



We analyse data because we need to answer questions. What did sales look like last month? Why was production low in a particular unit? How has the customer base responded to the latest promotion? The questions are endless, and so is the need for business analytics technologies and methods. Fortunately a wide range of tools and techniques are now available to us, and the data exhaust from transaction based activity fuels our analytical activities. Data from sales, production, procurement, web sites, social media, call centres - and so on, is available for analysis so we can address the uncertainties that are a part of everyday business life.

Analytical activities break down into a number of fairly well defined activities. These are the main headings:

Business Intelligence - often called descriptive analytics, is used to take a look in the rear view mirror. In its most familiar form it is scheduled reporting (daily, weekly, monthly), running queries against data (sales by region, by month), and more recently a whole raft of data visualisation tools mean business users can represent data in charts and dashboards as needed. Data warehouses have traditionally been used to format data in a way that makes it suitable for reporting and querying, but the advent of big data technologies, and particularly in-memory processing, means users can suck data from large data stores and process it directly in-memory. A new generation of tools characterised by Tableau, Qlik, and Spotfire support this type of processing, doing away with the need for IT to pre-process data into special formats for reporting purposes. The advantages of business intelligence technologies include ease-of-use, visual representation of data, and business user empowerment. The disadvantages relate mainly to the fact that we are always looking back, with no indication of how things might unfold in the future. It is also quite easy to misinterpret visual data - seeing trends that are no more than random accidents, and drawing conclusions from small samples of data. These latter issues particularly can be addressed through the next level of analytics - statistical analysis.

Statistics are used, by default, by many people in business. The notion of an 'average' is common to everyone, and other concepts such as variation, the median, a linear regression, quartiles, and so on are very commonly used. The main benefit of some level of statistical rigour is that users get to see how significant their data are. That three month rise in sales might have a forty per cent chance of simply happening by accident, and no business manager would want to act on a sixty per cent certainty. The problem with statistical methods is that they are unintuitive, complex, and can be misleading. The famous example of the man who walks across a river that is on average only one meter deep illustrates quite well - the 'average' doesn't tell him of the three meter pot hole. Nonetheless, statistical methods do add rigour, and allow business managers to make decisions with greater confidence. The use of statistical methods will change, and the growing use of Monte Carlo methods (where millions of iterations are performed to establish the probabilities associated with an outcome - for example the probability that a project will finish on time) is making a statistical approach available to a larger audience.

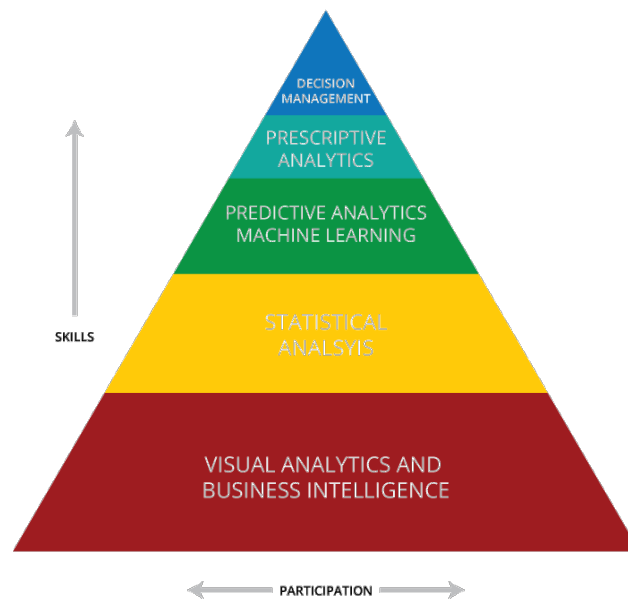


Predictive Analytics seems to promise a view into the future - otherwise it wouldn't be predictive. The essence of predictive analytics is very simple - analyse historical data for meaningful patterns of behaviour, test the patterns for reliability, and then apply them to new data. Many banks for example will trawl through their history of customer loans to find which customer profiles present a good risk. The resulting patterns will then be applied to new loan applicants - you don't get a loan if you already have two others, are unemployed and have more than two dependents. Data mining is one method used to detect such patterns, and this employs various methods for categorising data, for predicting numerical data (sales in August for example), for clustering data (which cluster represents people likely to buy a new product), and for finding associations (when someone buys eggs, will they also buy bacon). The advantages of well formed predictive models is obvious - we can target customers more precisely, predict when a supplier might fail to deliver, predict when a machine might fail - and a thousand other things. The downside relates mainly to the complexity of the methods and technologies, and to the generally unspoken fact that data mining will often throw out predictive models that look good on paper, but are pretty well useless in practice. There are ways to minimise this risk, but they are complex. And finally the terminology - machine learning refers mainly to the algorithms that are used in data mining - which is in turn used to build predictive models. It should be added that all analytical activity is in some way predictive. We wouldn't interrogate weekly reports unless we were trying to improve or fix something, which assumes that our new knowledge will be useful. And statistical methods are very often used to predict - a linear regression is a line drawn through a number of scattered points in the hope that it says something about the trend.

Prescriptive analytics is largely concerned with resource optimisation. The techniques used are actually quite old, and thirty years ago would come under the heading of operations research. The archetypal optimisation problem involves finding the best use of resources, given a set of constraints and a well-defined objective. Airlines use this technology a great deal. How do they route aircraft, decide on the split between various classes of seats, deploy employees, and so on, so that profitability is maximised? The relationship between prescriptive analytics and predictive analytics is quite interesting. The output from predictive analytics is very often probabilistic in nature - a seventy per cent probability that sales will grow over the next six months say. Once we have these probabilities we can employ a prescriptive analytics variant called stochastic optimisation, where the input is not concrete hard numbers, but probabilities. This is very powerful, enabling the best deployment of resources in the face of uncertainty. Once we have gotten over our "predictive analytic and big data" honeymoon, prescriptive analytics will receive much more attention. In summary, descriptive and predictive analytics may tell us what has happened and will happen respectively, but prescriptive analytics tells us how it should happen.

Take a break - get a coffee or read your twitter stream (does anyone actually do that?)

Decision management is where we tie the whole thing together. I've talked about these various modes of analytics as though they are happening in a vacuum - but of course they are not. Many industries have regulators, all businesses have auditors, every business has internal politics and unless we manage our analytical environment the whole thing will collapse under the weight of its own complexity and volume of activity. You can see what I call the Analytics Triangle above. It shows



quite well the relationship between participation and skill levels. As we rise up the triangle so the number of businesses possessing the requisite skills and awareness diminishes. Decision management disciplines, tools and methods tend to get used in highly regulated industries such as banking, and in the larger players in those industries - they have to. They cannot have a regulator turn up who asks why they are giving loans to so many dodgy candidates, without providing some justification via the models they are using. And so decision management - what enterprise resource planning (ERP) is to the world of transactional activity, is an overarching set of management disciplines, tools, infrastructure and techniques which ensure that. Only a handful of suppliers provide the resources to implement decision management - FICO, SAS and IBM, and possibly other I don't know about.

So this totally 'fat free' summary of business analytics - a thousand or so words instead of three hundred pages, shows how far we need to go. And this isn't the end of it. I've already spoken with suppliers who are starting to look at the use of game theory (John Nash died last week - what a loss) to optimise and automate decisions that are being made in real-time. Predictive analytics very generously delivers probabilities - once we have these there are many things we can do, and real-time optimisation of strategy is an exciting new frontier.

In summary, this is how things are panning out. We've spent half a century automating transactional activity and business processes - we are going to spend the next half century automating business decisions - and it is going to be a profound thing. Business intelligence is just a starting point, but no matter where your organisation is in this journey, it all needs to be managed. The latest eye-candy rich data visualisation tools are quite capable of giving six people, six different



interpretations of the data. The problems become more profound the further up the triangle we rise.

And finally a piece of advice. If you want marketable skills, get into analytics. It will eventually automate much that is currently manual - and we all know what that means!