DURLINGER CONSULTANCY

Steven Pauly
Paul Durlinger MSc

DURLINGER ESSENTIAL

Machine Learning

we make problems easy

DURLINGER
CONSULTANCY

DURLINGER
ESSENTIAL

Machine Learning

Steven Pauly
Paul Durlinger MSc

Contents

1. Introduction

There's no getting around the fact that the complexity of our logistic supply chain has increased exponentially in recent years. Where good business strategy used to be enough, it appears that maintaining control over the complex of cost factors will provide the main competitive advantage in the coming years.

In inventory management it may well become even more challenging. There is not only the complexity to overcome, but there are also the people who must be able to deal with it. This is a challenge. Many managers are not yet bothering about it, although arguably they should be. "Everything was better in the olden days". Well, in inventory management this might well be true. But just saying so doesn't provide a solution.

2. The power of machine learning

More than 50 years ago, knowledge about inventory management increased. The value of good inventory management turned out to be enormous and that led to many thick tomes being written on the subject. However, at the time there was a distinct lack of computing power and solutions were based on healthy common sense, albeit supported by the necessary mathematical analysis. This had the advantage that solutions were transparent, not too difficult, provided good insight and primarily focused on so-called '*quick wins*'. The *Economic Order Quantity* [EOQ], 'invented' in 1913 (see Harris) is a good example. If we know the parameters, we can calculate the batch size for a product on the back of a beer mat. Any manager reading this may well ask himself: then why do it any other way?

Well, there were many problems in those days, and many more arising today, that can't be solved through basic mathematical analysis and common sense. Where the back of a beer mat isn't enough.

3. A machine-learning example

An example of this type of problem is the calculation of 'out-of-stock' costs. Suppose a customer goes to the local store to buy a prize-winning bottle of Rioja Gran Reserva 2014 from the winery Marques de Caceres. But it is sold out. What will this customer do now? Will he buy a 2015 or 2016 Gran Reserva from the same winery? Or a Gran Reserva 2014 from a different winery? Or just some other red wine (but if so, which one?). Or leave the store without buying anything? Or no longer shop here for wine, just for other products? Or does he not come back at all? And what if a competing store is only 300 meters away, or 15 km away and the customer hasn't got a car? "The answer will depend on the situation". But what situation? Welcome to the world of machine learning.

A lot of effort was put into looking for more efficient solutions to overcome this complexity. Especially because a brute-force approach, in which we evaluate every possible outcome and then choose the best, didn't work well in practice. Given the growing interest in artificial intelligence, machine learning became the technique that received most attention.

"Solutions in inventory management haven't changed. The tools have."

4. Deep Blue Chess

Perhaps you have heard or read the words 'DBC system', 'machine learning', 'IBM' and the year '1997' in the same sentence. The Deep Blue Chess (DBC) system from IBM took centre stage in the story that took place in 1997. Through machine learning they managed to defeat Kasparov, the reigning world champion chess. Impressive, you might say.

This victory is not quite as clear-cut as it may seem, however. IBM staff claimed that the machine explored all possible combinations in each situation, taking into account the next possible moves and the patterns in Garry Kasparov play. This is more of a brute-force approach, of course, and would be particularly impressive as an exhibition of raw computer power. This claim raised Garry Kasparov's suspicions and led him to think IBM cheated. He demanded a rematch, but IBM refused and dismantled the machine soon after.

5. World champion GO defeated

In 2016, there was a true exhibition of the power of machine learning. Once again, the machine was pitted against the best human player in the world in a strategic game, this time the game of 'Go'. As before, they succeeded in their aim.

We should realise that 'Go' differs from chess in the number of possible combinations: in 'Go', there are 10^{174} possible board configurations. To give you some idea of what this means: there are 1 million trillion trillion trillion trillion more possible board configurations in Go than in chess. This makes a brute-force approach to the game virtually impossible.

6. What is Machine Learning?

So, what is machine learning? And what makes it different from brute-force and traditional mathematics? And why do we have to use it? And how? And what exactly does management have to do? These are all questions that managers must be able to answer in order to introduce machine learning into practical applications.

Machine learning differs from brute-force or traditional mathematics first and foremost by its learning component. This allows the machine to discover connections and patterns in a data structure without explicitly naming this. The machine learns the basic 'rules' of the problem. This allows it to tackle problems with high, underlying complexity and a high degree of uncertainty. Which is exactly what inventory management has to deal with more and more nowadays.

The fact that machine learning differs from other approaches creates new, valuable possibilities. This makes it possible to improve current forecasting techniques, for example, but also allows us to tackle many other issues that were not even considered a few years ago; for example, determining when exactly an item becomes obsolete. These are a type of multivariate regression problems with large numbers of variables.

But the quality of the solutions depends heavily on the relationships and the amount of reliable data available. The reaction of connoisseur Steven Pauly to the absence of the Rioja Gran Reserva 2014 will be very different from that of Paul Durlinger, who rarely drinks wine. This illustrates the potential power of machine learning, but it also shows where many managers may slip up.

7. Machine learning is a means, not an end

In essence, machine learning is nothing more than applied statistics with an emphasis on integrating current computer power. And given the increasing number of possible data sources, the further evolution of computing power and the growing complexity in our field, machine learning can be a tremendously powerful tool. But that is all it is: a means, 'just another tool in the box'. It should not become a goal for companies to get into machine learning. The goal in inventory management remains to translate uncertainty about all the present factors into an optimal decision process for the long term. Machine learning can play an important part in this, but it needs the right framework.

Every inventory optimization tool in a dynamic environment involves machine learning. The question that managers should ask software suppliers is not "do you incorporate machine learning?", but instead "*why* and *how* do you incorporate machine learning?"

The first question to ask is: "what is a good solution in inventory management?" Or rather, "what are the rules for optimisation?"

8. Four basic principles

Seven centuries ago the British philosopher William of Ockham introduced the problem-solving principle (1) known as '*Occam's Razor*' (Wikipedia). This principle states that solutions with fewer assumptions, and therefore simpler solutions, will work better than complex solutions.

A while later (well, a long time later actually, in 1906), another important optimisation principle (2) emerged. The great statistician Francis Galton (Wikipedia) observed a competition in 1906 in which the weight of an ox was to be estimated. He discovered that the average guess was extremely close to the actual weight. This turned out not to be a one-off observation. It is due to the fact that everyone considers a problem in a different way. Combining these various aspects eliminates the weaknesses in each approach.

Another principle (3) was that of the well-known management consultant Joseph M. Juran. He was the one who put the Pareto principle (the well-known 80-20 rule!) into a broader perspective. This principle is found to apply in optimization within inventory management: a small number of solutions

can result in a very large part of the savings for companies. These solutions are known as 'quick wins'.

The final principle (4) came up specifically in research on machine learning: the 'no-free-lunch' principle states that in the absence of prior knowledge, no optimisation algorithm can be considered 'best'.

9. Machine learning in practice

These four principles clearly describe the 'rules' within optimisation. In order to put machine learning to use as a powerful tool, one must therefore adhere to these rules. What follows are some important guidelines for management to do this.

First and foremost, management must determine what the desired outcome is before implementing machine learning. The idea is to look for those 'quick wins', and not just for an opportunity to use machine learning.

Management must also find out where and how machine learning can work significantly better than other, perhaps simpler solutions. "Why use machine learning?" is the question that needs to be validated at this stage. In other words, machine learning must be evaluated as a tool in comparison to other methods in terms of accuracy,

transparency and the effort required at the input and output sides. This should result in 'business rules for machine learning'.

Inventory management is driven by short-term decisions (ad-hoc decisions that define daily operations), medium-term decisions (predefined parameters that control the operation) and long-term decisions (a strategy that directs the tactical plan). Machine learning should therefore not only be used to determine parameters, but to provide insights across all levels.

These guidelines can be translated into a step-by-step approach that should lead to machine learning in its full strength as a means of inventory optimization.

10. Four questions about Machine Learning

This approach consists of four questions that need an answer:

1. Is machine learning the best possible tool for this problem in terms of results and efficiency?

2. Which tasks and in what order - keeping the correct objectives regarding inventory management in mind - will provide the most efficient way to arrive at the 'best' solution?

3. Which machine learning techniques and resources are required for the previous step?

4. Once a solution is found, how can we improve this in terms of simplicity, translating the acquired knowledge to other insights and improving future machine learning adventures?

"Machine learning can be a very powerful tool in modern inventory management, as long as we use it in an optimisation framework. This framework requires the right knowledge, the right process and thorough research."

11. Requirements: the DSP principle

Management must understand that machine learning can only be of value if used properly. And it comes at a price. There are some important requirements that must be met.

First the most important and influential variables must be identified. The demand for a certain type of ice-cream, for example, will depend on the temperature, the age (child, adult) and gender of the buyer, the place where it will be bought (beach, city, playground) and probably on many more factors.

Secondly we need a lot of data to establish relationships between these variables. Thirdly we need a lot of computer power, and fourthly extensive understanding of statistics. In short, *the DSP-principle* applies: Data, Statistics, People/ Power.

And even given these basic requirements it is not so easy. Google-maps can give you a quite accurate estimate of the time needed to drive from Amsterdam to Brussels if you will be leaving at 11 AM tomorrow morning. This is because Google has gathered the data of thousands of cars driving from Amsterdam to Brussels. But even then it is difficult to predict the actual time it will take, due to traffic accidents or other unforeseen events not (yet) included in the Google database. So the ETA is updated regularly along the way.

On the other hand, recommendations given by Amazon or Booking.com when you try to order or book can be rather poor. You can be sure there is enough computer power available and there are enough smart people in these companies. The problem is that there is simply not enough available data (yet). The recommendations I receive may improve if I buy more books or book more hotels. But to predict my reaction in the winery when that particular wine is out of stock, the computer will want to analyse thousands of my visits to that store. That would take several lifetimes, I am afraid. For this reason we expect only limited results from Machine Learning in the near future. Even so, this does not mean that we should not put in the effort!

12. *Call to action*

So what should management do now? The main thing is to look for opportunities within their own company. A good tip is to look for challenges that strongly influence the company's well-being, but where there is currently insufficient insight to be able to find a good response. Cases where the response depends on many variables and where sufficient data is available, examples of the outcome and the evolution towards it. Practical examples are optimising a total promotion policy, optimal sourcing with different channels in a general setting, providing forecasting and insights with respect to irregular and new items.

Management must know that Machine Learning can be extremely valuable if done properly. It is therefore highly advisable to set up a strategic company or a strong tactical level internally, with a great deal of knowledge and experience in this type of technology, the field of stock optimisation and the necessary preconditions to guide you in current and future innovations. Only in this way can an answer be given to current and future complexity that will strongly influence the competitive advantage of companies.

Let us conclude with these wise words from Charles Darwin:
"It is not the strongest or smartest of the species that survives, but the most adaptable."

13. Literature

Harris F. [1913]
'How many parts to make at once'
In: *Factory, The Magazine of Management*

Steven Pauly

As senior consultant and research scientist at Slimstock, Steven Pauly specialises in the mathematical side of stock optimisation.
He is also attached to the Slimstock Academy.

Paul Durlinger

Paul lectured at several universities and has twenty years of experience as a senior consultant. He sees it as his mission to make difficult concepts readily understandable.
He is also attached to the Slimstock Academy.

DURLINGER
CONSULTANCY

Nieuw Holsterweg 15
6061 EG Posterholt

☏ +31-6 224 07 919
Paul@durlinger.nl
www.durlinger.nl