

## STHANA

# Profitability Forecast and Situation Analysis for Automated Teller Machines

Cyril Way

ISoft

Chemin de Moulon

F91190 Gif sur Yvette, France

cyril@isoft.fr

### Abstract

The French credit card system makes it highly profitable for banks to have heavily used Automated Teller Machines (ATM). "La Caisse d'Épargne", one of the major French bank, manages 5000 ATMs all over France. The goal of the Sthana system is to capitalize the knowledge spread all over the company into a system capable of issuing recommendations for existing ATM's and capable of forecasting a new ATM's activity. Sthana uses Data Mining and Case-Based-Reasoning techniques so as to extract information from existing data (including economic, geographical and internal bank data) and from the bank's ATM experts. The system builds up classifications on high level descriptors from raw data and eventually indicates a measure of the ATM's activity and profitability, highlights factors which could lead to higher profitability or pinpoints the ATM's vulnerabilities. An object-oriented model coupled with an extremely modular system allows the data and rules to be customized for geographical units of the bank. Sthana has been deployed and customized for different geographical units, but the knowledge base is centralized in the bank headquarters in Paris.

### Problem Description

In France, ATMs help the bank save money but can also bring in cash. It is thus a very strategic decision to install an ATM. Lots of different factors can modify its profitability. Decisions are generally in the hands of experts. The goal of the Sthana system was to elicit these experts' knowledge and implement it into a system capable of analyzing ATMs' positions. The design of the system is based on years of data collected about ATM activity and on experts themselves. Data available about ATMs came from various origins:

- Geographical:  
Such as population density, size of nearby towns, etc.
- Economic:  
Such as unemployment rate in the area, average household income, distribution of the population in professional categories, etc.

- On-the-field data:  
Such as the ATM signs and lights, nearby shops and activities, position of the ATM in the street, competitors' ATMs, etc.
- Internal Bank data:  
Such as the number of withdrawals by the bank's customers, by other banks' customers, daily and hourly activities, etc.

For the system to be usable by end-users, some constraints were set:

- Geographical customization of rules and data:  
"La Caisse d'Épargne" is organized as a network of geographical units. Each of these units comprises between 50 and 250 bank agencies, thus maintaining a few hundred ATMs each. Distinctive local features of each unit must be taken into account. Seaside ATMs cannot be described the same way nor compared to suburban ATMs. Thus both data and reasoning should be customizable to answer specific units' needs.
- Nationwide deployment with centralized management:  
Each unit is independent with regard to hardware or software choices but a nationwide entity, called CENCEP ("Centre d'Étude National des Caisses d'Épargne Populaires"), helps synchronize developments and reduce costs by implementing a nationwide policy. The Sthana project was initiated by the CENCEP and proposed to geographical units as a service. Thus the knowledge base and the resolution rules would be maintained in Paris for all units.

The run-time system should be able to meet several distinct needs:

- Audit of existing ATMs:  
Weaknesses  
Risk factors  
Improvements to be considered
- What-If functionality:  
What if a competitor installs an ATM next to mine?  
What if the supermarket closes?
- Forecasting:  
We plan to install a new ATM there. What level of profitability can we expect?

---

Copyright © 1997, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup> In Sanskrit, "Sthana" represents the principle of position in the science of numbers.

Each of these elements should be considered along two dimensions ([Tocatlian 1996]), which are essential for evaluating the ATM's profitability:

- **External activity:**  
Each time a customer of another bank withdraws money from one of the bank's ATMs, this customer's bank pays a commission of 5 French Francs (approx. \$1). If the ATM is very successful and attracts numerous external cardholders, the ATM can make money by itself.
- **Internal activity:**  
ATM can be used for withdrawals by any cardholders but can also be used by the bank's own clients for other operations such as consulting one's account, transferring money between accounts, etc. These operations reduce the man-teller workload and can thus also generate profits. Evaluating this kind of profitability is harder than the previous type, but experts commonly agree that an operation other than a withdrawal generates on average 2 French Francs in savings. Proposing a wide range of internal services to one's customers is also a way of lowering the number of customers using other banks' ATMs, thus reducing the commissions paid to other banks.

This profitability should be compared to the ATM's costs so as to determine its overall profitability:

- **Installation costs:**  
ATM are generally installed downtown where the real-estate prices can be very high. The bank can either rent the location on a long lease or buy it. Installing the ATM always requires labour and this adds to the initial cost.
- **Operating costs:**  
The ATM must be loaded and cleaned regularly. This is generally not done by the bank itself but a security firm. This cost is recurrent.

The decision of installing an ATM really is a long-term decision, because of the initial costs, which can range up to 500.000 French Francs (approx. \$100.000) ([Pinsky 1996]), but also because of the dynamics of ATMs. Some of them can take a year and a half to find their customers. This is particularly true in residential areas where the people's habits play a strong role in their behavior.

## Application Description

### General Architecture Overview

The whole application is based on comparisons. This approach is close to that of a human expert reasoning by analogy: "This ATM reminds me of the ATM in this situation and of this one, except that in this case, the ATM has fewer signs advertising its presence. So we might

improve this ATM's profitability by increasing the number of signs in appropriate places". In order to make useful comparisons, the first step is to determine whether two ATMs are comparable or not. The first reason is that two ATMs in very different environments cannot be compared, and the second reason is that two ATMs in a given environment cannot be compared the same way as two ATMs in another environment. This comparability feature is assessed through various classifications into ATM "prototypes" uncovered during the design phase. One advantage of these classifications is that they give qualitative intermediate results to the user, who can then "follow" the reasoning of the system step by step. Once the ATM has been classified, it is compared to ATMs belonging to the same prototypes using the comparison scheme established for this kind of ATM. Various recommendations can then be issued and quantitative elements representing the ATM's profitability are given. The recommendations are derived from simple adaptive mechanisms such as k-Nearest neighbors.

The whole system was made modular so that each step of the reasoning can be customized. In figure 1, the objects in the right-hand part (High level descriptor methods, ATM prototype & comparison schemes and known ATMs) are considered as application resources. They can be changed at any time without disturbing the overall application workflow. Each geographical unit is responsible for ensuring that the various components they use make sense:

- High-level descriptor synthesis methods are based on decision trees and on simple aggregate functions. They can thus be redefined without special technical skills. A decision tree allows rules to be defined very fast and based on facts. Interactive tools with a graphical user interface make it really simple for the ATM expert<sup>2</sup> to create, edit and save any tree for temporary or definitive purposes.
- ATM experts see comparison schemes as weights put on attributes. Each of those weights simply represents the importance of the attribute in the comparison of two ATM in the same prototype. The definition of ATM prototypes is not available to ATM experts for changes.
- The known ATM base is the core layer on which the whole system is based. They are used to infer new decision trees for high level descriptors synthesis (see above) and, most important, they are used to make comparisons with new ATMs or with existing ATMs which need improvements.

---

<sup>2</sup> See next part for a definition of an ATM expert.

The diagram 1 illustrates the application workflow for prediction on a new ATM:

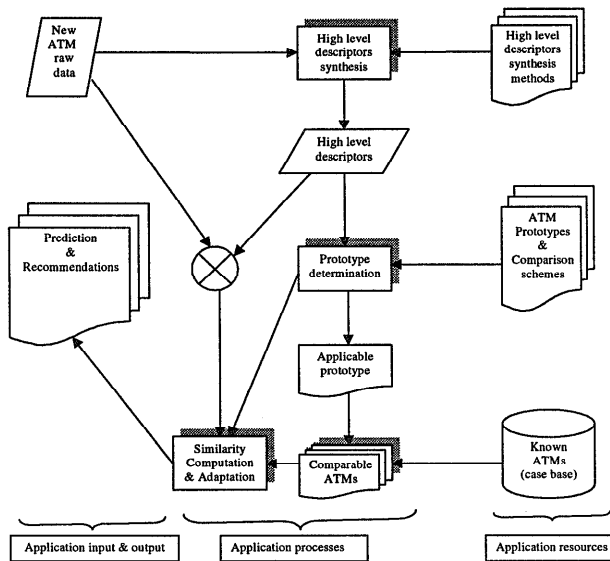


Diagram 1: Application workflow for prediction

### High Level Descriptors

Raw data about ATMs is synthesized in order to derive higher level descriptors. The descriptors are then used to derive higher descriptors and so on, till the final classification (which concludes at the ATM prototype) is reached. Intermediate results are stored because they are also very useful for the next phase of the process. All the information, from raw data to high level descriptors, is mapped and stored in an object-oriented model. This multi-layer classification allows easy customization of the reasoning process so that geographical particularities can be easily integrated and maintained into the system. For example, here are three high level descriptors that were used in the analysis:

- **Volume:**  
This is a qualitative measure of the number of people moving around the ATM, cardholders or not, adults, children, etc.
  - **Potential:**  
This is a qualitative measure of the proportion of people wanting to withdraw money on the spot among all the people moving around.
  - **Attractiveness:**  
A qualitative measure of the proportion of people withdrawing money in the bank's ATM among all the people wanting to withdraw money.
- It is of great interest to consider these three dimensions

because environmental elements have markedly different influences on them. The table 1 summarizes the influence of three elements on the three descriptors:

	Volume	Potential	Attractivity
A competitor installs an ATM-wall <sup>3</sup> nearby	No change	Up	Down
A supermarket opens in the same street	Up	Up	No change
Couponing <sup>4</sup> is installed on the ATM	No change	No change	Up

Table 1: Examples of influences on high level descriptors

For example, if a competitor installs an ATM wall, the location will be known in the neighborhood as offering easy withdrawals. The location will thus attract many more cardholders than previously. By definition, the potential of the location increases. Compared to the total number of people moving around, these cardholders are not very numerous, so the volume of the location does not change significantly. The fact that a competitor has installed several ATMs nearby will dramatically decrease the proportion of people using the first ATM; consequently its attractiveness decreases.

The high level descriptors were chosen because each of them represented a specific aspect of the ATM position, and because they were thought to be roughly independent. In the previous example, we can see that Volume, Potential and Attractiveness have very different responses towards three examples of events; they are independent. Of course, a volume always generate some potential (equal to the mean percentage of cardholders in the population of the location), but these dependencies are known and do not change over time with the same time-scale as other events considered in the study (they are stationary).

### ATM's Prototype Classification

The high order descriptors eventually lead to a classification. ATMs are then grouped into homogenous subsets, called prototypes. The goal was to isolate ATMs within the same environment, for which global influences could be suspected. ATMs in a prototype are not necessarily similar, they are just comparable.

<sup>3</sup> An ATM-wall is composed of two or three adjacent ATMs.

<sup>4</sup> The ATM can distribute reductions for parking lots or shops on the customer receipt

Here are three examples of prototypes, their description in terms of high level descriptors and the dynamics between them.

- **Strong & Risky**  
The location appears to be fairly good. There are no competitors, or our position out-ranges theirs. Volume and Potential are medium or strong; Attractiveness is excellent. This position might be subject to attacks by competitors, since it is very attractive.
- **Weak**  
The location is not a good one, so bad that there may even be no competitors. Volume is low; Potential is also low. The position is stable in the short term.
- **Competitive**  
The location is very good, and there are many competitors around. Volume and Potential are both quite high, and Attractiveness is medium or weak. The position is stable. Some ATMs in this prototype have quite a good profitability, whereas for others it may be very poor.

Of these three examples, the first prototype is the least common, since it is an unstable position. The third one is the most frequent, and the second represents one third of the first. An ATM may change prototype over time, in response to events in their environment. The following diagram illustrates such transitions between prototypes, triggered by various events.

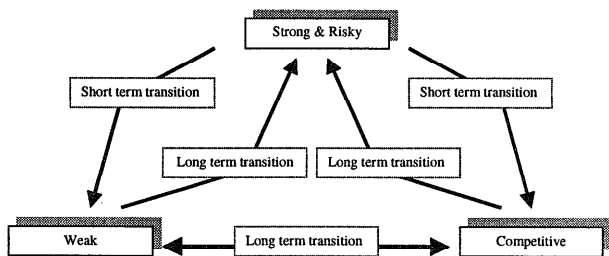


Diagram 2: Transitions between Prototypes

Examples of short-term transition include tactical changes in the position (a competitor's attack, different sign-postings, creation of parking lots, etc.) and minor external changes in the environment (a shop closes, security is reinforced, etc.). Examples of long term transition include major external changes such as when a whole shopping center closes, or when the traffic radically changes (creation of several lanes, change into Single-way Street, creation of a pedestrian area, etc.).

## ATM Comparison

Within a prototype, comparisons between ATMs make sense. During the design phase, comparison schemes were defined.

They assign a weight and a comparison function to each raw piece of information and high level descriptor. The function is used during run-time to assess the similarity between each descriptor's value for two ATMs and the weight is used to aggregate these descriptor similarities into an object similarity. Then the relationships between objects are used to compute an overall similarity between the two situations ([Way 1994]). The object-oriented approach makes it easy to compare. An OO model relies on hierarchies, which can be understood as hierarchies of comparable objects. The comparison between two objects (a supermarket and a small shop, for instance) is determined by the most specialized object being an abstraction of the two objects to hand (retail, in our example). When two objects' generalization is the root of the OO model, they are not comparable. This approach also allows easy extension of the model for geographical particularities. When new objects are defined, the comparison scheme needs only be specialized for the newly created objects, without changing other algorithms or other object comparison schemes.

## Underlying AI Technologies

The whole system is the result of three advanced AI technologies:

- **Object Orientation:**  
ATM descriptions must be customized to fit geographical particularities. Furthermore, treatment of data depends heavily on classifications performed on ATMs. Object Orientation fits these requirements perfectly and has numerous other advantages, which made its choice unavoidable. For example, geographical diversity makes it impossible to predict all the situations that will be needed when the application is deployed all over France. Object orientation allows one to postpone the definition of new objects describing these situations and all treatment for the objects without blocking the design or the run-time of other parts of the system. The whole system can then be deployed and enhanced incrementally.
- **Decision Trees:**  
Many different technologies can be employed to make classifications, such as neural nets, genetic algorithms (see [Fayyad 1996], [Mauth 1995]), etc. Decision trees had key advantages over these two techniques for the Sthana project. First, they can be understood by anyone, and this feature was absolutely necessary during the design phase when domain experts had an important role to play. Secondly, they give access to the final results but also to explanations, and this is crucial when the system has to issue recommendations

as well as predictions. Another key advantage is that they handle different types of data (continuous, symbolic, numeric) without hard-coded pre-treatment (see [Breiman 1984], [Quinlan 1993] and [Quinlan 1996]) which was very useful as data came from a number of different sources. Decision Trees are now widely used in the banking sector ([Way 1996]).

- **Case-Based Reasoning**

This techniques suits well in the domain experts' way of reasoning, which relies heavily on comparisons as we have seen ([Kolodner 1993], [Riesbeck 1989]). This approach was favored as the one giving the most understandable results for non-experts. Moreover, it facilitates tremendously the maintenance of the knowledge base, which is reduced to comparison schemes and raw data. Raw data constitutes what is called the Case Base and can thus be updated, replaced partially or entirely very easily. Comparison schemes are available to ATM experts for simple updates.

## Application Use and Payoff

Two types of users were defined:

- End-users are connected to Sthana in read-only mode only. They cannot update the case base, nor insert new ATM. However, they can submit new ATM positions for evaluation purposes. They cannot modify any of the application resources. These end-users are generally bank agency managers, who want to assess a new ATM position before submission to the ATM expert. For the moment, Sthana is not available directly in bank agencies; it is only available in geographical units via traditional PC workstation.
- ATM Experts can modify and customize application resources to fit their particular needs and adapt them to the geographical particularities. They are responsible for maintaining the application consistency, that is to say the adequacy between high-level descriptor synthesis, comparison schemes and the case base. They cannot redefine the nature of high level descriptors nor the core mechanisms for comparison and adaptation. These kinds of changes are considered as application design changes and can only be performed by developers at the request of ATM experts. There is generally one ATM expert per geographical unit, who is in charge of deciding whether and where to install a new ATM, at the request of bank agency managers.

This application was deployed in September 1996 in the first geographical unit, in the North of France (More than 150 ATMs). In October, Sthana was also deployed in the Paris unit (around 350 ATMs). By the end of 1996, the application was deployed in 15 geographical units (total

2500 ATM). It is being heavily used in the first time by agency managers to help them diagnose and improve their ATMs' profitability. The tool is also used regularly by the person in the geographical units in charge of deciding if an ATM should be installed and, if so, where. When Sthana was installed in the North of France, experts had just decided on the installation of an ATM in a very expensive location. The initial cost for installing this ATM was 500.000 French Francs (approx. \$100.000). When they used Sthana with this ATM, the diagnosis was that this ATM would indeed be very profitable, but would also be very risky. A slight change in the environment could reduce the ATM's visibility to zero, and would thus allow an easy attack from a competitor. As a result of this diagnosis, the experts re-examined the ATM's situation and decided not to install it. Nobody can tell what the exact result would have been but executives agreed that this kind of risky position did not correspond to their general policy concerning ATM installations even if highly profitable at first sight.

So as to determine the accuracy of Sthana, blind tests were performed with ATMs that had not been added to the initial knowledge base. The accuracy measured for these ATMs attained 700 withdrawals per month, while it is commonly accepted that human experts generally attain 500 withdrawals. Compared to the number of operations per month (which ranges from a few hundred up to 13000), this accuracy can be considered as pretty good. Indeed, when overall profitability was considered, Sthana obtained very good results, as shown in table 2.

Reality vs. diagnosis	Sthana diagnoses failure	Sthana suggests postponement	Sthana diagnoses success
Non-Profitable ATMs	16	5	2
Profitable ATMs	3	19	55

Table 2: Reality vs Diagnosis Results

One hundred ATMs were tested, for which the bank knew whether they were profitable or not. If they had been submitted to Sthana, 16 non-profitable ATMs out of 23 would have been rejected, 5 would have been postponed and only two would have been accepted. On the other hand, Sthana would have rejected only 3 profitable ATMs. This gives a 5% error rate, a 24% indecision rate and a 71% success rate. This means that, even if Sthana does not predict the exact figures very accurately, its final diagnosis as to installing an ATM or not can be considered as highly reliable.

## Application Maintenance

The whole application was designed to facilitate maintenance and favor the emergence of corporate knowledge around ATMs. New ATMs inserted in geographical units' case bases are centralized in Paris and distributed to all units. Thus, each unit can make comparisons with ATMs from all over France or from selected parts of France only. Application resources other than case bases are considered as volatile, and are retrieved centrally and stored for documentation purposes only.

The knowledge base is not explicit. Instead, it is disseminated through various case bases and decision trees, which can also be seen as an organization of data dedicated to decision making. The central Paris location maintains an "official" STHANA resource set, which can be distributed to geographical units on demand. This is useful when a geographical unit wants to synchronize its process after modifying and updating parts of the process too heavily to be able to restore it again.

The various technical choices were made with this distribution process in mind. Object Orientation, Case-Based Reasoning and Decision Trees greatly facilitate the easy maintenance.

## Acknowledgements

The author of this document would like to thank the following teams and individuals for their great help over the course of the project:

- ISoft Product Development team, for their constant support
- René Lefébure, La Caisse d'Épargne de Flandre, for all the knowledge he accepted to reveal
- Jacques Schuhmacher and Philippe Quentin, La Caisse d'Épargne CENCEP, for the trust they constantly put into this project and its authors

AC2 and ReCall are trademarks of ISoft S.A.

## References

- Breiman, L., Friedman, J. H., Olshen, R. A., Stone, C. J. 1984. *Classification and Regression Trees*. Belmont, CA: Wadsworth.
- Fayyad, U., Piatetski-Shapiro G., Smyth P., Uthurusamy, R. 1996. *Advances in Knowledge Discovery and Data Mining*. AAAI Press.
- Kolodner, J. 1993. *Case-Based Reasoning*. Morgan Kaufmann

- Mauth, R. 1995. *Endless Search*. Extra International, Byte Magazine, November 1995.
- Pinsky, D. 1996. *Flandre: l'Épargne fait son Data Mining*. Informatiques Magazine, October 1996.
- Quinlan, J.R. 1993. *C4.5 : Programs for Machine Learning*. San Mateo: Morgan Kaufmann.
- Quinlan, J.R. 1996. *Improved Use of Continuous Attributes in C4.5*. Journal of Artificial Intelligence Research 4 (1996) 77-90.
- Riesbeck C. K., Shank R.C. 1989. *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates.
- Tocatlian, É. 1996. *Optimiser les automates bancaires*. Banque Stratégie, February 1996.
- Way, C. 1994. *Similarité structurelle pour le raisonnement à partir de cas*. Rapport pour l'École Nationale Supérieure de l'Aéronautique et de l'Espace. Toulouse, France.
- Way, C. 1996. *Decision Trees for Credit Assessment*. In Proceedings of the First International Conference of Practical Aspects of Knowledge Management. Basel, Switzerland.