

# Cyborgorganisation: Machines and Humans Make Optimal Decisions Together



Alan Dormer

**Abstract** Business rules evolved from expert systems, a concept whereby expertise could be encoded and leveraged across an organisation. In the same way, artificial intelligence is now being used to replace human decision-makers. In both cases, the idea is to emulate and replace the decision-maker with a machine. But this approach could be considered as misguided for two reasons. Firstly, it is not possible to emulate a human with complete accuracy, and secondly, human decision-makers are fallible. Another approach is to consider how business rules and human experts can work together to maximise the expected profit of an organisation, creating a cyborgorganisation. There are several elements to this problem—the need to quantify the impact of different rules on the performance of the organisation, the accuracy of the human decision-maker on a case-by-case basis, and determine whether the machine and/or the human decision-maker makes the final decision. This paper considers a real example from a well-understood problem that of loan approval, but from the perspective of machine augmenting, rather than replacing, the human decision-maker. The results suggest that there are potential savings and increases in profit from this approach.

## 1 Introduction

A business rule is a rule that defines or constrains some aspect of a business. Business rules are intended to assert business structure or to control or influence the behaviour of the business [1]. Business rules are very common within industry, particularly the service sector. They enable consistent decision-making by individuals and a degree of automation using business rules engines. Business intelligence, expertise, and experience can be codified and used by many less expert and less experienced staff. Generally, rules add value to an organisation by reducing cost (either replac-

---

A. Dormer (✉)

Department of Information Technology, Monash University, Clayton, Australia  
e-mail: Alan.dormer@monash.edu

© Springer Nature Singapore Pte Ltd. 2019  
X.-S. Yang et al. (eds.), *Third International Congress on Information and Communication Technology*, Advances in Intelligent Systems and Computing 797, [https://doi.org/10.1007/978-981-13-1165-9\\_45](https://doi.org/10.1007/978-981-13-1165-9_45)

487

ing human decision-makers completely or by less experienced staff) and providing greater transparency and consistency.

As an example, in a mortgage application, the aim is to categorise customers based on data provided to maximise profit and minimise risk; the organisation wants to accept customers that will repay and reject customers that will default. The structure and parameters of the rules will have an impact on the accuracy of classification and eventual profitability. For example, if very strict criteria are applied, this may reduce the number of bad loans but, at the same time, refuse loans to some good customers. Conversely, relaxing the criteria will increase the acceptance rate, and potential profits, at a cost of higher losses.

Rules typically determine the flow of work and information in both services and manufacturing. In services, this can include the acceptance of new customers, offering different customers different service levels, accepting claims, etc. In manufacturing, there is acceptance, rework or rejection of components, rules on the acceptance of urgent orders, warranty claims, etc. In this paper, we are seeking to apply optimisation to a business process where

- Decisions are made by (imperfect) people
- Automating (some) decisions have the potential deliver cost savings and greater consistency
- Complete automation is too expensive, impossible, or undesirable

This applies to many sectors, such as

- Health: diagnosis and treatment decisions
- Social services: benefit claims and child protection
- Financial services: insurance and credit

We consider business rules optimisation (BRO) as determining that set of business rules that maximises the expected net contribution (e.g., revenue or profit) to the organisation that uses them. Hence, we are concerned with the structure and the parameters within the business rules insofar as they impact the business performance.

The following section reviews relevant literature in business rules and business process optimisation. In Sect. 3, we provide an informal and a formal definition of the problem addressed in this paper. The methodology and its application are discussed in Sect. 4, with results in Sect. 5, and conclusions in Sect. 6.

## 2 Literature Review

Business rules do have an impact on an organisation, and that they can be optimised to achieve certain outcomes or key performance indicators (KPIs), such as the percentage of customers accepted [2, 3]. Business rules also have an impact on organisational performance in retail [4] and inventory management [5]. Rules can also be optimised to allocate work to a range of human decision-makers based on

their skill levels and availabilities [6]. It is possible to go further and seek to optimise business rules to maximise the expected profit of an organisation [7]. Rules can also be combined with optimisation techniques on a case-by-case basis [8]. Here, we are looking at the ability to optimise the rules in advance so that the overall expected profit is maximised, with the potential for rules to operate alongside human decision-makers.

There is a large body of work in machine learning where the costs can be incorporated into a decision tree (which is essentially a set of rules). For example, [9] considers the cost of obtaining further information and [10] considers the cost of classification errors. Generally, this research is an extension of a methodology rather than solving a specific business problem.

Business process optimisation (BPO) is an area of research that considers a business problem that focusses predominately on the processes required to produce an outcome at minimum time and/or cost, and the way that tasks and activities are created, ordered, and interlinked, [11]. Rules, where cases or components are directed one way or the other, are not considered, and the idea of expected value is not considered [12].

### 3 Theory

The general formulation of BRO has already been proposed [7] where the initial assessment process was analysed. In this paper, we are analysing the second stage of the process and seeking to model the performance of the rules and the human expert rather than assume that either will be 100% accurate. The process that we consider is given below where there is a customer enquiry (a case) with associated data (the attributes) (Fig. 1).

Outcomes and expert decisions can be formulated using the LENS model [13]. This proposes a model of the outcome and expert decision as a function of the attributes or characteristic values of any case. For every case, we have four quantities:

- Predicted outcome
- Predicted decision
- Actual outcome
- Actual decision

This is a much richer set of relationships than a model of the expert decision alone, which is often the case in attempts to emulate the expert, and in emulating the expert, any fallibilities are taken into account (Fig. 2).

The essence of the LENS model is to create a model of the expert and the outcomes as a function of the attributes. We can use this concept by considering that we have:

- Rules, executed by a machine, that predict the outcome
- Experts, that can be modelled, also predict the outcome

**Table 1** Confusion matrix

	Judged good	Judged bad
Good outcomes	A	B
Bad outcomes	C	D

But to decide which is best for any particular case, we need to estimate the expected outcome if we use the machine compared to that of using the expert. We can use Baye’s rule to calculate the probability that the expert’s decision is correct. We have:

- $p(\text{good})$  and  $p(\text{bad})$  are the probabilities that a case is good or bad
- $Eg$  and  $Eb$  denote when the expert classifies the case as good or bad
- $p(Eg)$  and  $p(Eb)$  are the probabilities that the expert classifies the case as good or bad

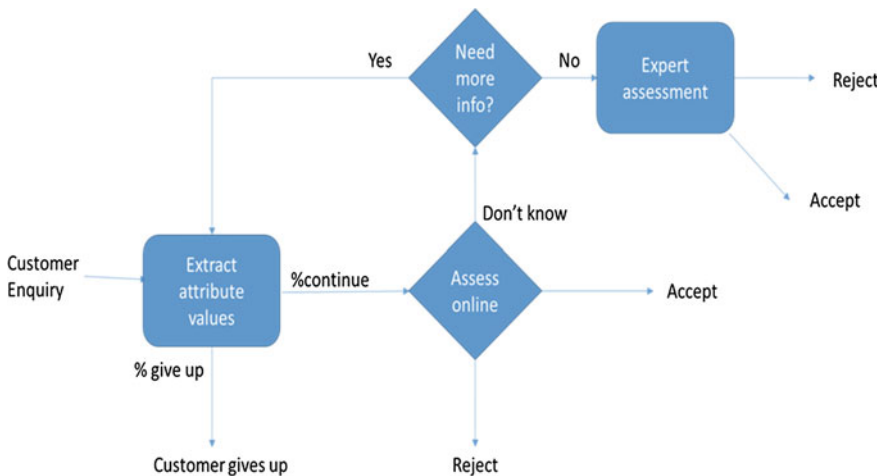
$$p(\text{good}|Eg) = \frac{p(Eg|\text{good})p(\text{good})}{p(Eg)}$$

We can use the confusion matrix to illustrate the point (Table 1). We have

$$p(\text{good}|Eg) = \frac{a}{a + c}$$


and

$$\frac{p(Eg|\text{good})p(\text{good})}{p(Eg)} = \frac{a}{a + b} \frac{a + b}{a + b + c + d} \frac{a + b + c + d}{a + c} = \frac{a}{a + c}$$



**Fig. 1** Generic business process diagram

Fig. 2 LENS model



We can use a similar approach to any estimate or calculation of the nature of a case, for example by the rules. Let us define:

- $p(Rg)$  as the probability that the rules determine a case to be good and
- $p(Rg|good)$  as the conditional probability

Then for any case, we can determine the expected net gain of giving a case to the rules or the expert. If  $I$  denote the potential gain, from a good case, as  $G$  and the potential loss, from a bad case as  $L$ , and the processing cost of the expert assessment as  $E$ , we have, for the cases processed by either the rules or the expert:

$$\text{Net gain from rules, } G_R = p(Rg)(Gp(\text{good}|Rg) + L(1 - p(\text{good}|Rg)))$$

$$\text{Net gain from expert, } G_E = p(Eg)(Gp(\text{good}|Eg) + L(1 - p(\text{good}|Eg))) - E$$

## 4 Application

In order to apply this method, we need to be able to calculate or estimate the probabilities and conditional probabilities identified above. For each case, we need to the following training data:

- Attributes
- Expert decision
- Outcome

And for any given set of rules, we observe that at the end of any rule sequence (e.g., a leaf on a decision tree or the end of an IF-THEN-ELSE process) there will be a confusion matrix and that there will be data in the training set. However, there is a potential complication insofar as we will have data on  $a$  and  $c$  (the outcomes of cases that are accepted) but will only know  $b + d$  (as we do not know the outcome for cases we reject). While this may impact the calculation of  $p(\text{good})$ , it does not make any difference to  $p(\text{good}|Eg)$  as above, but it will impact on the calculation

**Table 2** Rule decisions

Rules	Determined good	Determined bad
Good outcomes	40	40
Outcome bad/not known	20	

**Table 3** Expert decisions

Expert	Judged good	Judged bad
Good outcomes	45	50
Bad outcomes	5	

of  $p(\text{good}|Rg)$ . For my purpose, we can assume that all the cases that the expert judges to be bad will be bad, meaning that  $b = 0$  which is a worst case (that is lowest) estimate for  $p(\text{good}|Rg)$ . This will give the most conservative business process where the expert is called upon more frequently than strictly required.

So, given that we know the decisions of the expert within the training set, we can identify those cases that have been judged good and bad (by the expert) and apply that to the cases in each of the leaves at the end of the rule sequence.

For example, suppose we build a set of rules and one of the leaves consists of (Table 2).

And when the expert decides, we have (Table 3).

In this case, we get

$$p(\text{good}|Rg) = \frac{40}{60} = 67\% \quad \text{and} \quad p(\text{good}|Eg) = \frac{45}{50} = 90\%$$

One of the drawbacks of this approach could be the limited sample size in the leaves, with perhaps few or no examples of each classification. If this is genuine, it is not a problem as it suggests a high degree of purity, but we could also just be the training data. One way to avoid this is to use the LENS model more closely and use regression to build a model of the expert’s decision-making. Linear regression is typically used but our case logistic regression is more helpful as we can get an estimate of  $p(Eg|\text{good})$  and  $p(Eg)$  by fitting to the training set and thus calculate  $p(\text{good}|Eg)$ . Specifically, we look at the cases that the expert judged good and the cases that the cases where the expert judged a good case correctly. Similar considerations apply to the rules, where we can derive the relevant probabilities from inspection or regression, with my conservative assumption that all cases judged to be bad are, in fact, bad.

We have applied this approach to data from the Lending Club [14]. This is a very useful data set for our purpose as it includes data on judgements and outcomes. The data is typical for loan applications with attributes including credit score, accommodation, income, purpose of loan, etc. Outcome data includes loans that are paid up (good outcome) or in default (bad outcome).

Lending Club applies a two-stage approach. The first stage uses a reduced data set to decide whether further assessment is required. Once past the first stage, additional detailed information is requested, and a more detailed assessment made.

## 4.1 *The Process*

The reduced data set has already been analysed using a decision tree to determine how much of the acceptance process could be automated [8]. The results indicate that about 80% of assessments could be automated with a net benefit; the overall expected profit would go up slightly. In that case, we used a decision tree as did not need to estimate the accuracy of the expert; the assumption was that the expert was correct in every case. This was not an unreasonable assumption for the purposes of that exercise, but in this paper, we use data on decisions and outcomes.

Analysis of the detailed, loan status, data shows that the expert is not entirely accurate. Some 15% of loans become delinquent and were in default. This is a considerable cost, and we cannot therefore assume that passing the decision to the expert will result in an accurate assessment.

We then separated two-third of the data to fit regression models of the expert, rules, and outcomes, leaving one-third of the data to test the algorithm on data that had not been used to build it. We used the Weka package for this [15].

Here, unlike most previous research, classification accuracy is not the objective. we are looking for a way to determine which cases can be automatically classified and which ones could be given to the expert.

In order to test the efficacy of this approach, we have applied the following process to the validation set to determine what the decision would be and then calculate the costs and benefits. This first analysis was to determine if there was any merit in this approach, not necessarily to fine tune the process. Nor are we checking the accuracy of the regression; this has already been determined by cross-validation. What we are doing is running through a process on past data as an indication of how well it could work in practice.

Clearly, there is further potential to fine tune and validate with other data, or fit the regression with some of the data and test with the rest. However, at this stage we are more interested in whether it works at all.

The procedure is outlined below. Essentially, we determine the probability that either the machine or the expert will be correct when it makes a 'good' decision and then the expected net benefit thereof. The party with the higher number is given the decision.

For the purposes of the analysis, we assumed that potential profit and potential loss are equal to the size of the loan. The processing costs are a function of the amount of information, for example, a familiarisation time and a time proportional to the amount of information. we estimated this cost to be 5% of the loan value (Fig. 3).

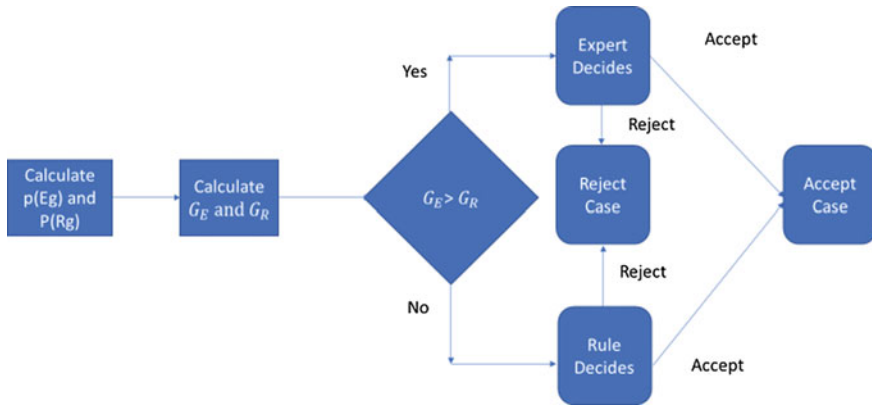


Fig. 3 Proposed rule execution flow chart

## 5 Results and Discussion

### 5.1 Results

The optimisation process yielded some interesting results. About 30% of the cases were deemed better given to the expert, the remaining 70% left to the machine. The overall benefit, as a percentage of the loan book, was between 4%. This is a result like the previous result of 80% of the cases to the machine and a 2% benefit, but that assumed that the expert was perfect. The benefit of this approach was impacted by the cases with unknown outcomes. We had to assume that, if the machine accepted a case that the expert had rejected, then the outcome would have been bad. That is not necessarily the case, and it does rather reduce the benefit. Even so, this approach (under the assumptions made) has some benefit. Overall there were potential benefits equivalent to about 4% of the loan book, which given that Lending Club is now at about \$24Bn and is worth at least further analysis.

Another point to note is the significance of labour cost saving. Relying on expert labour is a risk for any business, there is always variability that needs to be managed, and the business is not scalable. If this can be automated a large proportion of transactions and understand the trade-off between labour costs and net loan value, variations in applications (up and down) can be better accommodated.

Note that the assumption that the case with an unknown outcome is pessimistic, particularly that this data set does not include applications that are rejected without further analysis. If we had assumed that only half of these cases are bad, with the other half being good, the benefit increases to 7%. We could legitimately say that the potential benefit is between 4% (no unknown cases are good) and 10% (all unknown cases are good). Note that there are no unknown cases when they are passed to the expert; We know already that his or her decision is to reject. The table below



**Table 4** Results of optimisation

	Expert only	Machine	Expert
Cases (%)	100	68	32
Good	\$123,849	\$95,291	\$27,872
Bad	\$22,207	\$15,829	\$5,997
Unknown		\$1,480	
Expenses	\$7,302	\$0	\$1,958
Net	\$94,340	\$77,982	\$19,918
Total	\$94,340	\$97,900	
Saving		\$3,561	3.78%

shows the results of applying this approach to the data set, with all amounts in 1000s (Table 4).

This demonstrates that the savings come from a reduction in labour cost (27%) which far outweighs a 1.8% reduction in net loan value.

## 5.2 Discussion

The results indicate that there is potential benefit in this approach, consisting of savings in labour costs and higher profits. In this case, there are benefits of between 4 and 10% of total loan value; for Lending Club, this could run into \$Bns. There are many options in terms of further development and extension.

We have demonstrated and quantified a practical method for optimising the rules and determining the number of cases that are processed automatically and by the human expert, based on a historical data set with some simplifications and a fixed set of attributes. To apply this approach, we only need data that is normally available on attributes, judgements, and outcomes. But we have to accept that any data set is very unlikely to have outcomes for rejected cases. In this example, we have taken a conservative approach and assumed that all rejected cases are bad, which will mean that more cases than strictly necessary will be referred to the human expert. There may be ways to refine this approach, for example, by using the probabilities obtained by logistical regression on the good and bad cases where we know the outcome.

If this approach was to be adopted, it would generate a different set of data because not every case would be examined by an expert, yet there would still be the problem of not knowing the outcomes of rejected cases. This further suggests that having models (based on the training data) that deal with this missing data would be very helpful.

We could use the approach in this paper to address other problems that have the same features. Consider medical treatment, for example. There is data on symptoms of the patients (the attributes), the tests (additional attributes), the doctors and medical procedures (the experts), the costs of these, and the eventual outcomes. The deci-

sions about which tests and procedures can be made based on the incremental value (probability of recovery) and the cost in the same way that we decide on referring the case to the expert, and at each stage in the process, we can estimate this for any next stage.

## 6 Conclusion

Business rules are widely used with industry yet there is little research around how to create rules that in some sense optimise the expected value of their application or consider how rules can be complemented by human experts. Machine learning research has addressed some aspects of this problem, and business rules research has recognised that business rules can be used to shape business outcomes, but with neither addressing the complete problem nor keeping human experts as part of the solution.

On the other hand, BPO addresses a well-defined business problem but does not consider the rules.

We have defined the concept of BRO in the context of their use in the services sector and created the representative example of customer selection. The procedure we have developed employs ideas and techniques from machine learning and psychology and is complementary to BPO and could be used on a wide variety of business problems where decisions have costs and consequences, and neither the rules of the human experts are perfect.

Using real data on decisions and outcomes, we have demonstrated the potential incremental benefits of using rules and human experts together. Given the results, the potential of this approach is considerable.

We have used the term cyborgorganisation to reinforce the idea that an organisation, consisting of humans and machines (neither of which are perfect) working together by design, can achieve a better outcome when the relative performance of each party is factored into the overall business process.

**Funding** This research did not receive any specific grant from funding agencies in the public, commercial or no-for-profit sectors.

## References

1. Business Rules Group (2000, July) Final report, Revision 1.3
2. Wang O, Liberti L (2017, June) Controlling some statistical properties of business rules programs. In: International conference on learning and intelligent optimization. Springer, Cham, pp 263–276
3. Wang O, Liberti L, D'Ambrosio C, de Sainte Marie C, Ke C (2016, July) Controlling the average behavior of business rules programs. In: International symposium on rules and rule markup languages for the semantic web. Springer International Publishing, pp 83–96

4. Kunz TP, Crone SF (2015) The impact of practitioner business rules on the optimality of a static retail revenue management system. *J Revenue Pricing Manag* 14(3):198–210
5. Quinzanos JM, Cartas A, Vidales P, Maldonado A (2014, May) iDispatcher: using business rules to allocate and balance workloads. In: *DSS*, pp 110–119
6. Hegazi MO (2015) Measuring and predicting the impacts of business rules using fuzzy logic. *Int J Comput Sci Inf Secur* 13(12):59
7. Dormer A (2017, June) A framework for optimising business rules, presented at BIS. In: 20th international conference on business information systems, Poznań, Poland, pp 28–30
8. Taylor J (2011) *Decision management systems: a practical guide to using business rules and predictive analytics*, IBM Press, September 30, 2011, ISBN 0-13-288438-0
9. Turney PD (1995) Cost-sensitive classification: Empirical evaluation of a hybrid genetic decision tree induction algorithm. *J Artif Intell Res* 2:369–409
10. Zadrozny B, Langford J, Abe N (2003, November) Cost-sensitive learning by cost-proportionate example weighting. In: *Third IEEE international conference on data mining, ICDM*, IEEE, pp 435–442
11. Vergidis K, Tiwari A, Majeed B (2008) *Business process analysis and optimization: beyond reengineering (Applications and Reviews)*. *IEEE Trans Syst Man Cybern Part C* 38(1):69–82
12. Laguna M, Marklund J (2013) *Business process modeling, simulation and design*. CRC Press
13. Brunswik E (1985) *The essential Brunswik: beginnings, explications, applications*, New directions in research on decision making. In: *Research conference on subjective probability, utility and decision making*, Helsinki, Finland
14. Lending club, <https://www.lendingclub.com/info/download-data.action>, viewed 9/2/17
15. Frank E, Hall MA, Ian HW (2016) *The WEKA workbench. Online appendix for data mining: practical machine learning tools and techniques*, 4th edn. Morgan Kaufmann