



Case Based Reasoning

A reliable technique to automate manually controlled situations

Even as proliferation of automation increases, ever so often there are pockets in processes that are manually controlled. These pockets are generally exceptions to the standard process and hence require manual involvement. We illustrate in this paper how repetitive instances of manual decision making can be automated using the Case Based Reasoning technique, bringing a significant level of automation even in the otherwise manual activity. This solution has helped an organization go one step further towards maturity by strengthening the processes via increased automation.

Railroads run shipments on cars that are sometimes leased from private parties. The rates governing the lease payable for car movements are influenced by multiple systems and most payables are computed programmatically but there are exceptions due to various reasons. The programs record a default prescribed rate on the exception records.

Manual review is time-consuming and capping the number of people working on the reviews results in backlogs and sometimes calls for a less-than-hundred percent review coverage (certain records to be passed without review). With a case-based reasoning solution, we helped the organization significantly improve operational efficiencies by reducing backlogs, increasing coverage on reviews and automation even in exception situations. The solution also established a foundation to improve on consistency.

About the Author

Jatin Doshi

Jatin Doshi has worked in the IT sector for 10 years and largely associated with the Case Based Reasoning section of the Artificial Intelligence Group. Jatin has been significantly involved in the development of TCS' product Consult, which is a case based reasoning engine and is currently managing delivery of solutions around Consult.

Shobha Santosh

Shobha Santosh heads the Call Management Solutions Group based out of TCS, Pune. She has been instrumental in the formation of CMS group and is responsible for taking the solution offerings (CONSULT and CCM.Net) in this group to the large customer base of TCS. She has an overall experience of 13 years in IT and has handled a wide range of responsibilities during this time.

Table of Contents

1. Perspective of Manually-Controlled Processes in Railroads	3
2. Case Base Reasoning Solution-Reduction of Manual Intervention	4
3. Case Tagging	4
4. Inconsistency Check	4
5. Approval Process	5
6. Integrity Maintenance	5
7. Gearing up for ROI on Day 1	5
8. Retrospection	5
9. Results	5
10. Conclusion	6

Perspective of Manually-Controlled Processes in Railroads

Railroads run shipments on cars that are sometimes leased from private parties. The lease payments are made based on the distance covered by the car and generally these are prescribed per mile rates for different car types by appropriate regulatory mechanisms. Railroads also have contracts with their customers (or shippers) that determine the rate to be applied for each shipment based on parameters such as the origin, destination, commodity and also certain car-related parameters, for instance, the car type and owner. Often there are instances where the individual shipper contract could override the prescribed rates.

Railroads have to compute the payable amount to private owners on a regular basis. Car movement records are maintained and tracked to determine the payables. In our interaction with a railroad, we observed that existing software programs compute payables by looking at the prescribed rate and the shipper contract before determining the payable. A percentage of the car movement records are tagged by the payable computing programs due to certain reasons that include the following:

- a. Some mismatch between the car movement data vis-à-vis the contract master data.
- b. The details of the shipper contract are not available and hence the prescribed rate has been picked up.

These tagged records were manually reviewed. It was observed that the manual verification was done largely based on experience. Further, the car movement records were repetitive in nature given that most generally the railroad shipped the same material between the same origin and destination pairs. Thus, the action taken on a payable record reviewed once could be automatically applied to similar payable records on recurrence.

The significant aspect here was that the business totally relied on the ability of the manual review to determine the accuracy. The only form of validating the manual reviews was by adding another reviewer. However, since the nature of transactions being reviewed were highly repetitive, this would not have been a cost-effective option. Our observation was that there was a need for some cost-effective mechanism to validate the accuracy.

Car movement is a continuous activity and hence, every day there could be payables created. This also means exceptions could be recorded on each day, theoretically. In order to provide adequate time for the manual review, the railroad programs that computed the payables for private cars ran once in the middle of the month. This program processed several thousands of records and tagged the exceptions – normally to the tune of a few thousand records. The exceptions were handled manually for the rest of the month and whatever backlog remained on the last Saturday, was deemed accurate without review.

The objectives of our solution were:

- a. Increase the overall automation and reduce manual review
- b. Reduce the number of records deemed payable without review when a review was actually required
- c. Increase the overall throughput and thereby permit the date for the programs to run later than the middle of the month so that more records were tracked. At the same time, ensure that no backlog remained at the end of the month
- d. Bring in more consistency in the process. Track instances where the same record was reviewed differently over time.
- e. Ensure automatic integrity of data in the solution repository and the master data.

Case Base Reasoning Solution – Reduction of Manual Intervention

The solution is proposed using Case Based Reasoning, an Artificial Intelligence technique that is based on the principle that similar problems have similar solutions. Problem-solution pairs are recorded in the form of 'cases'. The cases are recorded in a repository called the case base. The payable records at the railroad concerned were repetitive and we found this technique could help reduce the manual workload drastically. To implement the solution, we used the TCS' proprietary CBR engine, Consult.

With Consult, each instance of a reviewed payable is recorded as a case. Cases are stored in a case base, which is the solution repository. Each case is comprised of a pre-image (the problem) and a post-image (the solution or action). Each time a tagged exception record was encountered, its details were 'matched' against those on the pre-image of cases. If a similar case was found, the corresponding solution was used to automatically review the instance at hand.

If a case was not located similar to the current instance, a new case was created with the current instance. This case would, in due course of time, be available for resolving other like-instances. Thus the solution provided the ability of encapsulating newer situations into the repository on the fly, thereby increasing its capability.

A significant item of focus was on identifying the parameters that would form the case in the proposed solution. This is critical since the parameters identified would be used to emulate a human decision-making process and hence had to be 'complete'. A series of discussions to determine the utility of each parameter finally resulted in an optimal list of parameters that would form the case in the case base.

Case Tagging

Each case in Consult was tagged to indicate what it was authorized to do. Since the proposed technology and solution were not run-of-the-mill, it was important to have a mechanism where the user controlled the behavior of the solution. Thus, each case was tagged to indicate what it was authorized to do. When a case was created, a 'review' tag was associated with it to indicate that it was a pending review. A case in pending review was not permitted to automatically take action on an input record.

When a supervisor reviewed the case and found it to be accurate, the tag was changed to authorized. Authorized cases were permitted to automatically review records.

If a case was found to be inconsistent or inaccurate, then a 'revoked' tag was associated with it. This could happen when there were two cases that looked alike but had different review actions against them. The supervisor after reviewing could approve one by tagging it as authorized and tag the other one as revoked.

Inconsistency Check

The existing process did not validate what was done manually. In theory, it was possible that the lease for the same kind of car movement was paid for on some occasions and not paid for on other occasions. The proposed solution implemented checks to indicate the presence of any such occurrence. Any instances that were alike but were reviewed differently at different times were hence tracked and brought to the notice of the supervisor. The supervisor would review and determine which action was accurate and 'turn-off' the other cases. Thus over a period of time, consistency was expected to improve.

Approval Process

As the proposed solution was operating simply on the basis of the manual action taken during reviews, the solution was as fallible or as strong as the manual review itself. If the reviewer made a mistake in review, it would still be captured in the case base and reused consistently. To avoid a runaway errant, we implemented an approval mechanism. There were two kinds of approvals put in – manual and automatic. We identified a role of a supervisor, whose job would be to review the manual reviews that were captured in the case base and if found accurate, approve the case for automatic review. This worked well since human confidence grows with the number of capable people reviewing the same artifact.

We were expecting high volumes and hence to avoid overloading the supervisors with cases for review, we put in a provision for a frequency count based automatic approval. When a case was created, the case was not approved for automatic review and was in a state we called 'review'. Any new car movement record came by would still be matched over these cases as well. If there were a match, the 'match count' would be incremented. If the match count exceeded a certain threshold, the case would be automatically tagged to 'authorized'.

Integrity Maintenance

The contract database was prone to changes from time to time and any change in the terms of the contract relating to payment of car lease would need to be carried over to the case-based reasoning solution proposed. A simple utility was provided to maintain data integrity and ensure that whenever a case was updated due to a contract change, it was moved back into a 'pending review' status.

Gearing up for ROI on Day 1

As the railroad already had past data available, we proposed to 'seed' the case base repository with the data for a few months before going live so that there was a possibility of getting a return on investment from day 1. Data was thus uploaded into the repository, reviewed, and certain cases authorized by the administrator before production move.

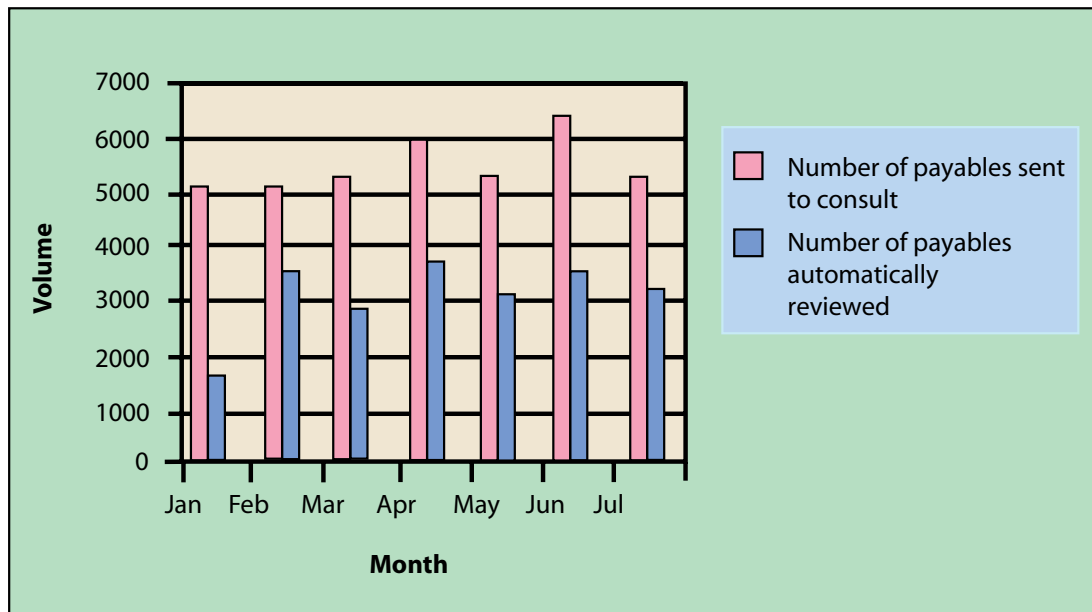
Retrospection

When we initially proposed a case-based reasoning solution, one of the questions raised was how it would be different from a rule-based system. The standard differentiator was that rule-based systems worked in well understood domains while case-based solutions worked well even in not-so-well understood processes. Rule-based systems are brittle and need data to be populated upfront. In comparison, case-based systems can be configured to be rigid or flexible. Data in case-based solutions can be captured on the fly, scoring a big point over rule-based systems.

Results

The solution was put into production with some historical data and during the first few months, we saw the automation range between 60-70%. This endorsed our observation that transactions were highly repetitive. At the same time, new cases were getting added at a rate of approximately 30-35% per month. This was indicative of new transactions consistently occurring.

Our observation of inconsistency being a possibility in theory was also vindicated. After a few months, we noted that 1% of the cases represented inconsistency in the manual review.



Conclusion

Many techniques have been used in automating traditionally and there has been a great deal of automation achieved. However, it is the seemingly small but nagging percentage of transactions that still call for manual involvement that can be now focused upon for automation with emerging techniques such as case-based reasoning. This technique does not claim to be a complete solution for 100% automation but does provide the promise to get you somewhere in the vicinity of that dream number.



About TCS' Call Management Solutions Group

TCS has an established Call Management Solutions Practice (CMS) to offer best-in-class call management solutions to its clients. The practice offers 2 solutions in the call management space namely CCM.Net – which is a helpdesk solution and CONSULT – which is a Knowledge Management solution for a helpdesk/ support domain. Providing cost-effective IT support in today's rapidly changing computing environments is a challenging, frequently frustrating, and sometimes seemingly impossible task. Call Management Solutions group from TCS has robust solution offerings in CONSULT to address these challenges. The focus of this group is to help organizations address these very business challenges by speedy call problem resolution. We offer Innovative solutions, Superior Service, expertise of core Consulting team and expertise in IT domain.

About Tata Consultancy Services (TCS)

Tata Consultancy Services (TCS) is among the leading global information technology consulting, services and business process outsourcing organizations. Pioneer of the flexible global delivery model for IT services that enables organizations to operate more efficiently and produce more value, TCS focuses on delivering technology led business solutions to its international customers across varied industries.

For more information contact

CMS Information Center
Tata Consultancy Services Limited
Godrej Millenium Bldg.
Pune 411001, India

Phone: +91-20-6608-7816
+91-20-6608-7873

Email: cms.enquiry@tcs.com
Website: www.tcs.com

TATA CONSULTANCY SERVICES

All content / information present here is the exclusive property of Tata Consultancy Services Limited (TCS). The content / information contained here is correct at the time of publishing. No material from here may be copied, modified, reproduced, republished, uploaded, transmitted, posted or distributed in any form without prior written permission from TCS. Unauthorized use of the content / information appearing here may violate copyright, trademark and other applicable laws, and could result in criminal or civil penalties.