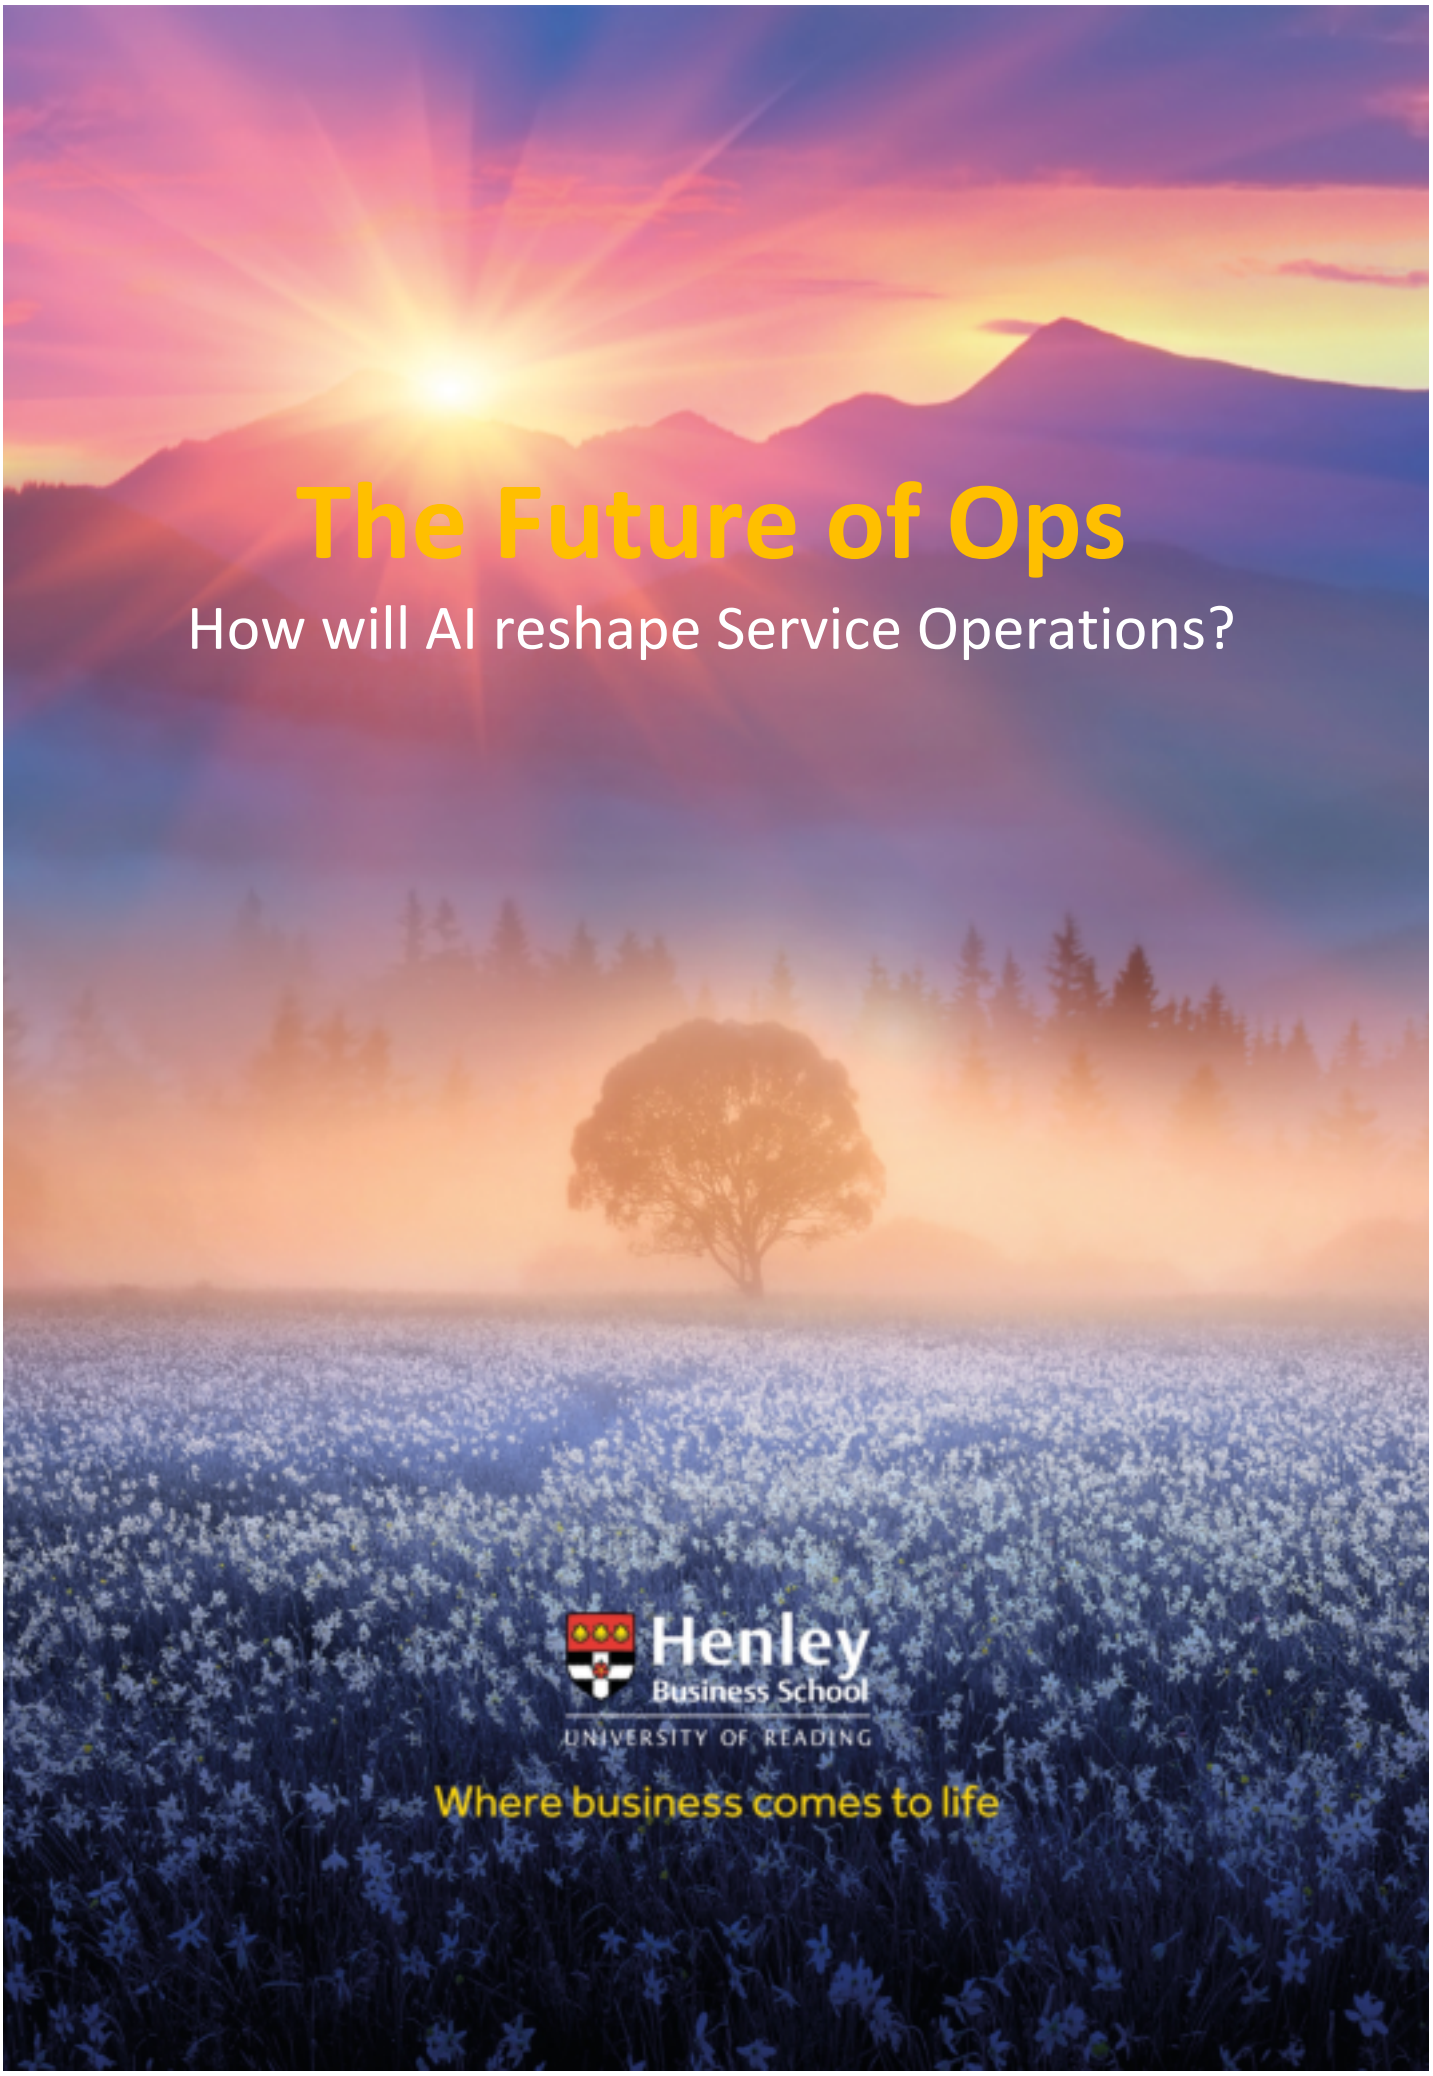
The Future of Ops
How will AI reshape Service Operations?
Henley
Business School
UNIVERSITY OF READING
Where business comes to life

Contents

Foreword

The avalanche of interest in generative AI is poised to entirely change the world of service operations. Many ops leaders are realising the potential for AI – both generative and non-generative – to ease the constant pressure to reduce costs and work more efficiently, even with fewer resources at their disposal.

While operations teams are already using technology such as robotic process automation (RPA) to automate certain tasks and streamline workflows, workloads have remained resilient – often because there is a limit to what can be done without human intervention.

The next generation of AI tools promises to change all of that. With increasingly 'intelligent' automation at their disposal, organisations will need fewer people and become far more data-centric, with AI generating prescriptive insights that will help organisations improve their decision-making.

In this white paper we examine the challenges service operations face to unlock the power of AI and explore the potential AI use cases that will change the way operations are run moving forward.

Dr. Mona Ashok, Associate Professor of Digital Transformation
Henley Business School

The AI opportunity for service operations

When OpenAI unveiled its latest version of its generative artificial intelligence (AI) tool ChatGPT late in 2022, the hype around AI went into overdrive as organisations rushed to reimagine how the technology could transform their business.

PwC has estimated that AI could contribute $15.7 trillion to the global economy by the end of this decade[1]. By next year, Gartner estimates that 40% of enterprise applications will have embedded conversational AI, up from less than 5% in 2020[2]. And by 2026, Gartner believes more than 100 million humans will engage with 'robocolleagues' to collaborate on work projects. Multiple industries are likely to face seismic disruption as a result.

While this promise of future innovation offers a tantalising glimpse of how advanced AI could transform businesses (though the potential ethical issues, would require a separate paper in their own right), it is easy to forget that AI and machine learning (ML) tech is already being used by many organisations to automate processes and reduce manual tasks. However, the key difference is that the previous generation of AI and ML tools had to be programmed to perform specific tasks, while the new generation of AI learns and adapts to different scenarios. The optimism around generative AI tools such as ChatGPT is likely to reignite interest in AI and ML more broadly, increasing adoption of existing applications to boost automation and improve customer engagement.

---

[1] https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html

[2] https://www.gartner.com/en/topics/generative-ai

There are a number of ways businesses are integrating AI into their operations. Take UK renewable energy business Octopus Energy. Almost half of its customer enquires (44%) are now handled at least in part by AI[3]. Ride-sharing platforms such as Uber and Lyft are using AI for operations management, for instance by optimising routes and enabling demand-based pricing. Many financial institutions are also using AI to improve fraud detection by reducing the number of false positives (when banks erroneously block legitimate payments).

According to a recent IDC survey, companies are likely to focus their AI investment in three key areas[4]. Top of the list is improving operational efficiency, followed closely by improving the customer experience and then improving employee productivity [Figure 1]. That desire for greater efficiency and productivity will be enhanced further by the advent of generative AI. For example, generative AI has the potential to completely transform the operations function because of its ability to analyse vast quantities of data but also give users access to more relevant and prescriptive insights, potentially identifying issues that otherwise may have been overlooked. This proactive element – AI being capable of serving up unsolicited recommendations – is a significant departure from previous iterations of the technology where AI would need to be given specific instructions to complete a task.

Figure 1 – IDC AI StrategiesView 2022

By using generative AI to support operations management, teams can unlock efficiency and productivity gains by optimising resources and planning with greater accuracy, so workers have less idle time.

All of this matters because service operations are viewed as a giant cost centre for businesses (for a bank, it would typically account for about 20% of costs), with operations departments under constant pressure to become leaner and more efficient.

As tech and AI advance, there will be a push from organisations to increase automation and decrease the size of their operations teams. However, for AI to power this shift, data and Decision Intelligence will be critical; an area where operations teams are significantly lagging.

[3] https://news.sky.com/story/artificial-intelligence-to-hit-workplace-like-a-freight-train-energy-boss-warns-12883712

[4] IDC AI StrategiesView 2022

The need for Decision Intelligence

Using data effectively means being able to analyse it and generate insights that can improve decision-making capabilities. This is what Decision Intelligence is: - the ability for companies to combine AI and human intelligence to deliver predictive and prescriptive insights for service operations to help them make better decisions.

The challenge that many organisations face is that they are still in the very early stages of their data journeys. A majority of service operations teams can't use their data to make real-time decisions. Many are still stuck in a world where processes are entirely manual, everything resides in spreadsheets, and where there is no common language for data terms. Others might have progressed to using basic dashboards, but the data is usually siloed across multiple systems, making it difficult to access. The data also tends to be backwards-looking and needs to be interpreted by humans (who often lack appropriate data literacy and skills).

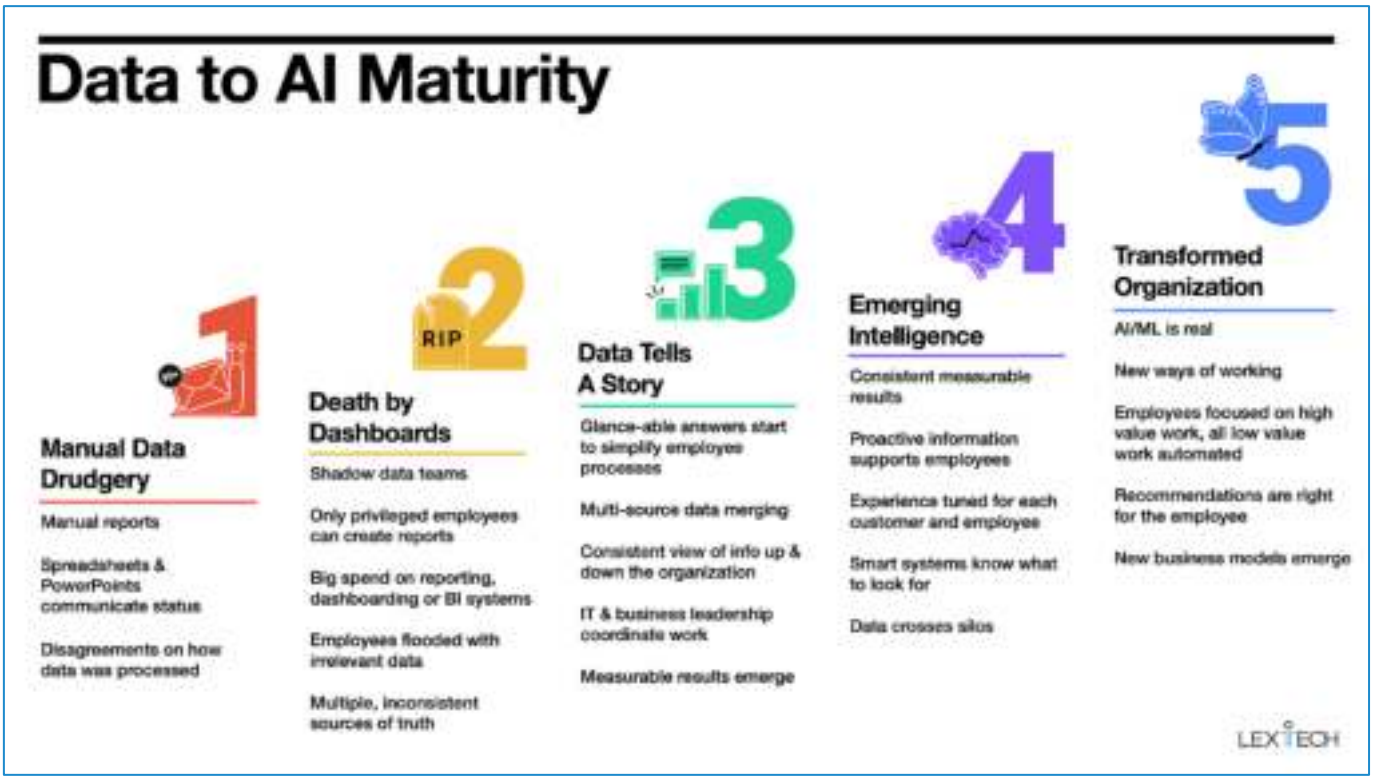

Figure 2 – Data and AI Maturity Curve, Lextech

There are typically four stages that will indicate where organisations are on their data maturity journey[5] [Figure 2]. The first stage is descriptive – data that tells users what happened. The second is diagnostic – data that tells users why something is happening. The third stage is predictive – data that tells users what might happen. And the final stage is prescriptive – data that gives users suggestions of what actions they should take.

To advance on this data maturity journey, organisations need to start making their data work for them, ultimately generating predictive and prescriptive insights that support real-time decision-making[6] [Figure 3]. This enables operations teams to make decisions based on forward-looking intel and data-driven advice with little or no manual

[5] https://towardsdatascience.com/how-to-measure-your-organizations-data-maturity-2352cbaf1896

[6] https://www.gartner.com/en/newsroom/press-releases/2014-10-21-gartner-says-advanced-analytics-is-a-top-business-priority

intervention needed. Only by getting their data in order first, can organisations unlock the power of AI for service operations – but many are nowhere near that stage of the journey.

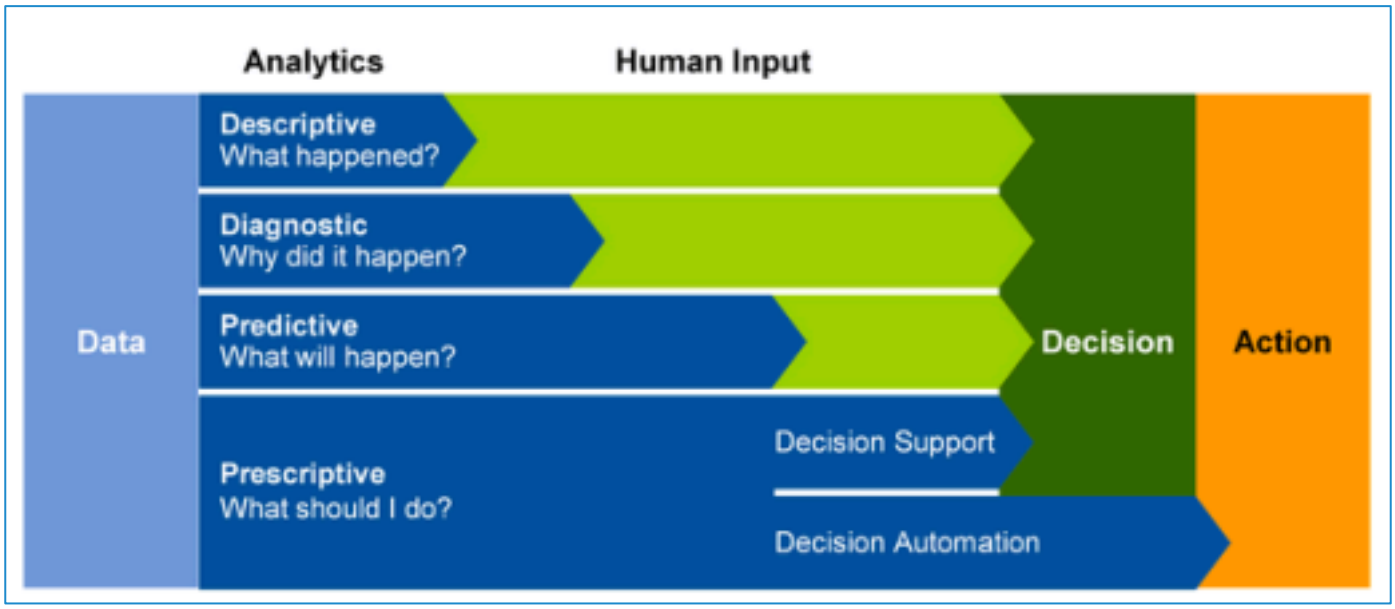

Figure 3 – Four types of AI capability, Gartner

Unlocking smarter Decision Intelligence

To access good Decision Intelligence, organisations need good quality data. Put crudely, if you put garbage in, you get garbage out. To avoid that, organisations must ensure they have access to the right data to plug into their AI systems to get accurate results. If they don't have their data in order, or the data they are feeding into the tool is wrong or biased, then the output or insights the AI generates will be misleading.

Boston Consulting Group's research shows that pioneers of AI typically commit a significant proportion of AI effort, (70 percent) to prepare the business and people (transformation) to adopt AI; 20 percent is committed to data and technology, and 10 percent to algorithms. Thus, organisational readiness to embark upon AI is essential.

There are four typical data challenges that organisations face: technical integration, a lack of data literacy, poor data practices and a lack of sufficient data investment in ops.

1. Technical integration

Many organisations, particularly banks, are saddled with legacy technology that has been built up over decades and operates in isolation – meaning data often sits in silos. This lack of systems integration makes data difficult to capture, access, extract and manage, increasing the risk that organisations won't be able to unlock the value of their data. Some organisations also lack a clear structure around who is responsible for maintaining and updating their data, including keeping it secure (a lack of clarity about who is in charge of managing the data can increase the risk of data breaches). Even if data is hosted by a third party (as is increasingly common with cloud-based applications) organisations are still responsible for the safety of that data. Customers also expect a consistent experience when interacting with a business. This means organisations need to be clear about who has the responsibility to help operations teams get the data and insights they need.

2. Data literacy and skills

Data skills are fundamental in any IT change so that organisations can benefit from new technology. Without appropriate data literacy or knowledge of how the tools work, organisations won't be able to integrate them into their workflows – yet organisations tend to lack these skills. Almost two-thirds (60%) of operations leaders say their teams don't have the necessary skills that service operations will need in the future, according to a Gartner survey[9]. A further 62% say it is difficult to find the right talent to fill emerging ops roles in areas such as process redesign, digitisation, automation, and data and analytics. Even if organisations are using external providers or consultants to help them adopt new technology, those organisations still need in-house capabilities and expertise to make sense of the data and ensure they are using the systems properly, so they know the insights being generated are reliable.

3. Data practices and insights

Even if organisations have access to good data and the right skills, sometimes they don't have the appropriate practices in place to make use that data. For example, organisations can often lack a common set of definitions for key metrics, which means different departments or even teams within departments might measure those metrics in different ways, making the data ineffective and skewing any insights that are generated. Organisations therefore need to ensure that definitions are consistent across the whole of their operations. Greater trust in data is also needed for operations teams to have faith in the insights AI is generating. This means ensuring there is transparency and clarity around how insights were generated (so ops teams don't have to double-check check the output is accurate).

4. A lack of investment in ops

Many organisations view operations purely as a cost centre. In an ideal world, some would want zero operations if they could get away with it, relying entirely on self-service. Therefore, operations are not seen on par with cash-generating parts of the business, suffering from not only a lack of investment but also a squeeze on their budgets as they are pushed to do more with less. But by investing in data within ops, organisations can better understand the impact the operations function is having on the rest of the business – both in terms of customer experience but also the impact on the bottom line.

Getting to grips with data is critical because, once it's been done, organisations can start to unlock real-time insights and the benefits that come with that. In service operations, real-time data enables ops leaders to react to shifts in workload faster, allowing them to flex resources by moving workers to different departments as volumes fluctuate. This can then allow operations teams to make real productivity and efficiency gains – but that is only possible if they can effectively manoeuvre resources in real-time and that hinges on where they are on their data journey.

Tips to get your data in order

If you want to get the most out of advanced AI for your organisation, then your data needs to be fit for purpose – here is the best way to get started on your data journey.

Definitions

The first step you need to take is to agree on your definitions and ensure there is consistency across the organisation. If one department records a particular data parameter in one way, and another labels it differently, then the data is not going to be consistent or accurate.

Understand what data you need

Organisations are sitting on a vast ocean of data – but not all of it is useful. By understanding the question you are trying to answer with your data, then you can work out what data you actually need, making it easier to find that data in amongst all the rest.

Understand where the data is

Once you know what data you need in order to generate the insights you want, you then need to figure out where that data is residing and how to extract it. Many organisations have multiple systems, with data often stored in hard-to-access silos, making it difficult to retrieve and process.

Practical steps

Getting your data in order is a huge, time-consuming undertaking. Breaking it down into manageable chunks will make the process easier and ensure organisations don't give up halfway through because it has become too complicated to complete.

---

[8] https://www.bcg.com/capabilities/artificial-intelligence

[9] 2020 Gartner Financial Services Operations Workforce Survey

The Killer AI Apps that will Transform Ops

Once service operations have a better handle of their data, they are then ready to start reaping the benefits of using advanced AI. In this section, we look at the key areas in running service operations that are ripe for AI: automating routine tasks (to eliminate manual drudgery), data analysis (identifying data trends that would otherwise be difficult to spot), generating predictive insights (to help forecast what is likely to happen in the future) and providing prescriptive advice (to enhance decision-making).

With those areas in mind, we have compiled a list of five killer apps that have the potential to reshape the operations function.

1. Next-level RPA / Intelligent Process Automation

Many organisations have already embedded RPA into their process flows, helping automate routine tasks and boosting the efficiency of their ops teams by allowing them to focus on more complex work.

How is it traditionally done?

To date, RPA has only been capable of handling basic tasks. In the US healthcare sector, for example, about 90% of insurance claims are handled using RPA without any human intervention, but the more complex 10% of claims still need to be processed manually. RPA 'bots' are also notoriously easy to break – if there are any changes to how the data is presented, the bots will hit a brick wall and won't be able to complete the task because they won't recognise the changes (they can only execute tasks according to rigid rules-based commands).

What are the challenges with this?

In the case of complex claims, the 10% of cases that demand manual intervention still require thousands of workers to process them. And if an RPA bot fails to complete a task because it can't understand any changes to the source data, a human will be required to reprogramme the bot, increasing time needed for ongoing maintenance.

How can AI solve it?

Advanced AI has the potential to make RPA more sophisticated by enabling bots to adapt and learn when encountering unexpected data parameters without the need for manual intervention. As bots get smarter, they will also be able to handle more complex casework, allowing operations teams to become leaner and scale their robot workforce up or down as work volumes dictate.

Benefits

- Enables increased automation as RPA bots can 'intelligently' respond to events through self-learning
- Reduces failure rates and avoids unnecessary human intervention, improving speed and turnaround times
- Less maintenance expense and increased operational efficiencies will help reduce costs

2. Real-time skills catalogue

Building a skills catalogue to track the competencies of operations staff is essential for improving flexibility and responding better to shifts in demand across different operations departments. This ability to share resources across teams is critical for organisations to remain agile and better manage peaks and troughs in workloads.

How is it traditionally done?

Historically, maintaining a skills catalogue has been a very manual and static process.

What are the challenges with this?

Almost from the moment a traditional skills catalogue is created, the information in it goes out of date. As a result, operations leaders are never sure if a particular employee still has the ability to carry out a certain task productively or not (maybe because it has been several months since they last performed the task and are now too rusty to be effective). And because a skills catalogue is generally managed manually, any attempt to keep it up to date and therefore reliable takes a massive amount of time and effort.

How can AI solve it?

By using AI, organisations can build a live skills catalogue that tracks tasks being performed and by who in real time. With this information, the organisation gains an up-to-the-minute view of how suited a particular individual is to perform a specific task. The catalogue can record not only how recently an individual has completed a task, but how

fast they completed it and to what standard. An AI-powered skills catalogue can could even automatically assign tasks based on competency without any manual input, and ensuring ensure employees get regular exposure to tasks so their skills don't atrophy.

In addition, AI can identify potential skills gaps and recommend areas for upskilling based on an organisation's priorities (customer needs, regulatory requirements, and so on). This ultimately builds agility levels and creates cross-skilling opportunities that can improve job satisfaction and increase career progression opportunities.

Benefits

- A live skills catalogue helps organisations increase their agility by sharing resources across teams more efficiently to better manage fluctuations in work volumes
- Greater agility helps unlock spare capacity and boost productivity
- Cross-skilling staff means more work can be done with fewer resources, reducing operating expenses
- Cross-skilling staff also increases employee engagement, improving career prospects and wellbeing
- AI can track what work is completed without manual intervention, ensuring skill levels are always up to date

3. Capacity planning

To run operations effectively, ops leaders need 'shared planning tools' that give every department a consistent view of estimated work volumes over the coming week. This enables more effective resource planning, so the right number of workers are assigned to particular tasks according to expected demand.

How is it traditionally done?

Frequently this process is carried out in a slightly different way by different managers. Some will use Excel, while some may use other systems or software, meaning there is no consistency in the way resource planning is managed. The level of ops leaders' experience can also impact planning accuracy – some managers are better at anticipating workloads and allocating appropriate resources than managers who may be slightly less experienced.

What are the challenges with this?

Often managers will plan for the worst-case scenario, which means they tend to overcommit headcount for the amount of work they actually need to do. The easiest way for organisations to unlock spare capacity is to ensure resource planning is optimised according to expected volumes so teams have sufficient work and aren't overcommitted or understaffed. If ops leaders don't have a consistent view of what is going to happen in the week ahead, organisations end up with wasted capacity and additional costs. Manual planning also takes time, decreasing efficiency and increasing the risk of inaccuracies.

How can AI solve it?

AI can enable smarter planning by building up a vast database of tasks and measuring how long each component of that task takes to complete, providing a consistent benchmark for what is needed to process a particular piece of work. The AI can then look at expected volumes for the week ahead (based on historical data) and then assess resources needed based on the tasks that will be involved and how long those will take, ensuring plans are aligned with and will meet any service level agreements (SLAs). This provides operations leaders with early insights to make decisions ahead of time rather than reacting as events unfold. By using AI-generated forecasts, organisations can ensure they are basing those decisions on accurate and consistent information. It also enables much faster planning and improves efficiency (removing the need for managers to block off time to work on their plans every week).

Benefits

- AI can help improve the accuracy of weekly plans by 30% to 40% when compared to manual planning
- More accurate planning can help unlock capacity gains of around 3% to 4%
- AI ensures plans are always consistent and not subject to an individual's bias or inexperience
- Better capacity planning helps reduce costs because operations teams are no longer over-staffing for the worst-case scenario

4. Predicting case outcomes

Certain types of casework (insurance claims, for example) don't advance in a linear fashion, with different tasks running concurrently with and no clear timeframe attached. That makes planning problematic – with so many variables, that it can be almost impossible for a human to predict the likely outcome.

How is it traditionally done?

Operations managers can give a rough timeframe for work in progress based on teams of people feeding data into spreadsheets and making calculated projections, but the complexity of case outcomes (given the thousands of different directions any particular case could go in) means it is an imperfect and time-consuming process.

What are the challenges with this?

Organisations need real-time case status visibility because they need to monitor how cases are progressing and what will likely happen next, helping operations leaders better track if SLAs will be met or breached. Without a clear idea of how long a case will take or how challenging it will be to complete; ops managers won't know where to allocate their resources. These risks can result in 'zombie' and 'vampire' cases – zombie, because the cases just meander aimlessly without advancing, and vampire, because they suck up a disproportionate amount of time and resources.

How can AI solve it?

By looking at the path taken by hundreds and thousands of cases, AI can predict the potential journey of a case on what has happened with previous cases that share the same or similar characteristics. The AI can then plot what is

likely to happen next at any stage of that journey and how long it is likely to take. This will enable operations leaders to allocate appropriate resources to cases that would otherwise risk breaching SLAs, ensuring the right people are doing the work (so organisations aren't spending money on unnecessary resources) and avoiding cases becoming zombies or vampires, helping ops teams become more efficient.

Benefits

- Better predicting the likely path a case will take can help avoid over-staffing, therefore improving capacity management, boosting productivity and reducing costs
- Understanding where a case is heading next can also reduce risk because operations leaders can identify problem cases sooner and avoid breaching SLAs by allocating resources as necessary

5. Prescriptive decision-making

Organisations always strive for consistency in their decision-making, but the quality of those decisions will ultimately hinge on the experience and competence of the individual in question. Inconsistent decision-making is a problem for organisations, potentially creating operational inefficiencies and increasing risk.

How is it traditionally done?

Different people make decisions in different ways. Someone with decades of experience under their belt will likely make better decisions than a new starter, which inevitably means the quality of decision-making will vary across operations teams.

What are the challenges with this?

Inconsistent decision-making can cause a wide range of issues for operational performance, from wasted capacity (poor planning that results in overstaffing relative to work volumes) to low productivity (workers are overstretched because not enough capacity is made available).

How can AI solve it?

AI will enable organisations to improve the consistency and quality of their decision-making by combining the collective experience and shared knowledge of operations professionals worldwide – surpassing what is possible even with their best employees. By giving operations teams access to AI-generated insights and prescriptive advice, decision-making will be consistent regardless of the experience or competence level of an individual. Microsoft, for example, has introduced an AI-powered assistant for 365 users called Copilot – akin to a supercharged version of its classic 'Clippy' office assistant. Having a similar virtual expert on hand can help operations professionals make better decisions – either by following AI-powered advice or potentially by fully automating the decision-making process, with the AI assistant identifying problems, choosing the right solution and then automatically solving it without any human intervention.

Benefits

- AI can help achieve a higher level of consistency with operational decision-making
- Using prescriptive advice can help reduce costs by avoiding wastage caused by poor or inconsistent decision-making
- Prescriptive advice can also reduce the operational risks associated with bad judgement or human error if decisions are being made by inexperienced staff

Conclusion

The advance of generative AI technology like ChatGPT has the potential to reshape operations departments by helping ops leaders work smarter, improve efficiency and boost productivity.

With ops departments under increased pressure to cut costs and do more with less, the need for more automation is clear. This in turn will shine a spotlight on the importance of having high-quality data.

Without high-quality data, the insights generated by AI tools will be unreliable. As organisations become more dependent on prescriptive AI-generated insights, the consequences of poor data will grow.

This means organisations must get their data in order if they want to unlock these benefits of AI – and that often involves tackling four key issues:

Technical

Many organisations have sprawling legacy tech systems that mean much of their data is stored in silos, making it difficult to locate and access.

Data skills

Organisations need to ensure their staff have sufficient data skills so they can actually make use of new technology.

Data practices

Organisations must have robust data practices in place so that end users can interpret AI-generated insights in the context of their business.

Underinvestment

Organisations have often shunned ops department investment, viewing operations as a cost to be cut. But by investing in the right tech and data, organisations can start to properly understand the value that ops bring to the wider business and the benefits of greater optimisation.

By solving the data issue, organisations can start to adopt advanced AI tools in a meaningful way and reap the benefits that prescriptive insights can bring to operations management, increasing agility levels, powering efficiency gains and ultimately boosting productivity, helping organisations get the maximum value out of their workforce.

About the Author

Dr. Mona Ashok, Associate Professor of Digital Transformation
Henley Business School

Mona is a Fellow of the UK Higher Education Academy, a member of the Institute of Chartered Accountants of India, a Certified Software Quality Analyst, and a Six Sigma Master Black Belt. She has 20 years of corporate leadership experience in global IT and accounting firms; worked in Asia, Australia, Europe, and North America. In Higher Education, she has led post-experienced postgraduate and undergraduate programmes. Mona's research and practical work have received tremendous commendations. She has been invited to deliver seminars-keynotes by international think tanks (Asian Development Bank Institute), Government bodies, professional bodies, organisations, and universities in the UK and internationally.

About Henley Business School

Henley Business School is a triple-accredited business school and part of the University of Reading. With campuses, offices and partnerships around the world, over 7,000 students from more than 100 countries and over 87,000 alumni from 160 countries, we are a truly international institution. Our courses are enriched by up-to-date knowledge, research and commercial experience, and aimed at students and professionals at every stage of their career – from undergraduate through to postgraduate, PhD, MBA, DBA and executive education.


Henley
Business School
UNIVERSITY OF READING
Where business comes to life