DataTechnology

Gaining The Edge

Empowering your business with AI

Today, we're spoiled by the sheer amount of data available to us. It's transforming the working world as we know it, putting more power into businesses' hands than ever before. For the better. From this pool of data, businesses can tap into all kinds of processes across customer experience, sales, HR, finance – the list is endless. But doing so requires new technologies, like artificial intelligence (AI) and Machine Learning (ML).

In this report, we look at ways to leverage these technologies, paying close attention to the business challenges they address, best practices you can apply them to, and your role in maximising the ROI from AI- or ML-driven solutions.

Chapter 1:

Translating business problems to Machine Learning

More and more businesses around the world are starting to make the most of Machine Learning, with thousands of companies already reaping the benefits of predictive analytics to increase revenue growth and boost efficiency. With all the talk of it in the media, and with artificial intelligence starting to permeate into our day-to-day lives, companies are turning to data scientists and engineers to help them make the most of this exciting technology. What many of these companies don't understand is that the success of data-science projects relies on far more expertise than just a team of data experts.

Building a Machine Learning solution that delivers business value involves both business and technical experts. This collaboration is essential from the very inception of use-case ideation so the solution not only produces accurate predictions, but that these predictions can also be used in business processes to help the business achieve its goals.

When a company looks at its performance, there are a number of KPIs to judge its successes. Behind each of these KPIs lies processes it depends on. By examining these processes, you can start to find potential use cases. Then – and only then – can you begin to find ways of improving them.

So, how do you decide on a Machine Learning use case to improve a process? For a start, it's critical that both business experts and data scientists work together to make sure they both understand the problem they're trying to solve. Predictive models empower a company to gain information that could help improve a process; information that was previously unavailable at the point of action. Only by working with business experts can you understand which information would be valuable, and therefore, what we should aim to predict.

Take a retention team, for example. Their aim is to keep customers purchasing from their company. They need to know things like, what the current retention strategy is. Do they have a way of assessing who's a flight risk? Are they just sending blanket promotional emails to their customers with the hope that some will find it in their junk mail? Are they going through a list of high-value customers and phoning them to see if they can find these flight risks by sheer force? If teams are following these processes, or similar ones, a model that can predict who's a flight risk starts to look extremely valuable. In fact, it could completely reshape how these teams work for the better.

If there's a current process that aims to find who's a possible flight risk, how well does it achieve these goals? If you can quantify the accuracy of this current method, you can then use it to compare performance with a predictive model, which gives a clear indication on whether or not a new solution is going to improve the accuracy. If not, you'll need to take a deeper look into the underlying retention processes, as integrating predictions into the retention strategy is going to represent a significant change in the way teams are currently working.

Once you've identified a use case and its potential value, you might think that sending your data science team away to get going is the right thing to do, but you've only scratched the surface of how the use case could be formed. Sending a team away here with only part of the story is likely to lead to a solution that doesn't integrate well with the business process, and can also lead to lower accuracy in models. With your new potential use case, always ask yourself the following questions about your business.

When do you need the predictions?

Timing's key for any business process. So it's really important you make predictions with enough time for the business to act. You need to answer this question as soon as you start developing a solution, as it not only dictates the result, but it also influences the criteria the models must be scored on. A highly accurate demand forecast for the next week, for example, is of no use at all if the business only has a month to react to the forecast.

What do business professionals think drives behaviour?

Data scientists are experts in the underpinning algorithms, statistics and mathematics that make up Machine Learning. Of course, these skills are invaluable. But what they often lack is the domain knowledge needed to get the best results in the use case related to the business knowledge. Industry experts will often understand what factors affect customer behaviour, or the output of a process. Business experts must communicate with data scientists early on – whether its increased sales and footfall on days with local events or good weather, or an increase in defects in a manufacturing process when using materials from a particular provider. By doing so, the data scientists can begin to find ways of representing these factors within the data set that will train the model. Once you have a data representation of these factors, the model can then learn the factors' true relationship with the outcome, with none of the bias that industry may currently hold.

How will the predictions be used by the end user?

Looking beyond gathering data and building the model, how should the resulting predictions of the model be used? Creating a model that accurately predicts the outcomes you're interested in is one thing, but making sure they're used to improve a process is another. Do you need to make predictions in batches on a time schedule? Will a non-technical user need to manually produce predictions in real time? Does an end user need all predictions back to do their job, or will this just prove more confusing than just feeding back specific predictions based on a criteria?

As the end user of a model's predictions is often someone who isn't well-versed in the world of data science, it's really important that, when they come to interact with the model – either to create predictions or action based upon them –they can do so with confidence. If you don't do this, it not only increases the risk of the end user making errors, but it will also likely lead to a redundant system, with teams doing their best to steer clear of the solution and regressing back to their old ways of working.

With all Machine Learning projects, companies are stepping into the unknown when they discover a use case. No data-science team can be certain the results of a project are going to be an accurate model until they come in from the data. Companies of all sizes still have such a wide variety of data quality and volume that often they find themselves trying to run before they can walk in their data journey. That said, if they make sure their business and technical experts collaborate from the very start of a project, they can be confident that, when they explore building a predictive solution, their vital experience has been used when they present the model's performance, and that when a model is seen to have good predictive capability, it can be integrated in a way that delivers real business value.

And one a final thought. Beyond the relationship with business and data professionals, there's much to say about the relationships between the software and data teams. How they need to work in unison to make sure that in a productionised model, the data that feeds the model can be accessed at the right time, without affecting other operational systems. But that's a subject that deserves its own post.

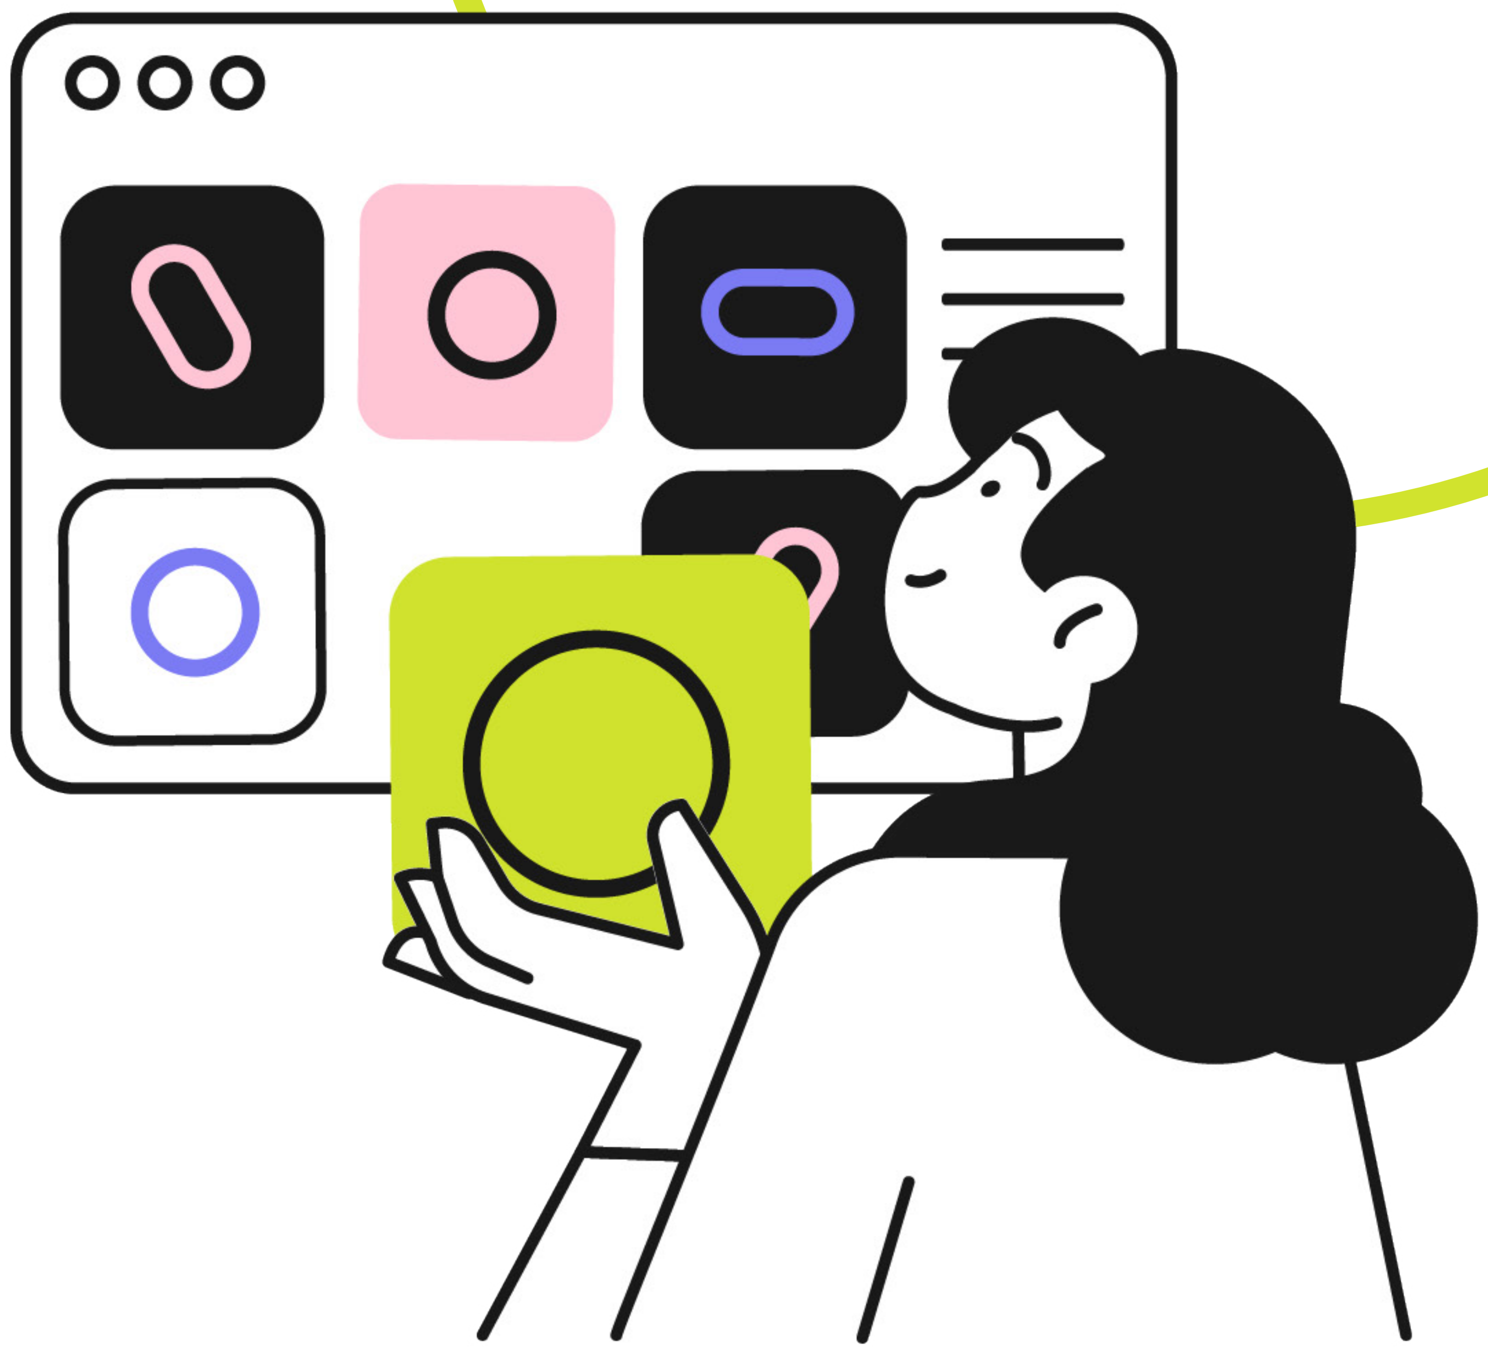

Chapter 2:

Mastering the art of Machine Learning

To be a successful data scientist in business, you need a combination of communication skills, along with both technical and domain expertise. Data scientists making the move from academia into the world of business will be faced with a range of new challenges. This post, I've taken my experience of implementing Machine Learning to achieve business outcomes to create some key considerations every data scientist should bear in mind when trying to deliver value in business.

Keep the business outcome in mind

When building Machine Learning solutions for business, always look beyond just the accuracy of models to gain a view on whether these models deliver real business value. The underlying aim of a solution is usually going to be increasing profit margins. Quantifying the amount that can be saved based on the performance of a model is essential when winning stakeholders over in favour of a new project – making sure time's not being spent on projects with no business value.

Keep the end productionised model in mind

When building and developing models for use in the real world, considerations need to be made in the pipeline that goes from data to model, to predictions to end user. Where's data coming from? Is this going to be accessible when you need it? Are you risking slowing down operational systems? All of these questions need to be considered when conceptualising a use case.

This means making sure the right teams of data engineers, database admins and so on are involved in planning – to build an accurate timeframe for development and make sure that a promising model doesn't get shelved due to constraints on the external systems that it relies upon. This also crosses over with keeping the business in mind, especially when you consider the costs associated with running and training models. Often, model training can be computationally expensive, and it leads to high bills from cloud providers each month. Getting the right balance of retraining for accuracy and achieving results that benefit the business is critical to get the best overall results from a model.

Keep the end user in mind

It's crucial you produce the output of the model in a format that the end user can work with. Often, teams resist change, especially when they're used to ways of working that have been around for a long time. Bringing these teams in for discussion on how the end result is formed makes sure the models you've worked so hard to produce actually get used. While end users may resist change to their ways of working, they're the people who know the process better than anyone else, and will likely have some valuable input that can help form your solution. This is a great way to give accountability to the end user to make use of the model, as they feel they've played their own role in forming the solution.

Communicate data issues early

The world of business doesn't have perfectly formed, readily available data sets. Data can vary in quality dramatically, both in terms of its accuracy, completeness, and accessibility. Communicating these issues early on in the process gives you the best chance to address and overcome these issues going forward. Communicating these issues makes sure stakeholder expectations are set with a timely explanation, keeping trust with the data-science team further down the line. It also allows companies to reform their data strategy so that, while a project may be possible at present, steps can be taken to give a better chance of success in the future.

Assess your model beyond accuracy

Accuracy is the key metric. But your model should be the first port of call when assessing its performance. And a further assessment should be made on your model's performance, which is linked back to the overall business outcome you're trying to achieve.

Say a demand forecast for non-perishable goods is predicting with 90% accuracy, but when assessing the bias of a model, you see that it tends to underpredict demand rather than overpredict. If you were to blindly follow this prediction when making stock orders, the business would quickly find it's failing to meet the demand of its customers. Not a good outcome at all. Solving this problem could involve going back to the modelling stage to try and rectify the problem through bias correction, or perhaps a simple post-processing addition to the prediction would suffice. If you're considering the predictions of perishable goods, then the assessment of what's best becomes a bit more complex, and you need to quantify the costs of missed sales vs waste, and maybe even unquantifiable outcomes such as your business's reputation.

For instance, take a model prediction that aims to predict customer churn. Initial accuracy might be low, where many more customers are predicted to churn than actually do. The temptation of a data scientist here may be to say that the model shouldn't be used. However, when looking at the business process, it could offer real value. Suppose the model's recall (the number of customers that were predicted to churn correctly) is very high, and the previous business process had no view on the likelihood of each customer to churn. In this case, you've now managed to narrow down the list of customers to reach out to significantly, allowing your team to be much more effective in its retention strategy.

The world of data in business is often a chaotic one, riddled with legacy systems, data-entry errors and siloed visibility on the operation as a whole. While strong technical skills are essential for the success of data-science projects, navigating the challenges outside of the mathematics of Machine Learning is just as important.

Chapter 3:

Driving ROI from AI

Whenever a business is looking to implement Machine Learning, a return on the investment from the development and maintenance of models should be a key consideration. This return is going to be gained primarily by either increasing revenue or improving efficiency. And these two fundamental objectives lie at the centre of growing businesses' bottom line.

Once you've identified a use case, and you're in the ideation stage, you need to quantify a few things. Firstly, examine the process that the model is aiming to improve. Perhaps, you're trying to find a way of making sure adequate stock is in store for purchase (increasing revenue), or conversely reducing the amount of waste that comes from out-of-date stock (improving efficiency). Start by quantifying the performance of a baseline system. How accurate is the current forecasting process? How often does this lead to the unavailability of a product or waste? What's the cost associated with these scenarios? You'll only be able to assess the performance of your new system once you have a metric that quantifies the performance of your current one.

If you can quantify the costs associated with either lost business or wasted stock, you can then start looking at how improving forecast accuracy can reduce these scenarios, and therefore how much you could save by implementing Machine Learning. On the assumption of a perfectly working supply chain, a perfectly accurate demand forecast should lead to a zero-waste and zero-lost sales scenario. Of course, a perfect demand forecast isn't achievable. But a relationship can be drawn from the accuracy of a forecast, and the number of missed sales or wasteful scenarios you encounter. This will allow you to get an estimate on how much you can save based on the model's accuracy.

Now, when data scientists start building prototype models to assess the potential accuracy gains when using Machine Learning, they can gain a rough estimate of how much real value this can deliver to the business.

Beyond these metrics for success, your business could also significantly improve its reputation – both a more reliable supplier and as a waste-conscious company. While these improvements aren't easily quantifiable, remember to consider them when you're deciding on whether to go forward with a new solution.

Both overarching improvements can be approached in several ways. So let's explore how you can apply Machine Learning to improve these factors.

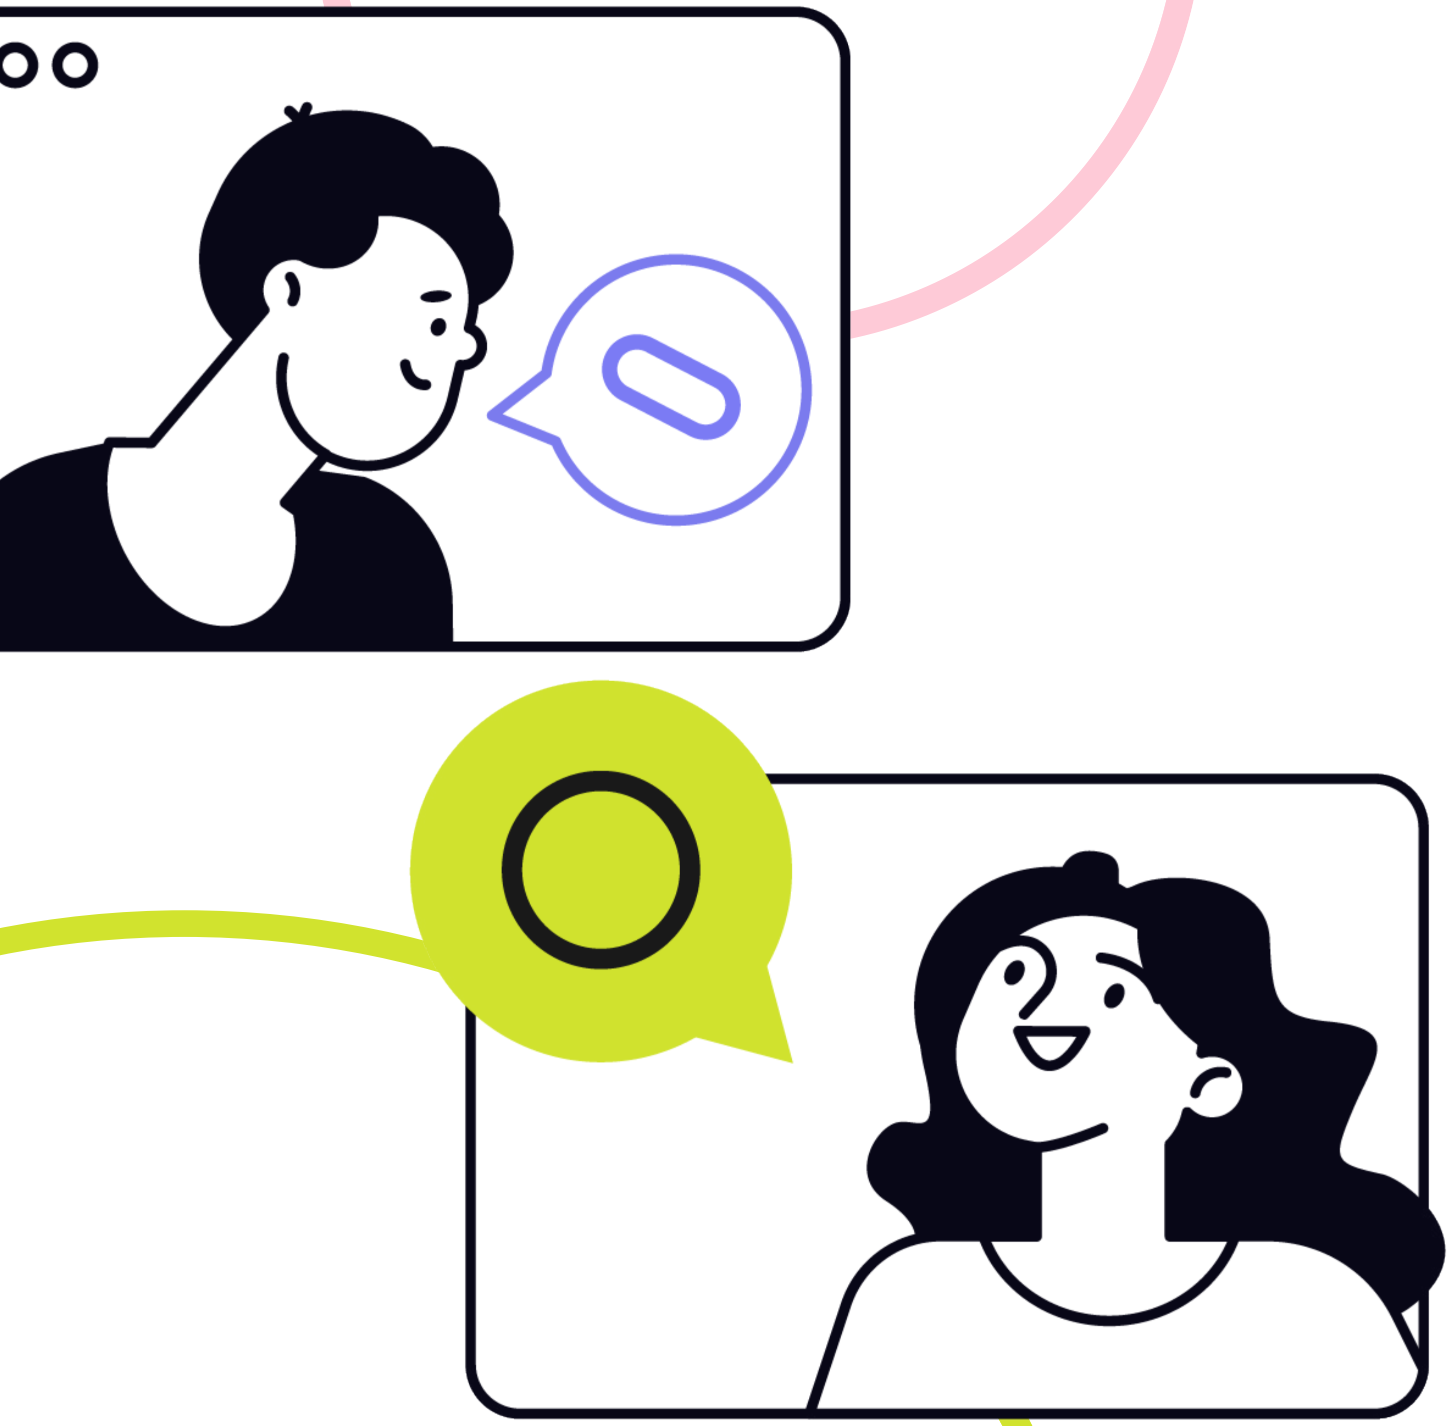

Generating new custom

Growing the pie is the most obvious way to increase revenue.

Believe it or not, it all starts with improving the way you target existing customers. Winning new customers is a great way of increasing both the volume and robustness of your revenue stream. But expanding the products and services used by existing customers is a proven method of effectively increasing revenue. Businesses will often send promotional information to their existing customers through communication channels such as email, or even by picking up the phone and calling. Without a method of predicting what each individual customer is likely to be interested in, such methods risk wasting valuable time in the case of calling, or for the customer to send your promotional emails straight to their junk folder once they get tired of irrelevant spam.

When you're trying to quantify the gains in response to promotional material, it's difficult to do it ahead of time, as a comparison in promotions can only be accurately done after it's concluded. A better way may be to quantify the expected efficiency of the process when you look at the hours given by marketing teams. How long does it take to create an effective list of products and people to target? How quick can this process be done with a Machine Learning algorithm to the same complexity? This is an easier thing to quantify ahead of time, and it may give the motivation needed to start the journey.

Machine Learning can help identify the products customers are most likely to take an interest in, as well as the best channel of communication to get their initial engagement. This leads to far more effective marketing campaigns, which result in wider sales to the customer, in much less time.

Increasing efficiency

The other side of the balance sheet lies costs. And reducing costs is just as important to growing your bottom line as increasing revenue.

In the manufacturing space, an understanding of when a machine is likely to fail, and the type of failure, is likely to be a game-changer when improving the overall efficiency of the manufacturing process. With sensor data on both machinery and the environment that these machines operate in, along with the time and nature of past failures, Machine Learning models can help improve the maintenance schedules of manufacturing plants, making sure that adequate stock of replacement parts are kept in stock.

Quantifying these gains should be based on the costs associated with machine failure, and time being out of action. The historical data of machine failure and maintenance downtime should allow a company to understand the costs associated with it. A reduction here is going to lead to more uptime and greater efficiency, so by comparing the expected reduction in downtime – due to predictive maintenance – should be able to give you an indication of what gains can be made.

Understanding the ROI for implementing a Machine Learning problem is extremely helpful when you come to qualifying the usefulness of a Machine Learning solution. At times, this may be easy to do, but often it means you need to implement a system to see its performance. In these scenarios, creating pilot groups to A/B test a new method over an amount of time may be the best you can hope for. Whether it's simple or complex, understanding where the gains are theoretically going to come from should be clear. And you always need to qualify a cool sounding idea.

Chapter 4:

How Machine Learning can help cut customer churn

Who do you put first: new prospects or existing customers?

You might not know it, but the relationship between you and your existing customers is everything. After all, the moment they lose interest in what you're selling, or they become unhappy, they'll move on. And your loss is your competitor's gain.

Marketing Metrics claims that the likelihood of selling to an existing customer is 60-70%, while the chance of selling to a new prospect is only 5-20%.[1]

With that in mind, it's fair to say customer churn and retention is a pretty big deal. It's a topic that comes up in pretty much every business meeting, with colleagues discussing ways to analyse their customer-churn data faster and more accurately, giving them the best chance of reducing churn.

Take a business that wants to improve its retention strategy. Currently, it reaches out to customers with surveys with a poor return rate, and sends blanket promotion emails to groups based on basic historical analytics, or to its entire customer group. The business also has capacity for a more focussed effort from the team's capacity. But with no idea how to prioritise who to reach out to, the best way forward is to work through the long list of customers, which – even with the best efforts – only puts the smallest dents in the entire list.

By taking the data of customer demographics, purchase history and historical contact, a Machine Learning model can help identify those who are the biggest flight risk, giving the team a focus around who is best to contact. Taking it a step further, analytics on the responsiveness of different outreach methods can also find the optimal way to contact these customers to maximise the chances of a good outcome.

Beyond this improvement to the process of the retentions day-to-day team, analysing the importance of the model's feature also allows the team to explore the complexities that create flight risks at a much quicker pace than traditional analytics, helping to guide the business in its long-term strategy around the goods and services it provides.

That's why, today, it's so important to truly understand your customers' behaviours and how they feel about your business. And that's something traditional analytics can no longer tell you. Even if it could, you wouldn't get the results when you need them.

With Machine Learning solutions, you can delve deep into big data and identify customers who are thinking of leaving – before they get out the door. This gives your customer team the chance to target at-risk customers with their pain points in mind, and better prepare to significantly reduce churn rates and keep them by your side.

Many businesses are already putting Machine Learning solutions to good use. In fact, a report by Harvard Business Review found that almost half (48%) of organisations surveyed use Machine Learning and artificial intelligence in their marketing and sales processes to gain better insights into both their existing customers and potential prospects.[2]

Marketing Metrics claims that the likelihood of selling to an existing customer is **60-70%**, while the chance of selling to a new prospect is only **5-20%.**

[1] https://www.superoffice.com/blog/reduce-customer-churn/
[2] https://hbr.org/resources/pdfs/comm/microsoft/artificialIntelligencefutureofsalesandmarketing.pdf

Chapter 5:

Using Machine Learning to meet customer demand

It's no secret that without customer demand there's no business. That's a given. And it's why demand forecasting has become more important than ever.

Ask yourself: how much are you going to sell this quarter? How many customers are going to walk in store, or go online? They're big questions that need answering. Something an accurate, Machine Learning-powered forecast can do for you. And it can have a huge impact on your bottom line and profits, too.

In sales, Machine Learning forecasts give you greater control over inventory management, meaning you'll have products readily available on the shelves, without overspending on things like storage – or worse, risking any perishable goods going to waste. In fact, PwC cited inventory and delivery management, as well as anticipating customer demand, as two of the three key areas where Machine Learning and AI has significant potential.[3]

The same goes for managing your workforce. Specifically, having the right people, in the right place, at the right time – ready and waiting to serve customers and meet their demands. We call it smart scheduling. Where employee downtime comes down, and productivity levels come up.

Having accurate forecasts across inventory and workforce management is what your Finance team needs to gain the full picture of your business's projected performance. They can then analyse whether you're on course to meet (or even exceed) the goals you set for the year. And if you're not, you can proactively address the situation and get back on track before it's too late.

Predicting customer demand is never going to be perfect. There will always be factors in the world outside of our control. No model for example could have quantified the effects we saw in the Covid-19 pandemic. But Machine Learning solutions can bring you closer to that goal than ever before – giving you a much greater wealth of information to consider when creating a forecast and removing the human bias inherent in manual intervention.

For example, a supermarket requires not only a great supply of products to make sure it meets its customers' demands, but it also needs to keep a keen eye on wastage that eats away at the bottom line. Traditional methods of forecasting may look at sales from the previous year and apply a small increase or decrease depending on how the business thinks it's performing that year.

But how well does this really match up with the variance in consumer demand? Things like weather, local events and close proximity competitors would typically have been handled by manual forecast adjustments, costing time and incorporating the bias of the data analyser making these adjustments. A Machine Learning model can help eliminate this bias by observing at scale how demand changes to these events, and finding their true relationship to the outcome.

[3] https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf

Chapter 6:

Your workforce, powered by Machine Learning

Many people hear the words Machine Learning and artificial intelligence and instantly think of robots replacing humans. But that's simply not the case. It's about trusting technology to help you manage your existing workforce so you can get the most out of them. And they can get the most out of their roles.

That goes for all areas of the business: the shop floor, the warehouse, your management teams, admin – the lot. Every single employee needs to be in the right place, at the right time. And Machine Learning solutions can help you do just that.

Traditional workforce management processes can only get you so far. And they're heavily time-consuming, and quite often, they provide inconsistent steps. Not really the dynamic workforce scheduling you need, given all the economic uncertainty we're now facing.

The good news is technology's advanced significantly in the workforce space. Advanced analytics, artificial intelligence and Machine Learning have empowered people teams in all kinds of businesses – giving them actionable, predictive insights. No more counting heads. No more error-riddled spreadsheets. No more gut-based decisions.

The results speak for themselves. Research from IBM found that companies leveraging these technologies for workforce planning outperformed those that weren't by 37%.[4] And by 2035, it's predicted that AI will increase business productivity by up to 40%.[5]

Machine Learning and data-driven workforce planning might be the way forward. But don't forget about the important role your people still need to play. These projects aren't easy to scope, develop or deliver. They require a diverse range of skills and collaboration between data specialists, project managers and engineers. Only then will you fully maximise the rewards of Machine Learning-driven workforce planning.

Balancing workforce management has analogous goals to demand, making sure that work can be done with the minimum amount of people needed. Let's imagine a retail brand that prides itself on customer service and good stock. The needs in the warehouse will fluctuate with the needs on the shop floor, and both have to run efficiently to make sure the company operates as intended. Footfall data enhanced with external factors like economic events and local competitors can help build models that transform the workforce management process for the better on the shop floor. And, at the same time, data around the types of deliveries and historical time can move the different classes of stock – giving you a far more nuanced understanding of what it takes to run a warehouse.

The additional dimension to workforce management is, of course, staff. Right now, all kinds of industries are experiencing staff shortages. That's why your staff's satisfaction and happiness is more important than ever. Having a future-focused view of workforce requirements allows you to build more robust rota schedules earlier on, eradicating last-minute shifts being assigned to or taken away from your staff. And that's going to create a far more attractive workplace.

By 2035, it's predicted that AI will increase business productivity by up to 40%.

[4] https://www.shrm.org/executive/resources/people-strategy-journal/Fall2019/Pages/moore-bokelberg-feature.aspx
[5] https://draup.com/talent/blog/ai-based-strategic-workforce-planning-unlocking-the-benefits/

DataTechnology

Get in touch

No matter your business challenges or objectives, we can help you make the most of Machine Learning solutions to gain that all-important edge in your industry.

Get in touch today on **0845 6588 480**