
Software review: Rules-based engines or statistical optimisation: The intelligent way forward

Received: 8th July, 2002

Shaun Doyle

is Vice President, Intelligent Marketing Solutions, SAS. In this role he works with various parts of the SAS organisation to develop business-oriented solutions for marketing, particularly the SAS Marketing Automation (MA) Solution.

He was the co-founder and chairman of Intrinsic, a campaign management software vendor acquired by SAS in March, 2001. He has extensive experience in the design, development and implementation of customer-focused database marketing systems in a number of business sectors, and has helped design and build more than 150 marketing databases for blue-chip enterprises in financial services, retail and telecommunications, and other industries as well.

He is well known in the marketing industry and a regular participant in leading industry conferences. He is also the author of a number of papers on the application of database marketing that have been published in the UK and elsewhere.

Earlier in his career, he co-founded Berry Consulting, a firm specialising in database marketing.

Abstract The paper explores the advantages and disadvantages of rule-based engines and optimisation techniques in supporting marketing communication decisioning.

INTRODUCTION

In the past, determining which marketing communication a customer received was easy as the choice of channel was very limited. With the growth in active (eg e-mail) and passive (Web) communication channels and the increase in the volumes of campaigns facilitated by marketing automation technologies, however, this simple choice is becoming more difficult to make.

To address this complexity two new classes of technology are emerging — optimisation and rules engines. This paper explores some of the advantages and disadvantages of rules-based engines.

RULES-BASED DECISIONING

This type of technology breaks the communication decision into one or more simple rules that when applied determine the most appropriate marketing communication(s) for a

customer. It does not use any specific statistical process, although one or more of the variables used by the rule(s) may be behavioural model scores.

DEFINITIONS

The following section describes rule-based engines in more detail:

A rule is the smallest unit of business logic. It generally captures a test condition and identifies the action to take, based on the test evaluation, eg if a customer is male then send communication.

A rule set is a collection of rules, generally grouped together because they relate to a common task, eg if a customer is male and aged between 20 and 30 years then send communication.

A rule flow shows the order in which a particular rule set or series of rule sets are executed. These rule flows are normally represented graphically. The

Shaun Doyle

VP Intelligent Marketing Solutions, SAS, 63-64 Shrivensham Hundred Business Park, Watchfield, Swindon SN6 8TY, UK.

Tel: +44 (0)1628 486933;
Fax: +44 (0)1793 784748;
e-mail:
shaun.doyle@suk.sas.com

order in which rules or sets of rules are applied can significantly affect the outcome. For example, if gender is male, age is 25 to 30 years, marital status equals married (rule set 1) then if product code equals loan, and loan status equals open (rule set 2) send standard cross-sell communication.

Many of the new rule-engine technologies are based on object-orientated programming techniques, in which case the concepts of *classes and objects* are used to represent various business entities and ease rule definition. They also allow the rule developer to connect the rules with some outside calling application and pass information back and forth.

Variables are used to maintain and pass information within the rule, eg gender.

Enumerations allow the rule developer to define a limited set of values for some object. In some cases, there may be some unit of business logic that involves executing a series of steps in a rigid, pre-defined order. While this could be implemented as a rule, the function entity is designed to capture this form of business logic.

BUSINESS APPLICATION

The following section explores some of the business applications of rules-based engines in a customer relationship management (CRM) context.

Perhaps the most common use of business rules-based engines is the processing of applications for products and services, where there are a large number of complex business rules that have to be applied and where these rules change frequently. For example, processing complex lending rules as part of personal loan application.

For delivery of context-specific information the business rules are used to determine when context-specific

information is to be provided to a client or employee as part of a standard process. For example, providing context-specific data in a telemarketing script.

In order to provide personalised content on websites the business rules are used to determine the specific content to be presented to a customer as part of an interaction on the website. This personalisation could be based on the customer's segment, attributes profile or behaviour on the site.

For supported selling the business rules are used in combination with product purchase propensity score to determine the 'next best offer' as part of a supported selling process at the point of sale.

In the case of lead allocation in channels the business rules are used to determine how a sales lead is to be allocated to a specific channel or sales agent within the channel. This approach facilitates a consistent treatment of a sales opportunity while ensuring effective use of channel resource. For example, allocation of mortgage sales leads between a retail branch, direct sales force and third-party agent where the individual is not account managed.

Rules-based engines work best when the:

- process is well defined, complex, uses a large number of rules or rule sets and is executed in real time
- business policies or procedures change frequently
- business users want to control changes in the business rules.

BENEFITS OF RULES ENGINES

Modern business rules-based engines have the following benefits, they:

- allow non-IT users to create, maintain and monitor the impact of business rules on processes

- allow organisations to automate and streamline processes
- allow organisations to ensure consistency of process
- allow organisations to increase the efficiency of the processes
- provide a clear audit trail of the decisioning processes.

ISSUES WITH RULES ENGINES

The following section explores some of the problems associated with rules engines.

Performance

Where the business process requires the application of a large number of rules or rule sets, performance of the decisioning process can be significantly affected. This is particularly true where the rules engine is being used in real time and has to access data from an associated data repository (eg operational data store) rather than from the customer interaction (eg Web page interaction).

Batch processing

Most of the recent rules engine technologies have been developed with real-time application of the rules in mind: one of the consequences is that they have been poorly designed for large volume batch processing.

Perhaps one of the obvious potential applications of rules engines was as a core component of campaign management technology, but their success to date has been limited.

Understanding implications of change

Most business rules engines allow the user to change individual rules with ease but it is difficult to assess the

impact of these changes. A simple change in one rule or rule flow can radically change the outcome. This can cause significant problems where the rule engines are being used to drive real-time core business applications, eg instant credit processing.

In some cases the ability to test the impact of changes in the rules, rule sets or rule flows has been achieved by allowing the user to change the rules and to process the same series of events or data. The user can then identify the 'ideal' set of rules and rule parameters.

Effective change control and/or authorisation procedures are essential if the business is not to be adversely affected by any changes.

Integration with other business systems

For rules engines to be effectively deployed they often need to be integrated with other operational systems, eg call centre or branch teller systems. In some cases this can be quite difficult, although this issue has been recognised and the newer technologies have addressed the integration issue well. A range of different deployment models is becoming available.

Many rule engines can be run as a *service* (ie using threads provided by the container application in which it resides) or as a *server* (ie where it needs to use its own threads since there is no container application). They can be deployed in most environments that support Java including: EJB; J2EE; LDAP, MTS; RMI; CORBA; JSP, ASP and MQS series.

A few vendors supply standard deployment APIs that handle the interaction between the rule engine server and the calling application, eg call centre application.

THE USE OF RULES ENGINES FOR OFFER OPTIMISATION

A number of vendors have positioned rules engines as optimisation technologies for managing customer communications, following in the wake of statistical optimisation products like those from niche player Market Switch and mainstream players SAS and SPSS.

With the wider availability of channels that can be used to communicate with and receive communication from customers, has come a significant increase in the complexity of the communication–decision process. The key thing that the optimisation technologies are trying to do is to maximise mathematically the profit generated from marketing and other communications, removing the guesswork from the campaign management process. The problem is that with dozens of potential offers, many communication channels and thousands or millions of potential customers, there could be billions of different offer–customer–channel combinations to choose from.

Optimisation technologies use mathematical algorithms to determine the best combinations of customers, products and communication channels to maximise profit, while at the same time taking into account real-world business constraints such as limited budgets or channel capacity. The output of the optimisation process can then be used to drive communication with a customer. Analytical and modelling tools can be used to understand the customer and build behavioural models. For a given marketing offer, these models can be used to determine a customer's likelihood of responding to that offer.

The use of product or offer propensity models is becoming commonplace. But these single dimensional models do not take into account the impact of concurrent

offers, channel or business constraints. This is where statistical optimisation techniques can help to resolve the problem of maximising return on investment (ROI) across all planned communications. The better solutions take into account interdependencies of financial goals, business constraints and customers' needs to deliver the offer that is best for the business and satisfies the customers.

The optimisation technologies allow the user to define simply the business constraints, input customer response data and specify which offers and communication channels to consider. The software automatically selects the appropriate combination based on specified objectives, such as maximising profit, then rolls up the projected results and scores the database for offer distribution via the correct channel.

Business rules engines and other 'best guess' solutions that were invented to try to address this problem do not mathematically maximise the profit. These techniques do not consider the complicated interdependencies of multiple offers, are not easily updated with new information and cannot scale to the magnitude of the problem.

The following example compares the results of a rules-based approach to offer optimisation and a statistical optimisation approach. In this scenario:

Customer 1 has two potential offers:

Offer A through call centre with an expected value of \$100

Offer B through direct mail with an expected value of \$75.

Customer 2 has two potential offers:

Offer C through call centre with an expected value of \$90

Offer D through direct mail with an expected value of \$20.

Due to channel capacity constraints

only one call-centre contact can be attempted.

Using a typical rules-based approach the following would result:

Customer 1 would receive Offer A through call centre with an expected value of \$100.

Customer 2 would receive Offer D through direct mail with an expected value of \$20.

Hence the expected return from a rules-based approach would be \$120.

Using a statistical optimisation approach seeking to maximise the expected value, the following would result:

Customer 1 would receive Offer B through direct mail with an expected value of \$75.

Customer 2 would receive Offer C through call centre with an expected value of \$90.

Hence the expected return from the statistical optimisation approach would be \$165.

Within the rules-based approach, the process selects the first customer from a ranked list; this does not necessarily deliver the optimal solution. Within the statistical optimisation approach, the process selects the best combination of offers to maximise the optimisation function.

Hybrid solutions for communication optimisation

The emergence of hybrid solutions has been seen where rule-engine technology has been combined with other technologies such as statistical optimisation or neural networks. This approach seems to take the best of both approaches and is gaining in popularity.

CONCLUSION

The use of rules-based engines is gaining ground in a number of areas as they allow business users to manage the evolution of the rules and are simple to understand. Their application in communication decision making is being proposed by a number of vendors and in some cases these technologies are being positioned as optimisation solutions. They do have a role to play in a number of areas but they come a poor second when compared to statistical optimisation for communication decision making.

Acknowledgments

MarketSwitch Corporation, 108 Powers Court, Dulles, VA 20166, USA, telephone: +1 (703) 444-6750, website: marketswitch.com; SAS Institute Inc., SAS Campus Drive, Cary, NC 27513-2414, USA, telephone: +1 (919) 677 8000, website: sas.com.

© Shaun Doyle (2002)