A criterion defines a binary relation on outcomes: it includes all the pairs of outcomes whose X features are related by R. The following notation is introduced for criteria:

$$\mathrm{CRIT}(X,R) = \{(o_1, o_2) \text{ s.t. } o_1, o_2 \in I_F, (o_1(X), o_2(X)) \in R\} \quad (3)$$

The preference relation for a particular use case is defined by a set of preference statements. A preference statement is a binary relation on F, defined as the intersection of criterions, one criterion for each feature:

$$P = \mathrm{CRIT}(X_1, R_1) \cap \ldots \cap \mathrm{CRIT}(X_n, R_n) \quad (4)$$

With $F=\{X_1 \ldots X_n\}$ and each R_i denotes a relation on Domain $(X_i) \times$ Domain(X_i).

To ensure that preference statements are acyclic, it must be imposed that one at least of the criteria is irreflexive. To follow the spirit of CP-theories at statement elicitation, the default relation for a given variable is equality ("ceteris paribus assumption").

From the pilot perspective, decision assistant functions have to support the pilot in diverse diversion situations ("use cases"). Our preference knowledge base is naturally structured accordingly, i.e. it can be represented as a mapping which associates a pair {Fu, Qu} to each use case u, where Fu is the list of features relevant for u, and Qu is the set of statements to be applied in case u.

To reduce computation times in operation, the decision assistant uses a pre-processed knowledge base of preference statements. This knowledge base is computed offline, from the initial set of explicit statements augmented by its transitive closure.

More formally, the pre-processing step works as follows:

Let R_1 and R_2 be binary relations on the same domain D, $R_1 \bullet R_2$ is the product relation defined as (Bouyssou, 2005):

$$\{(x_1,x_3) \text{ s.t. } x_1,x_3 \in D, \exists x_2 \in D \text{ s.t.} (x_1,x_2) \in R_1 \wedge (x_2,x_3) \in R_2\} \quad (5)$$

Because preference relations are transitive, a new preference statement can be obtained by the product of two preference statements. The new relation $P_1 \bullet P_2$ is also a preference statement in our framework because it can be rewritten into the standard form, thanks to the following property. With $P_1 = \cap_{i=1..n} \mathrm{CRIT}(X_i, R_{1i})$ and $P_2 = \cap_{i=1..n} \mathrm{CRIT}(X_i, R_{2i})$:

$$P_1 \bullet P_2 = \cap_{i=1..n} \mathrm{CRIT}(X_i, R_{1i} \bullet R_{2i}) \quad (6)$$

The product preference statement is obtained by the conjunction of product relations at the level of each feature. Then it becomes possible to derive all the relevant preference statements by transitivity, based on the products of binary preferences.

The computing process requires that product operations have been defined to combine the binary relations for each feature. This supposes that binary relations on each feature can be combined to derive new valid relations. This property obtains easily for qualitative features as soon as any binary relation can be represented by a boolean matrix, and the product of relations corresponds to the product of their matrices (Bouyssou, 2005). For quantitative features, we have to restrict the relations used in preference statements to the disjunctions of basic =, <, > relations.

4 EXAMPLE APPLICATION ON DIVERSION ASSISTANCE

Reasoning about pilot preferences is only one module in an information processing chain whose objective is to push informed decision proposals towards the pilot. CRP are used here to formalise and to reason about pilot explainable rules to select a diversion airport among several candidate solutions. For example, if a passenger is sick, the following statements will apply: "a safe diversion flight is always preferred to a flight whose safety level is degraded". "Among two equally safe diversion flights, I will prefer an airport with medical services, provided the flight time is not much longer than to the other one", "among two reachable airports, if none of them have medical services, I will prefer the shortest time to get to the nearest hospital".

The functional architecture works as follows:

When a diversion is required, a short list of candidate airports for diversion is selected (typically, the few closest airports, including the ones which have been identified as possible diversion airports at flight preparation).

Flight plans are calculated for the airports of the short list to evaluate quantitative variables (time, distance,...), and diverse descent strategies.

The features needed to reason about the different solutions for the particular use case are calculated (e.g. quantitative to qualitative conversion, when relevant).

The solutions are ranked by using the logical framework described above.

Justified recommendations are sent back to the pilot, who can accept the first proposal or another one in the list, or ask for explanations, or ask to consider additional solutions.

In many real situations, a few more interaction loops will be needed (question answering, what if questions, requirements for more airports…), which do not change the principle of this functional architecture.

The Decision-Making analysis should be as close as possible to natural reasoning of pilots in operation. For example if a passenger is sick, the aim of the diversion decision is to land as fast as possible to an airfield where the passenger will be quickly attended by medical services. The pilot must first ensure flight safety, a safe diversion solution will always be preferred to an option where safety margins are significantly degraded, whatever the other features. Among the solutions where safety is ensured, the pilot will prefer airports with adequate facilities to take care of the sick passenger. Among the safe diversions to airports where the passenger can be attended, the diversion flights with minimum travel time will be preferred. The example also shows how the facilities about passenger handling can be taken into account.

In the case of engine fire, flight time has to become the dominant criterion as soon as the safety margins are degraded. Different descent strategies might also impact the final choice, which results in proposing the less bad solution from a compromise between degraded solutions.

In case of closure of the destination airport, beyond safety, the decision assistant has to take into account a different set of features, including more commercial and economic aspects. The preference statements to be invoked here consider the availability of ground support teams, the availability of services and commodities for passengers, and the impact on airline flight schedules.

5 DISCUSSION

We proposed a framework to model and reason about preferences which is derived from cp-theories (Wilson, 2011). In our approach, preference statements are handled as conjuncts of feature level criteria. This approach is suitable for the development of a pre-processing step of the knowledge base to improve on-line processing times. The new framework had to fulfil specific requirements for our application. In particular the language for statements is more expressive: it does not require to focus on a single feature for each statement; disjunctions are allowed in feature level criteria; preference statements are not limited to qualitative (or propositional) criteria; they admit limited usage of quantitative comparison. It can be demonstrated that our language for preference statements is more expressive than preference statements in cp-theories (cp-theories can be reformulated in our language).

How do CRP compare with classical numerical approaches to multi-criteria decision making? Those methods usually fall into two categories: the compare and aggregate approaches, and the aggregate and compare approaches (Gonzales, Perny 2020). Our approach could be classified as a compare and aggregate one: each preference statement operates a comparison at feature level; then the result is aggregated to decide if the statement triggers or not. Nevertheless, our approach differs from multi6criteria decision techniques on several aspects: each preference statement is an independent module: it uses its own rules to compare the features, and it considers only the few features which are relevant (the remaining features are assumed to be equal or indifferent). This modularity cumulated with the property that the language used is mostly qualitative, facilitates the elicitation of preferences by human experts. Quantitative multi-criteria approaches nevertheless have a strong competitive advantage:

they are suitable for automated learning of numerical functions for comparison and aggregation. But our problem is not well fitted for automated learning because of the diversity of use cases (each with its own list of relevant criteria), the scarcity of accurate and documented diversion data, and the dependency of decisions to airline policy and aircraft types. Moreover, the decisions proposed must be justified to the pilot. In summary, the framework presented in this abstract is operationally well adapted because it privileges knowledge elicitation with expert pilots, performance at run time, and explanation of the ranking of the solution proposed.

In the following steps, we will finalise the use cases and their preference statements with test pilots; we are also working on improving the explanation processes and developing the capability to customise the knowledge base, including by taking into account pilot's feedback during the interactions with the assistant.

REFERENCES

Bienvenu M., Lang J., Wilson N. (2010). From Preference Logics to Preference Languages, and Back. Principles of Knowledge Representation and Reasoning: Proceedings of the 12th International Conference, KR 2010.

Boutilier C., Brafman R., Domshlak C., Hoos H., Poole D., (2004). CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements, Journal of Artificial Intelligence Research (JAIR).

Bouyssou D., Vincke P. (2005). Relations binaires et modélisation des préférences. 2005.

Brafman R., Domshlak C., Shimony S. (2006). On Graphical Modelling of Preference and Importance. J. Artif. Intell. Res. (JAIR). 25. 389-424.

Geißer F., Povéda G., Trevizan F. W., Bondouy M., Teichteil-Königsbuch F., Thiébaux S.(2020) Optimal and Heuristic Approaches for Constrained Flight Planning under Weather Uncertainty. ICAPS 2020: 384-393.

Gonzales C., Perny P. (2020) Multicriteria Decision Making. In A Guided Tour of Artificial Intelligence Research, I, Springer, pp.519-548, 2020, Knowledge Representation, Reasoning and Learning, ed. Marquis, P; Papini, O.; Prade, H.

Kaci S., Lang J., Perny P. (2020). Compact Representations of Preferences. In A Guided Tour of Artificial Intelligence Research, I, Springer, pp.519-548, 2020, Knowledge Representation, Reasoning and Learning, ed. Marquis, P; Papini, O.; Prade, H.

Khannoussi A., Olteanu A.-L., Labreuche C., Narayan P., Dezan C., Diguet J.-P., Petit-Frère J., Meyer P. (2019). Integrating Operators' Preferences into Decisions of Unmanned Aerial Vehicles: Multi-layer Decision Engine and Incremental Preference Elicitation.

Mueller S., Veinott E., Hoffman R., Klein G., Alam L., Mamun T., Clancey W., (2021). Principles of Explanation in Human-AI Systems.

Shoham Y. (1997), Conditional Utility, Utility Independence, and Utility Networks Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI 1997).

Wilson N. (2011). Computational techniques for a simple theory of conditional preferences Artificial Intelligence Volume 175, Issues 7–8, May 2011, Pages 1053-1091.