

# Meeting a Competitive Challenge in the Frequent Flyer Competition

*Brian Ebersold*

In January 1990, a major airline made a promotional offering to its frequent flyer members: "Fly three round trips from a hub city, and earn a free round-trip pass." American Airlines, wanting to gain a competitive edge, reacted rapidly by offering its frequent flyer customers a more progressive offering. Marketing offered a free round-trip domestic pass to any AAdvantage customer flying three round trips anywhere in the entire American Airlines route network, the awards to be mailed in May 1990. Round-trip identification had not previously been solved by a computer automation method because of complexities introduced by inconsistent or missing data. Historically, airlines have based the calculation of frequent flyer awards on the number of segments, segment miles flown, or specific (that is, hub) city activity. Offering round-trip promotions held the potential for opening a new era in frequent flyer competition and tipping the edge toward the airline that could make the offer.

In late January 1990, the Sabre Computer Services department initiated a joint effort (traditional data processing staff and knowledge based systems staff) aimed toward developing a system that could

identify round trips. This system, a response to a competitive market situation, was intended for use against the frequent flyer customer database. The effort consisted of the traditional application development support group defining and developing the methodologies for the extraction of data and the compilation of award certificates, and a knowledge engineer from the knowledge system group developing an expert system solution for the round-trip identification problem.

Ultimately, an expert system that resolved the round-trip identification task was designed, developed, and implemented by the May 1990 deadline. The problem required the pattern processing of, and inferencing on, nearly 10 million data records. The objective, automating the identification and classification of a variety of round-trip types and minimizing a need for human intervention, was achieved.

### **Problem Statement**

The marketing department is supported by a dedicated data processing support staff that in the past has provided exclusive service to the user's data processing requirements. In January 1990, the staff members were faced with a problem of enormous proportions, with a short time in which to respond. In short, they needed to develop and deliver a system that could process 10 million customer records with greater than 98 percent accuracy in less than 3½ months. Within 3 months, the promotion would be over, and customers would be anticipating the delivery of their travel awards. The initial forecast by the marketing department was that 7 to 9 million flight segments would be traveled by frequent flyer members during the 3-month period (the actual number of segments flown was 10 million). Within this short time period, a suitable round-trip strategy would need to be defined, developed, and implemented.

The data processing staff supporting this marketing group is a conventional application development group of medium size (10 to 12 people). Typically, staff member time is spent developing methods to perform extracts from mainframe DB2 files and compiling the results in other DB2 or flat-file output. The output files generate the activity and award reports that are used by the frequent flyer program and sent to the customers. The current promotion would utilize the majority of the programmer resources in developing the extraction and award software. Data processing estimated the time required to design and build the round-trip identification and categorization software at 10 months but only if resources (staff) were available to support the development effort. A decision was made to proceed with a radical departure from the past and incorporate an expert system solution.

A developer from the knowledge system group was assigned to exam-

ine the round-trip identification and classification problem. Initially, a similarity with another expert system, one that identified markets, was assumed. An initial prototype was completed within 24 hours and a determination made that although round trips were not similar to markets (in the context of frequent flyer travel), the problem showed the potential for incorporating a rule-based solution. A second prototype, completed within 2 weeks, confirmed that an expert system was capable of identifying and classifying the majority of the segments in question. This prototype ran against actual frequent flyer member data and classified the segments into 1 of the 4 or 5 round-trip types that were then currently known. Another important result of the prototyping effort was discerning that a delivery time of 2½ months was feasible. The primary issues raised during prototyping were processing requirements and hardware capabilities.


The definition of a round trip is not as simple as flying from point A to point B and back again. A flight segment in the airline industry is often referred to as a *leg*. A pure *round trip* is one in which all the legs flown during an itinerary can be linked in a sequential order, leaving from, and returning to, point A (figure 1).

The *open jaw* (from the appearance of the itinerary when drawn on a piece of paper) is a second type of round trip that is ambiguous and less self evident: A customer visits a number of cities during the course of his(her) itinerary, and the possibility exists that either the data for one or more of legs are not reported in a timely manner, or a flight originally booked was not taken (that is, one or more segments of a pure round trip is missing).

Linking the open jaw segments posed the greatest challenge in the round-trip identification problem. Specifically, when a customer had a number of open jaw itineraries that occurred sequentially in his(her) account record, it was difficult to determine when one trip ended, and a new one began. Dates could not necessarily be used as a constraining factor because of the wide range of travel dates that were eligible to qualify for the round-trip promotion (for example, travel out on the first day of the promotion and back on the last day of the promotion could qualify as a round trip). Reservation numbers (record locaters) also could not be used exclusively to constrain the problem because the possibility of a customer booking a flight out and back at different times (with different record locaters) was real.

The following distinct characteristics of this problem lent themselves toward using an expert system solution:

First, the problem had a well-defined and limited scope (classify all flights for an individual account activity from 15 January to 15 April). In addition, there was a drop-dead date (15 May) limiting how far the



*Figure 1. Round-Trip Types.*

customer and developers could press the breadth and depth of the knowledge modeling.

Second, distinguishing patterns were evident in the data when flight legs were linked.

Third, there was a distinct concept of what the answers should look like

(although some round-trip characteristics were formulated just prior to deployment).

The primary issues encountered were (1) the accurate capture of the business rules (because this problem was new, previous specific expertise didn't exist); (2) processing time (the volume [10 million] of records would require deployment in a mainframe environment, and the knowledge system group had no previous experience with mainframe expert systems); (3) high accuracy (a 1 percent error rate could generate 10,000 inquiry calls; operations would potentially be incapable of handling a rate greater than 2 percent); (4) fragmented data resources residing in a number of different environments, including real-time and commercial processing mainframes (figure 2); and (5) rapid development to leave adequate time for system test and integration.

## Approach

This system was developed on an IBM/PS2 Model-80 personal computer and ported directly to an IBM 3090 mainframe. The development profile fits that of the rapid application development model. For example, four distinct prototypes were developed, all somewhat similar but using significant strategy changes or enhancements. The customer was constantly revisited with the processing results for knowledge validation and refinement. The business rules modeled during these iterative processes became more distinct and simple as the knowledge-acquisition and knowledge validation phases proceeded.

The ART-IM expert system shell was chosen as the development and, ultimately, the deployment environment. The knowledge engineer who developed the expert system was familiar with the ART-IM development environment. The ART-IM tool seemed to be a fairly good match for the data representation (frames) and inferencing technique (rules). Also influencing the decision was the fact that Inference Corporation had a mainframe tool available that allowed portability from the personal computer development environment.

By working interactively with the customer, the business rules were modeled. By comparing the results they expected with the results given by the expert system (running against a control data set), the differences narrowed. The criteria for the types of round trips that hold the greatest significance were prioritized; this information was then incorporated in the rule base so that the expert system would seek the patterns of greatest significance first.

Recognizing the (itinerary) topographies (figure 1) that would possibly qualify as round trips was the single greatest problem in solving the

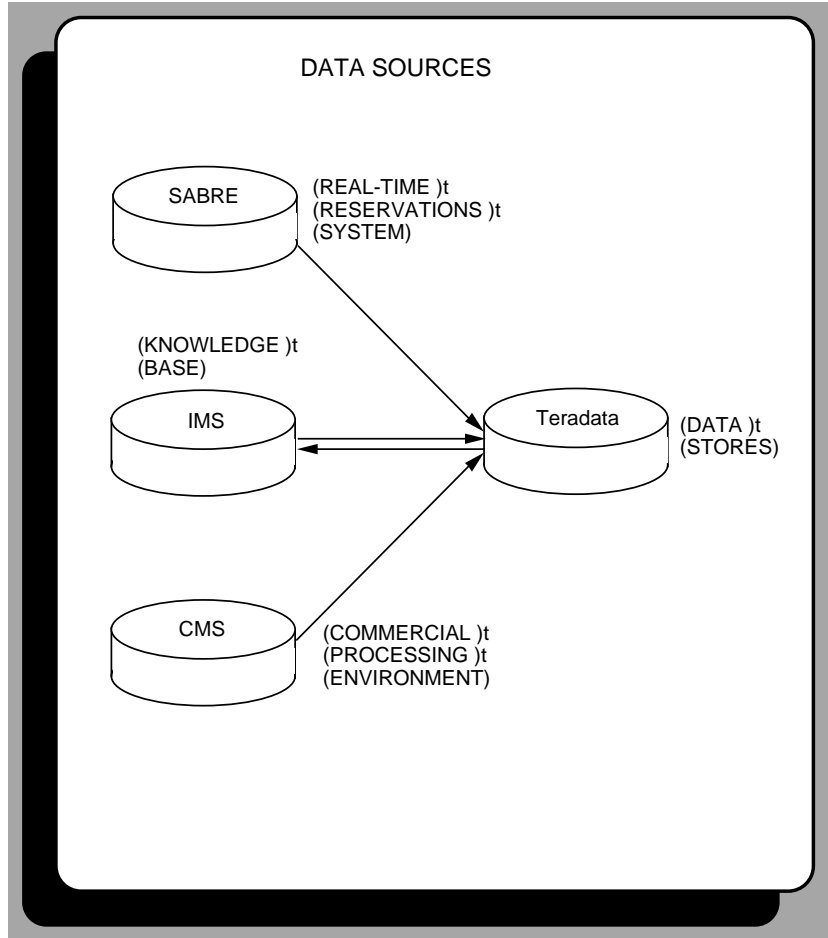


Figure 2. Data Resources.

identification and categorization problem. The two distinct superclasses of round trips, pure and open jaw, contain eight subclasses, which were eventually identified by the customer and the expert system. Each account was individually resolved using a combination of date, reservation number, chronology, and airport constraints. The result was an account activity log for the flyer that categorized and classified all the flight activity.

The system was designed to use the minimal number of rule firings to resolve each flyer account. Rule firings were at a premium because of the magnitude of data that required processing. A total of 20 rules were deployed, many of which were complex and multifunctional.

Most of the configuration parameters are user definable and picked up by the system as variables. This approach empowered the user greater flexibility to modify system behavior (either liberalize or constrain the expected solutions).

Fixed-format data files residing on the mainframe provided the simplest way to interface the data to the knowledge base. Data initially came from a variety of databases and a real-time reservation system, then were compiled into the knowledge base data-input files. A joint decision was made by knowledge systems and data processing to create files with a working set size of 450,000 customer records each, making a total of 22 input files.

### **Analysis**

The analysis is performed in four basic phases: The first phase establishes a chronology for the occurrence of events (flight segments) (figure 3). The second phase ties the plausible events based on origin, connection city, or destination points. The third phase classifies the event strings based on the topography by using pattern-matching heuristics. Pure topographies have the highest confidence, thus the highest priority. Open jaws have lower confidence and precedence values (figure 4). The fourth and final phase disconnects illogical solutions and reinstates the second phase using an altered group of events. When inferencing is concluded, all flight segments in the account are grouped or categorized.

### **Development Cost**

The internal development costs were calculated at approximately \$62,000 for the knowledge engineer to develop and deploy the expert system and \$150,000 for the data processing programmers compiling the data sources and developing the software to perform the award identification. The total cost was approximately \$212,000.

### **System Use**

This application has been in use since May 1990. Its initial run was against a full 10 million customer records and took approximately 122 hours of central processing unit time on an IBM 3090 computer. It was subsequently run as a monthly batch process until the promotion expired, allowing the expert system to reevaluate customer accounts with previously missing or unaccounted-for flight activity data. The reevalua-

E622712AA 893430232LAX	DFW	F	0100 0547	BAIHW
E622712AA 900030924DFW	CLE	F	0828 1204	NWMBWA
E622712AA 893381289DFW	PDX	Y	1600 1816	KPABWB
E622712AA 900050901CLE	DFW	Y	2058 0032#1	NWMBWA

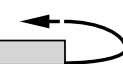


Figure 3. Sequencing.

Round trip #1 :					
AA 1289	DFW PDX	4DEC89	1600 1816	O4	Y KPABWB
AA 232	LAX DFW	9DEC89	100 547	O4	F BAIHW
Round trip #2 :					
AA 924	DFW CLE	3JAN90	828 1204	P	F NWMBWA
AA 901	CLE DFW	5JAN90	2058 0032	P	Y NWMBWA

Figure 4. Categorizing.

tion runs were only applied against frequent flyer records with effected data and typically ran in minutes.

## Payoff

The revenue implications for this effort are difficult to calculate. One would need to determine the revenue damage that would have occurred if the competitive offering had not been responded to. However, American Airlines did post its best first-quarter revenue results in the company's history, and the number of segments flown during the qualifying period was 10 percent more than had been forecast by the marketing department.

Planned future use seems assured. Airline marketing is cyclical by nature, and it typically makes use of all the available marketing options at its disposal. Knowledge of what a round trip is has now been codified, and the expert system can be called on to assure this knowledge is applied (on demand) in a consistent fashion; otherwise, the application of the round-trip policies would be subject to the individual interpretation of each customer service agent. Marketing has a unique and powerful new tool that they can leverage to offer travelers attractive new incentives for brand loyalty. Provisions for future enhancements, notably an account fraud-abuse strategy, were also incorporated into the system design. This benefit also results from the pattern processing and rule techniques used.



## Conclusions

This mainframe expert system project produced the following significant results for American Airlines: (1) Round trips can now be identified and categorized for the frequent flyer program. (2) The use of an expert system shell was proven to be a dramatic productivity enhancer, with a reduced personnel requirement to deploy and maintain an application. (3) An expert system was successfully developed on a workstation platform and deployed on the mainframe (the development could have been accomplished entirely on the mainframe). (4) The system required minimal maintenance interventions by the developer or data processing staff. The scope of the knowledge domain was correct, and the flexibility to manipulate the results was incorporated into the design. (5) Marketing can offer more competitive incentive programs now that the automated round-trip system is in place. A situation occurred in late 1990 when another competitor made a promotional offering to its customers. This promotion was based on six completed segments of travel being awarded a free trip for future travel. American Airlines was able to respond with a promotional offering that included round trips or one-way trips (a subset of round trips), the award to be based on the most favorable conditions for the passenger.

## Prolog

The round-trip identification expert system continues to be used in support of marketing's promotional offerings in the frequent flyer program. No significant changes or modifications have been made to the core knowledge base, although some heuristics were modified to support different strategies applied during the account processing. At the time of the initial writing of this chapter, the application had been used to support one promotion. Since this time, in the spring of 1991, another competitive response promotion implemented by marketing made use of the round-trip expert system.

## Acknowledgments

My thanks go to the following individuals: Chuck Harvey, Data Dallas Inc.; George Geotz, director, AAdvantage Customer Services, American Airlines; and Mark Reynolds, Tim Robe, and Pippa Hearn, Sabre Computer Services, American Airlines.