

AI FOR BUSINESS

04

SUSTAINABILITY

AI has enormous potential to make firms more sustainable

07

LANGUAGE

How NLP and AI are being used to close the digital language divide

08

REGULATION

Brussels lays the groundwork for future AI regulation

DATA

Can synthetic data overcome the challenges of AI bias?

As data-hungry machine learning models demand increasing amounts of information, the market for synthetic data continues to grow. But is it as good as the real deal?

Sam Forsdick

One billion photos were used to train Meta's latest photo-recognition algorithm, a powerful demonstration of the current appetite for data. For those companies without access to platforms like Instagram, there is another answer: synthetic data.

Synthetic data is artificially created by a computer, rather than collected from the real world. These computer-generated images can be automatically annotated by the machine that creates them. Annotation is an important part of AI training and is a process where important points in a photo, such as people or objects, are labelled to help the machine learning models understand what the image depicts. They also avoid any compliance or privacy-related issues by virtue of being an original picture that doesn't feature real people.

Such technology spares companies the challenge of sourcing and collecting thousands of real-world images, while avoiding issues around privacy, GDPR and copyright.

"AI's biggest bottleneck is the scarcity of privacy-compliant, real-world data," says Steve Harris, CEO of UK-based synthetic data startup Mindtech Global. "Even a simple image recognition application needs up to 100,000 training images, and each image needs to be privacy-compliant and perfectly annotated by a human." Sourcing, annotating and cleaning of real-world data is "a monumental task", he says, which can occupy up to 80% of a data scientist's time.

Marek Rei is a machine learning professor at Imperial College London. "Collecting manual data is time-consuming and expensive," he says. "If you're able to generate data from scratch, you can essentially create endless amounts of it. For some rare events, obtaining even 10 real examples can be difficult, whereas synthetic data can potentially provide unlimited examples."

Thanks to these benefits, 60% of the data used for the development of AI and analytics projects will be synthetically generated by 2024, Gartner predicts, leading the consulting firm to describe it as "the future of AI".

With previous AI models, the development process involved collecting the data, training the model, testing it and making any necessary changes before testing it again.

The issue with this method is that the data used stays the same, according to Ofir

Any biases that are present in the data generation process – whether intentionally or unintentionally – will be picked up by models trained on it

Chakon, who is a co-founder and CEO of synthetic data company Datagen.


"The increase in performance that you get from this model-centric approach is relatively low," he says. "To get a significant improvement on the performance of your AI algorithms, you need to change your mindset. Instead of iterating on the model's parameters, you need to iterate on the data itself."

Datagen produces synthetic data for a range of AI applications, from facial recog-

GROWTH IN SYNTHETIC DATA

Synthetic data as a share of total data used in the development of AI projects

Gartner, 2021



niton technology to driver monitoring systems, security cameras and even gesture recognition. Chakon believes such applications will be increasingly popular as more companies expand into the metaverse.

To produce the computer-generated data for a facial recognition system, Datagen scans the faces of real people from a range of ages and demographics. Based on this 3D information, its AI learns the composite parts of the human face so it can generate images of new people. "From scanning 100 base identities, we can create millions of new identities," Chakon says.

For example, with enough information, the generative model could be asked to create a face of a 30-year-old white male with brown hair; it will spit out a completely new image each time.

"Based on what it learns from the real-world scans and the conditions that are put in, it can generate a completely new identity that isn't at all related to what was in the original collection of faces," Chakon says.

Proponents of synthetic data say this can help reduce the bias that often infiltrates algorithms at the training stage. "Biased training data can result in technology solutions and products that reinforce and perpetuate real-world discrimination," Harris says. "For example, AI systems have on many occasions been found to be poor at recognising darker skin tones. This is because the AI in question has been trained on datasets lacking diversity."

With synthetic data, it is theoretically possible for AI developers to generate an endless number of faces of people of different ethnicities to train its models, meaning gaps in the AI's understanding are less likely than has been known to occur historically.

Harris claims that some of its customers use Mindtech's AI training platform Chameleon to generate diverse data from scratch, while others use it to address the lack of diversity in their existing real-world datasets.

"By using computers to train AIs, we're removing the biggest roadblock to progress: human bias," he says.

There are, inevitably, issues with using computer-generated images to train AI for real-world applications. Rei explains: "Synthetic data almost never gives the same results as a comparable amount of real data. We normally have to make some assumptions and simplifications in order to

model the data-generation process. Unfortunately, this also means losing a lot of the nuances and intricacies in the real data."

This is easy to identify from a cursory glance at some of the faces that have been synthetically generated – they're unlikely to fool a person into thinking they're real. Datagen is currently investing in its photorealism capabilities, but Chakon argues that realism isn't crucial for every application.

"If you are developing a blemish detection AI for makeup application, having the detail is important," he says. "But if you're developing a security system, it's much less relevant whether you can identify small details on a person's face."

Synthetic data also isn't a silver bullet for AI bias; it relies on the people generating the data to use such platforms responsibly. Rei adds: "Any biases that are present in the data-generation process – whether intentionally or unintentionally – will be picked up by models trained on it."

An Arizona State University study showed that when trained on predominantly white, male images of engineering professors, its generative model amplified the biases in the dataset, meaning that it produced images of minority modes less frequently. Even worse, the AI began "lightening the skin colour of non-white faces and transforming female facial features to be masculine" when generating new faces.

With synthetic data programmes giving developers access to unlimited amounts of data, this has the potential to drastically exacerbate the issue of bias if errors are made at any point in the generation process.

If used correctly, synthetic data may still help to improve the diversity of some datasets. "If the data distribution is very unnatural – for example, it doesn't contain any examples of people from a particular race – then synthetically creating these examples and adding them to the data can be better than doing nothing," Rei says,

adding: "But it will likely not be as good as collecting real data with a more accurate coverage of all races."


While synthetic data can make the process of creating AI models quicker, cheaper and easier for programmers, it still comes with many of the same challenges as its real-world counterpart. "Whether synthetic is better than real-world data is not really the right question," Harris advises. "What AI developers need to do is find or create adequate amounts of appropriate data to train their system."

Using a combination of real and artificial data may be the solution. ●




of tech policy professionals in the US and Europe named AI algorithmic bias and transparency as a priority for new technology regulation

Clifford Chance, Milltown Partners, 2021



of AI projects through 2022 will deliver erroneous results due to bias in data, algorithms or teams managing them

Gartner, 2018



Distributed in

Published in association with

Contributors

Jon Axworthy

Journalist, specialising in health, tech, science and the future, with work published in *T3*, *Wearable* and *The Ambient*.

Sam Forsdick

Raconteur's staff writer, with particular interest in technology and the future of work. He has previously written for *I-CIO*, *NS Business*, *Press Gazette* and *New Statesman*.

Memuna Konteh

A London-based Sierra Leonean multidisciplinary writer with special interest in the intersections of identity culture and politics.

Oliver Pickup

Multi-award-winning journalist specialising in business, technology, sport and culture.

Daniel Thomas

Writer and editor, with work published in *The Telegraph*, *Newsweek*, *Fund Strategy* and *EducationInvestor*, among other publications.

raconteur reports

Publishing manager

Chloe Johnston

Managing editor

Sarah Vizard

Deputy editor

Francesca Cassidy

Reports editor

Ian Deering

Sub-editors

Neil Cole
Gerrard Cowan
Christina Ryder

Head of production

Justyna O'Connell

Design/production assistant

Louis Nasse

Design

Celina Lucey
Colm McDermott
Samuele Motta
Sean Wyatt-Livesley

Illustration

Kellie Jerrard

Design director

Tim Whitlock

@raconteur

raconteur-media

@raconteur.stories

raconteur.net

/ai-business-2022

THE AI DATA COMPANY

Business Success in a Data Driven World

ddn.com/ai-for-business





How to successfully build AI into your business

AI can help organisations unlock innovation and gain a competitive advantage, but delivering it at scale can be a challenge

There is no doubt that the potential of AI is exciting. Whether it's driverless cars or faster drug discovery, AI is changing our world and how we live in it. In the corporate world, financial analysts are able to use AI to improve the quality of their market analysis. Using AI, organisations can pull together and analyse millions of pieces of data from company tweets to regulatory filings and even media appearances, to identify patterns of language and behaviour. In other words, AI can tell you whether a company's good news is really good news. Still, despite the exciting and futuristic applications of AI technology, many early adopters have struggled to deliver any ROI or business value outside of small, experimental projects. Gartner reports that, in 2022, 48% of CIOs have already deployed or plan to deploy AI and machine learning technologies

in the next 12 months. However, they also report that many of these organisations won't be reliably using AI in production until 2025. The problem isn't usually with AI itself, but rather with delivering AI at scale, says James Coomer, vice-president for product management at DDN, which specialises in storage systems for AI. Having solved the scalability challenge for some of the largest global AI projects, DDN now helps organisations plan their AI journey to production. These challenges are more often linked to the large-scale implementation and deployment of AI, as organisations look to deliver greater value and innovation to their customers and shareholders. For example, some use cases like online retail and consumer preference analysis can work well in modest deployments, using datasets that are not too large in size. On the other hand, natural language processing

and real-time video analysis need datasets that are well beyond the experience of traditional IT teams. "Truly delivering the benefits of AI means dealing with volumes of data that are orders of magnitude higher than anything we've seen before," says Coomer. "Not only is the volume of data higher, but much of this data is rich, unstructured information like video, audio or images." For some organisations, the ability to gather and process very large volumes of data can unlock the next level of innovation. For example, DDN recently worked with Recursion Pharmaceuticals to use AI to accelerate the process of drug discovery. "Using AI rather than conventional methods can massively speed up our ability to learn and develop new vaccines or medications," says Coomer.

Truly delivering the benefits of AI means dealing with volumes of data that are orders of magnitude higher than anything we've seen before

with DataOps, ModelOps and DevOps to create a consistent approach to accelerating AI development.

An AI factory framework might define agreed processes and approaches to things like how security is applied to data, how the data platform is managed, or how silos are reduced to improve overall performance. "It's a reference architecture that reduces the risk and work involved in building AI at scale," explains Coomer.

By adopting standardised processes and governance, the risks associated with AI are massively reduced. For organisations that are new to AI, Coomer suggests thinking about AI projects in terms of four essential requirements:

- First, you need to be able to identify what data is needed; what is the right data source and how will you capture and organise that data?

- Second, do you have appropriate governance in place? If you're building an AI model that will make recommendations about insurance cover, how do you prove that the model is trained and using data reliably?
- Third, it's important to have the right people with the right tools. To build knowledge, make it easier for data scientists and data teams to share knowledge and collaborate.
- Finally, ensure you have the right platform to process data efficiently. Many projects fail because of data bottlenecks or because your network and devices cannot process data at scale.

According to Coomer, many AI systems drown in data as soon as they move out into the real world because traditional approaches can't manage the huge volumes needed to feed large-scale AI models. "AI systems need to ingest data at incredibly high rates, and you need to make that data accessible and secure to enable the kind of deep learning needed to make critical decisions." When it comes to the fourth requirement, Coomer advises companies to think carefully about the capacity of their storage systems. "In our experience, people think about AI projects in terms of GPU and computing requirements, but don't consider the data ingestion, and the volume of data that will be involved once an AI model goes into production," says Coomer. "You could have the best data scientists and models, but if your storage can't ingest the 250 petabytes of data that's supporting your AI model, then your project will fail."

Businesses should prepare for much larger data requirements by adopting AI-optimised systems and processes. These are systems developed with AI in mind, so that adjustments are made at every stage of the data journey to improve the speed and efficiency of data ingestion and transfer. In its State of AI in the Enterprise report, Deloitte found that the top priority for IT leaders investing in AI was modernising the data infrastructure for AI.

These days, AI does not necessarily need to be the complex, painful learning curve that early adopters have travelled, says Coomer. Companies embarking on AI projects today have access to an increasingly large number of AI-optimised tools and platforms that make it easier and faster to deploy AI at scale. Deloitte found that 80% of technology leaders are buying at least as much AI technology as they build, and the longer a company has used AI, the more likely they are to buy commercial platforms and tools.

To achieve quick wins from AI, organisations must invest in all four pillars of AI – data, people, process and platform. "They're all important, and what we focus on at DDN is solving the platform part of the issue, whether you need a twentieth of a single rack or something that can handle petabytes of data without increasing the complexity," says Coomer. "The key thing to remember is not to fall into the trap of doing what you've always done, and rather start with something that's built to solve the problem in front of you."

To find out more, visit ddn.com/ai-for-business



Q&A Unlocking the promise of AI

Alex Bouzari, CEO of DDN, talks about how enterprises can get their AI projects right

Q Who is DDN, and how do you help organisations on their AI journey?
A DDN is a market-leading global organisation which provides powerful data storage solutions, and facilitates the AI- and data-enabled digital transformation journey for enterprises and research facilities. With over 10,000 customers and thousands of deployments worldwide, DDN is making it easier, safer and faster to implement AI and reap massive rewards in healthcare, autonomous driving, finance, retail and many other sectors. We de-risk the customer journey to a successful AI transformation: AI systems need enormous amounts of data to unlock meaningful business value: at DDN, we accelerate their digital transformation journey, by delivering

the fastest path to implementation of AI, helping them move seamlessly from prototype to production. This means that organisations can focus on business value, instead of managing technical complexity and having to make compromises on data and scale. **Q Many business leaders see AI as a daunting, complicated undertaking. Is that a fair assessment, and how are things changing?**
A I think when any new technology comes along that has the potential to transform how we do something, people get baffled. The market is maturing so rapidly and there are now so many more organisations who can help you to implement AI strategies without delay or complexity.

The important thing is to interview and select vendors and partners carefully. Have they done this before? Do they know your market? If you get the right people and they have the right experience, you will get success. **Q What are the key steps a CIO must start with to unlock business value and benefits from AI?**
A I firmly believe to succeed with AI you need the right infrastructure, and that means having platforms that can deal with the massive amount of data and then help you to find the virtual needle in that haystack. People wasted a lot of time trying to get their existing infrastructure to support AI, but you need new data repositories that are

optimised for AI, and they need to be integrated with the business. That integration process has matured significantly in recent years, and there are now many more bridges between AI platforms. **Q What's the most exciting application of AI you've seen recently?**
A We recently helped build a number of AI systems to support critical pharmaceutical and life sciences research. This is an exciting growth area, where companies are using AI supercomputers to discover new types of therapies and to improve the design of antibodies to help treat inflammatory diseases such as asthma and COPD. Digital biology is attracting both established organisations and emerging


startups because it changes the economics of drug discovery. **Q What would you say is the biggest mistake companies are making with AI right now?**
A I think it's trying to use existing technology for new purposes. A lot of technology exists with the purpose of keeping the lights on, managing the database, and taking care of operations. It's not often seen as an enabler. Now things have changed, because AI is a business enabler, and the volume and type of data required to support that enablement is completely different to what we've seen before. It's not that it's much harder, it just needs different tools and platforms. Don't take my word

for it. Do the research. Go and see the successful implementations, see who has delivered, and ask them how they achieved those better outcomes. **Q Over the next five years, what's the biggest change we'll see in how we approach AI projects?**
A I think within five years we will start to see customised implementation of AI tools and platforms because organisations will want AI that meets their specific needs. Beyond that, I'm excited to see AI used in intelligent homes. I don't just mean automated smart lights or music, but things like medicine built into intelligent dwellings to enable us to live healthier and more active lives.


HOW FINANCE USES AI

As AI grows in sophistication, it has become increasingly attractive to the financial services industry. In an industry that relies on accurate data analysis, financial services firms are using AI to assess risks and automate processes. Here, we highlight some of the most common use cases for AI in financial services and explore how firms are using AI to combat money laundering and fraud

AI ADOPTION IN THE FINANCIAL SERVICES INDUSTRY WORLDWIDE, BY FUNCTION




AI CAPABILITIES EMBEDDED IN STANDARD BUSINESS PROCESSES IN THE FINANCIAL SERVICES INDUSTRY WORLDWIDE



FRAUD DETECTION IS THE NUMBER-ONE USE CASE FOR FINANCIAL SERVICES

AI use cases in the financial services industry worldwide



Up to \$2tn is laundered each year globally
SAS, 2021

£159.7m was lost to online banking fraud in the UK in 2020
UK Finance, 2021

56%

of managers with compliance responsibilities, globally, said that AI and ML tools significantly help with financial crime prevention

Refinitiv, 2019


USES FOR AI IN AML INVESTIGATIONS

Top three AML investigation processes that financial services professionals think could be improved with AI



IBM, 2019

SHARE OF COMPANIES WORLDWIDE USING AI AND ML TO DETECT FRAUD



Algorithmia, 2021

SUSTAINABILITY

Machine learning vs climate change

AI is increasingly being recognised as an important tool for businesses in the pursuit of their net-zero targets – although its own energy requirements are not inconsiderable

Jon Axworthy

There has been a growing realisation among businesses in recent years that becoming environmentally sustainable is a must, not a choice. Customers, investors and employees and industry regulators are all putting pressure on them to act before the climate crisis worsens to calamitous levels.

Alongside this, the willingness of companies to publicise their progress in reducing their ecological impact is increasing. More than 3,400 organisations, with a combined market cap of £21.4tn, have registered their support for the Task Force on Climate-Related Financial Disclosures since it published its first reporting recommendations in 2017, for instance.

AI has a key role in helping firms to hit the ambitious net-zero CO₂ emissions targets they are setting themselves. The *Global AI Adoption Index 2022*, IBM's latest annual survey of uptake, found that two-


“We’ve identified a multitude of practical applications and decision points where AI can make a meaningful impact

thirds of the 7,500 IT chiefs it polled were either using AI to achieve sustainability goals or planning to do so.

Among other things, they’re looking to the technology to automate data collection and help them provide verifiable information on their firms’ environmental performance. But, according to the research, by far the most popular use for AI in this field is to make operations more efficient and, by extension, more environmentally friendly.

“When it comes to the successful adoption of AI, it’s crucial for companies to identify a clear challenge that it’s well suited to address, then focus on implementing the tech in a way that fits with their work streams,” says Dr Kareem Yusuf, general manager of IBM Sustainability Software and leader of its AI applications business. “For example, you could use AI to capture a product’s carbon footprint data across complex supply chains more easily and feed that into your sourcing and procurement decisions.”

It was a supply chain problem that first prompted Moosejaw, a Canadian retailer of outdoor recreation clothing and equipment, to turn to AI for a solution. Given the nature of its products, the firm was particularly interested in showing that it was constantly reviewing its environmental performance.



the *Enterprise 2022* research report has highlighted a potentially problematic lack of governance. The study revealed that more than half of the organisations it polled that were using AI didn’t have a governance plan in place. It found that the biggest problem was a lack of in-house expertise to get at the necessary data.

The race for the data to help organisations reduce their emissions and strive for carbon neutrality also appears to be fuelling an unprecedented acceleration in the uptake of AI. Tortoise Media’s *Global AI Index* research has tracked a significant upsurge in adoption over the past two years. Half of the firms in its most recent survey revealed that they had recently stepped up their AI usage. “We’ve reached a point in AI maturity where we’ve identified a multitude of practical applications and decision points where AI can make a meaningful impact,” Yusuf says. “For example, companies can look to automatically link their energy usage with their physical asset systems to identify predictive maintenance opportunities to improve their environmental performance. Or, when it comes to the operation of data centres, firms could use such insights to optimise workloads and schedules based on their energy consumption.”

The operation of data centres highlights a dichotomy at the heart of the relationship between AI and sustainability. For all the enthusiasm that firms have shown for the technology’s environmental applications, this must be tempered by the knowledge that using AI itself consumes significant amounts of energy.

Nowhere is that more evident than in the humming global data centres of Google, a company that’s responsible for a significant proportion of the world’s electricity output, particularly the energy needed to keep its servers cool and thereby reduce the risk of a devastating outage.

Since September 2020, Google has been working to achieve its pledge of running entirely on carbon-free energy by 2030. It’s been applying AI to this challenge, using it to predict the combined effects of various procedures in its data centres on energy consumption and identifying where efficiencies can be achieved.

Professor Ong Yew Soon, chief AI scientist at Singapore’s Agency for Science, Technology and Research, believes that there is further potential to cut the energy consumption of AI systems.

“Usage can be saved at the ‘training phase’, so we are looking to reduce the need for massive amounts of training data,” he says. “Instead of training an AI system on a massive web of information, researchers can pick and choose according to the system’s application, which would save tonnes of CO₂.”

It’s just another example of the link between AI and sustainability. If the technology can find a way to minimise its own carbon footprint, it will close a virtuous circle, helping its users to hit the net-zero targets that matter so much to all stakeholders in the climate crisis. ●

Moosejaw identified that many of its customers were practising a common online behaviour known as size sampling. This occurs when a consumer orders the same garment in more than one size – with the intention of trying them all on at home and sending back those that don’t fit. The task of dealing with these unwanted items when they enter the reverse supply chain generates a lot of unnecessary greenhouse gas emissions. Research published by tech firm Optoro has projected that handling returns – in

the US alone – could be responsible for putting 23 million tonnes of CO₂ into the atmosphere in 2025.

With the help of AI, Moosejaw identified that almost 15% of online purchases returned by its customers could be attributed to size sampling. It set about minimising this behaviour by using a data-driven personalisation platform called True Fit.

Sarah Curran-Usher, managing director of True Fit in EMEA, explains how the system works: “When a shopper places multiple sizes of the same item into their shopping


cart, a change in the user experience prompts them to create a True Fit profile. The platform can then pair that consumer’s data with data in its fashion genome to recommend the best fit. Using this AI, Moosejaw has been able to reduce the rate of size sampling by 24% in one year.”

Whether it’s deployed at the customer interface or on a purely operational level, AI can help firms to extract precious gems of insight from the mountain of data they’re sitting on. But there are barriers. For instance, O’Reilly Media’s *AI Adoption in*

IS AI THE KEY TO GOING GREEN?

Capgemini, 2020

AI’s potential contribution to reduction of economic emissions intensity (volume of emissions per unit of GDP) by 2030, by sector



QuantSpark

Are you ready for AI or are you a data ad-hocracy?

Yesterday’s analytics meant manually refreshing an Excel spreadsheet.

At QuantSpark we partner with our clients in Private Equity and Asset Management to support their journey from Excel to AI. Our approach to transformation through analytics and AI starts with business process analysis and ends with the development of bespoke analytics solutions:

1

Strategy & Process Analysis

2

Data Engineering

3

Advanced Analytics

4

AI & ML Capabilities

SCAN ME

Take the first steps on your AI and analytics journey: www.quantspark.com/financial-services

COMMUNICATION

Sales talk: how conversational AI can win over customers

Conversational AI can mimic human interactions. With today’s consumers turned off by the hard sell, the technology holds strong potential for businesses

Oliver Pickup

When it comes to sales, businesses should reverse Elvis’s famous advice: a little more conversation and a little less action, please.

The secret to success with today’s consumers revolves around small talk and adopting a long-term approach. Direct approaches – seeking to add notches to the sales equivalent of a bedpost – are a huge turn-off for customers.

Happily, so-called conversational AI is now mature enough to assist adroitly with the more mundane topics, enabling humans to enter the chat room later, at the most appropriate point.

Conversational AI refers to tech solutions such as chatbots or virtual agents that use vast volumes of data, machine learning and natural language processing to imitate human interactions. Businesses today must adopt the technology as a matter of urgency, with laggards likely to lose out.

Over 70% of customers expect conversational service, meaning human-like interactions – complete with emojis, gifs, images and videos – whenever they engage with a brand, according to Zendesk. But only 40% of businesses can deliver this successfully.

Little wonder the global software-as-a-service company recently announced new capabilities for its Sunshine Platform, a customer relationship management service, including conversational automation via bot technology. The upgrade enables organisations to expand automation to messaging apps such as Facebook Messenger and WhatsApp and allows them to build and train custom bots to address common issues.

“A quick conversation can resolve most things in life,” says Matthias Goehler, Zendesk’s chief technology officer in EMEA. “Embracing advances in AI to deliver conversational exchanges with customers easily is a natural direction for customer experience teams to take.”

One of the most significant benefits of conversational AI is that all customer communications are retained, Goehler adds. This means a more complete picture is achieved, allowing businesses to understand people’s personal preferences better and enrich their experience. It facilitates a personalised, data-driven service, removing some of the burdens on human agents and empowering them to do more for the customer in less time, he says.

“A conversational approach makes interactions more informed – built with the context of the customer’s history. When done right, it can even help increase a customer’s spending with you by making useful and simple recommendations to purchase from within a chat.” Katie King is the author of two books about AI for sales and marketing and a



member of the government’s All-Party Parliamentary Group Taskforce for the enterprise adoption of AI. Companies that embrace conversational AI will charm employees and customers alike, she says.

“Many of the queries that cross the service agent’s desk are frequently asked questions with simple answers,” she says. “While these queries might be easy to answer, they still take up valuable time and limit the agent’s capabilities to handle some of the more complex issues. It’s overwhelming and leads to faster employee burnout and potential staffing issues for the company.”

Conversational AI can help tackle this challenge, so appeals to many organisations, King notes. “AI can cut out that first

step of the process by engaging the customer and potentially resolving their issue without human intervention,” she says. “Additionally, technology doesn’t have working hours like a human employee does, meaning time customers can gain the help they desire on their terms.”

With the surge in energy prices, concerned customers of E.ON – the largest energy and renewable electricity supplier in the UK – have certainly wanted help. Conversational AI is easing the load.

Nikolai Berenbrock is the company’s head of conversational experiences. He says the company currently has more than 50 conversational AI solutions across the group, serving customers and employees and covering about 30% of demand. “This has enabled us to offer a better customer service experience and a massive reduction in our operational costs,” Berenbrock says.

E.ON uses AI to automate repetitive tasks so that agents are “available to jump in where they can make a valuable difference”, he adds. The technology “allows us to scale our customer service in a location and time-independent way, so that we can be where our customers are by offering our service on our website in a LiveChat channel, WhatsApp, Facebook Messenger, telephony channel, etc, whenever they need us, 24/7.”

Jason Costain is head of fraud prevention at NatWest, which serves 19 million customers across 12 banking and financial services brands. He offers another example of how conversational AI is being utilised.

“Using voice-biometric technology, we’re building a clear picture of our customers’ voices and what criminal voices sound like,” he says. “We can detect when we get a fraudulent voice coming in across our network as soon as it happens. Using a combination of biometric and behavioural data, we now have far greater confidence that we are speaking to our genuine customers and keeping them safe.”

Demand for conversational AI isn’t limited to customer experience, says Goehler. “We’re seeing huge demand from companies using our solutions for employee experience, with tickets filled by corporate employees jumping 31% last year – nearly double the rates seen by customer-facing support teams at B2B and B2C companies,” he says, signposting the direction of travel.

Despite the clear advantages of conversational AI and the momentum behind the technology, Goehler sounds a note of caution to business leaders who, to quote another Elvis song, can’t help falling in love with the technology. “While just over half of EMEA companies report that chatbots are becoming more human-like, AI can’t – and shouldn’t – be a 100% solution,” he says.

Zendesk research indicates more than 60% of customers will walk away after one poor experience – up 22% from last year. Perhaps this shouldn’t be a surprise. After all, who can blame their suspicious minds? ●

Q&A

“I didn’t want AI leaving people behind, so I decided to do something about it”

A celebrated tech leader and design pioneer, **Robb Wilson** built a conversational AI platform that is sweeping the marketplace. Robb has poured decades of learning into his new book *Age of Invisible Machines: A Practical Guide to Growing a Hyperautomated Ecosystem of Intelligent Digital Workers*



Q Your career began in the film industry. Did that influence your later work?

A Definitely. In the early '90s, I worked in sound design at Warner Bros., where many of my colleagues had been in the film industry for nearly half a century. At the time, the film world was transitioning from analogue to digital. I realised the wisdom these sound editors brought to their craft was in danger of being lost if the industry started hiring people who were simply better at using computers.

There’s a very distinct rhythm to editing film, and a lot of that comes from the physical act of cutting and splicing celluloid, so I designed a digital editing interface that mimicked the controls of an old school sound editing system. The older, experienced editors were able to acclimatise to the digital editing process and younger editors learned the rhythms of working film by hand. The experience showed me how important it is to bridge the gap between humans and machines.

Q From film you moved into experience design, enjoying considerable success. How did that shape your perspective on conversational AI?

A I don’t like waving my accomplishments around, but it’s safe to say that almost everyone on the planet has touched a piece of technology I had a hand in developing. I was tapped to design one of the first iPad apps and, over the years, I’ve worked with many of the largest technology companies in the world.

I’ve also collected more than 100 awards, which I only bring up because roughly half of them are in technology and the other half are in design. This balance is important, because while conversational AI involves the orchestration of a whole host of complex

technologies – like NLU/NLP, code-free programming, RPA, and machine learning – its success is entirely dependent on adoption, which can only be achieved by focusing on experience design.

Q Is that balance, between technology and design, central to your work at OneReach.ai?

A Definitely. I’ve been exploring conversational AI for more than two decades, and what first drew me in was the fact that, of all the experiences people had with technology, conversational ones were uniformly the worst. Once conversational AI experiences stop sucking, adoption will skyrocket. Conversation is an interface that requires zero training, which gives anyone, anywhere, the ability to leverage powerful problem-solving machines. Basically, I didn’t want to see technology – specifically AI – leave people behind, so I decided to do something about it.

Our platform allows organisations to take charge of their automation strategy and put themselves in the position to become self-driving. Our approach works, and OneReach.ai scored highest overall in the inaugural 2022 Gartner Critical Capabilities for Enterprise Conversational AI Platforms report. For me and my team, it validated our philosophy that the only way to really succeed with conversational AI is to have the ability to orchestrate all the associated technologies using code-free tools on an open platform; one that’s flexible enough to incorporate the best products the marketplace has to offer.

Q Are there other things you’re doing differently at OneReach.ai we should know about?

A Our no-code creation tools make it so that the people building

“OneReach.ai scored highest overall in the inaugural 2022 Gartner Critical Capabilities for Enterprise Conversational AI Platforms report

solutions don’t need any of the skills typically associated with a software developer. This has completely changed the way we hire solution designers. We aren’t looking for unicorns who can write code while understanding the nuance and intricacies of design. When building with our platform, the software is, in effect, the direct result of the design process. This frees us up to hire people who are passionate about solving problems.

One of our most successful lead solution designers came to us right out of college with a biology degree. It made no difference to us that she didn’t have a background in technology or design. We saw her as a crack problem solver, and in a few short years on our platform she’s probably designed more unique conversational experiences than entire teams at enterprise-level companies. And I don’t think it’s a coincidence that more than half of our designers are women. I was raised by strong women who got stuff done – they were outcome-driven, just like the designers on our team.

Q How did growing up under the tutelage of Canadian philosopher Marshall McLuhan inspire you?

A It was interesting because he was a friend of my mother, and he lived down the street so was more like my lovable, grumpy uncle. I didn’t realise how renowned he was outside of my block until years later. He instilled in me a deep reverence for the power of technology. One thing he said that I still think about all the time is this: “We shape our tools and thereafter our tools shape us.” I see evidence of this everywhere.

When engineers created algorithms with the positive outcome of getting the peak number of eyeballs on social media posts, they probably didn’t realise how evolved versions would imperil the mental health of an entire generation – or that news feed algorithms would undermine the mechanics of democracy.

We are at a critical impasse right now. Conversational AI can quite literally help humans fix the world, but only if it’s used in thoughtful and responsible ways. This will be a tall order as companies the world over scramble to keep up, but I have hope we can make technology work for humanity in profound and positive ways.

Download a free chapter of Robb’s new book at invisiblemachines.ai

OneReach.ai

Can emotion AI really boost sales performance?

New AI technology could help sellers analyse customers' emotions. But how will people react if their facial expressions and body language are tracked?

Sam Forsdick

The sales rep might not be the only one watching you on your next virtual meeting. Some companies are employing 'emotion AI' - a subset of AI that can detect human emotions - to monitor individuals on sales calls, provide feedback on their reactions and highlight the most engaging parts of the pitch.

The market for such sales enablement platforms is growing, with consultancy firm Verified Market Research claiming it could be worth \$7.3bn (£5.82bn) globally by 2028.

Zoom is one of the latest entrants to this burgeoning market, launching Zoom IQ for Sales last month. Described as a conversational intelligence software, it claims to offer "meaningful and actionable insights from customer interactions to improve seller performance and enhance customer experiences".

But the announcement met with opposition from human rights groups and privacy campaigners. An open letter to Zoom CEO Eric Yuan, co-signed by the American Civil Liberties Union and the Electronic Privacy Information Center, called for the video communications company to halt its plans to advance the feature, adding: "This move to mine users for emotional data points based on the false idea that AI can track and analyse human emotions is a violation of privacy and human rights." Zoom declined to be interviewed for the piece.

According to Uniphore, one of the companies behind the technology, its emotion

AI assistant, Q for Sales, can "help sellers 'read the room', sense emotional cues, and improve engagement".

Patrick Ehlen is vice-president of AI at the company. He thinks the fears over emotion AI come from a misunderstanding of the technology.

"I'm not a huge fan of the term 'emotion AI'," he says. "There seems to be a concern that when people talk about emotion AI, they're talking about a computer system

size of the emotion recognition and detection market worldwide

Markets and Markets, 2022

that can read your internal emotional state with a high degree of accuracy. The truth of the matter is there's no AI system that can do that."

Psychological research backs up this assertion. The way that people communicate emotion is not universal and facial expressions are much more contextually



sensitive than previously thought, according to a 2019 paper from psychology professors at the California Institute of Technology, Northeastern University, University of Glasgow and University of Wisconsin-Madison.

Lisa Barrett, professor of psychology at Northeastern University and one of the paper's authors, believes the findings "call into question the scientific justification used by some emotion AI technologies".

Ehlen is therefore reluctant to make any bold claims about the technology's capabilities. He maintains that "it's an unreliable technology to be making very big, critical decisions". However, he thinks it could prove useful on sales calls. In situations where a salesperson may be speaking to multiple people on a video it can be hard to determine who is engaging and who isn't.

"You're looking at people in this tiny little window, so it's not as easy to read the room and see what people's facial expressions are," he says. "If you can have a machine gauging their reactions and determining at what point the CPO seems interested, it can be a helpful tool."

To do this, the visual AI software analyses a number of elements of the conversation, including facial expressions, tone of voice and people's gestures. "From that information, we're able to get much closer to having a 360-degree view of what people are doing when they're in conversation to better understand them," Ehlen adds.

The platform can then give salespeople real-time feedback on sentiment and engagement to help them adapt their responses and, in theory, improve their sales conversions.

Sybill is another platform which claims to use emotional intelligence to accelerate the sales process. Co-founder and CEO Gorish Aggarwal is reluctant, however, to say the programme can identify people's emotional state. "We consider it as a behaviour AI, which is different from emotion because emotions are subjective," he says. "You cannot tell whether a person is fearful, angry or contemptuous just by looking at their face."

The software instead looks to identify an individual's body language or facial expression to highlight key moments, Aggarwal says. For example, whether someone is nodding along to a conversation or smiling. Salespeople can then revisit the recording to see when people were most engaged and which parts they should follow up.

Although Aggarwal claims its AI can determine whether someone is nodding their head with up to 95% precision, the technology's impact in the sales context is yet to be proven. Both Uniphore and Sybill's platforms launched earlier this year and are both currently conducting valuation studies to determine to what extent their AI programmes can improve sales performance.

Jason Bell is associate professor of marketing at Said Business School and works on

You cannot tell whether a person is fearful, angry or contemptuous just by looking at their face

AI models for computer vision and natural language processing. He believes there is a lot of 'overclaim' in the emotion AI market. "Many webcams and front-facing smartphone cameras are low quality, so even if you have a great predictive model, the signal that it's getting is not great," he says. "In principle it's possible, but I haven't been convinced by the current technology."

Relying on this technology in a sales context could lead people down dead ends, according to Bell. His biggest concern centres on the technology's accuracy. "It dramatically oversimplifies emotional states," he adds. "If you're monitoring one or both parties involved in the sales process and categorising their emotions, it can be quite crude ... You could create more complications for the sales process without adding a tonne of value."

Concerns around privacy, as raised in the open letter to Zoom, have also made

people wary about implementing the technology. Uniphore and Sybill are both highly conscious of this concern.

"We're very aware that some people see AI as being somewhat creepy," Ehlen says. "So, we've tried to build in as many protections and safeguards as possible to make people feel comfortable."

As a result, Q for Sales is an opt-in experience. With Sybill, users are notified that "the call is being recorded and notes are being taken". Although there is no reference to AI unless the user decides to reveal this.

Aggarwal doesn't recommend using the programme for internal calls or in situations where there is a power imbalance. For example, using it as part of an interview process would be "ethically not correct as they wouldn't be able to refuse".

There is also a risk that people change their behaviour when they're made aware that they are being recorded. Ehlen admits that this could mean "the AI is not going to work as well as it did before" and will have to go through a process of retraining.

"It's an iterative process," he says. "As these technologies make their way into the mainstream, people will become more comfortable with them and this will become less of an issue."

Whether these conversational AI programmes become a standard part of the sales process remains to be seen, but convincing customers to be monitored by AI might be the hardest sell. ●



LANGUAGE

Joining the global conversation

To avoid indigenous populations dropping out of the world's conversation, steps are being taken to narrow the digital divide – using current tech and new methods

Memuna Konteh

It's easy to think that anyone with an internet connection can access the information, networking platforms and convenience-enhancing tech found online. Unfortunately, this is not the case. Of the world's 7.151 languages, more than half don't have any digital footprint. This phenomenon is described as the digital language divide and unless it is addressed, could accelerate the extinction of thousands of languages.

World language resource Ethnologue reported that about 40% of languages are now endangered, often with fewer than 1,000 speakers remaining. Meanwhile, just 23 languages account for more than half the world's population. If a language falls out of use, so does knowledge of the rich cultures and histories they describe.

Tech startup Derivation aims to address the problem, using language insight solutions and pioneering AI. Stephen Jones, acting CEO at Derivation, explains: "Our solutions include measuring every living language to assess their progress towards increased and sustainable relevance in the digital age. Linguistics informs us that languages thrive on a 'use it or lose it' basis – if no one is writing, reading, signing or speaking a language, then over generations its usage declines and eventually it falls out of use. As the world relies more on technology, that use-or-lose equation is shifting into the digital arena."

A lack of digital infrastructure (such as keyboards, fonts, operating systems) supporting a language usually indicates real-world vulnerabilities are affecting a language and its native speakers. Among the most vulnerable both online and in the real world are non-written languages – those used in areas that have low literacy rates, insufficient internet or mobile phone coverage or areas with slow economic growth. Derivation research suggests that languages which are used in an official capacity (such as in government, legal, economic, health and education systems) are more likely to be digitally supported than non-institutional languages.

Jones notes: "The digital language divide is felt less in Europe, where there are fewer regional-specific languages, and seen more in Africa and Asia, where there's generally a higher concentration of indigenous languages. "If you speak an 'official language', particularly one in Europe, you are much more likely to find digital options available in your mother tongue. But if your first language is an indigenous one, and particularly if it is one that doesn't have official recognition in Asia and Africa, you're more likely to face digital language exclusion and forced to operate in a different language."

The exclusion of these languages from the digital sphere is both a humanitarian

and an economic issue, as native speakers of vulnerable languages have little or no access to educational materials and social interaction, while businesses lose out on marketing opportunities and valuable data from within these communities.

In his co-authored research paper, Sustaining Language Use, linguist Dr Gary Simons says the preservation of indigenous language is also a matter of public health and human rights. "Members of a community which is in the process of losing its language and culture experience significant amounts of disruption and stress in all areas of life. A child growing up in a community which is viewed with disdain develops a self-image that reflects that experience."


"As social norms are abandoned, they frequently are not replaced all at once or with adequate equivalents. This can lead to social tensions, divisions in the community, disruptive and harmful patterns of behaviour, and even violence. There is evidence that communities experiencing such a transition may have elevated levels of alcoholism, drug addiction, HIV/Aids, and suicide."

He goes on to cite the United Nations Declaration on the Rights of Indigenous Peoples (Article 13.1), which states: "Indigenous peoples have the right to revitalise, use, develop and transmit to future generations their histories, languages, oral traditions, philosophies, writing systems and literatures, and to designate and retain their own names for communities, places and persons."

Aside from the human-centric arguments, Simons notes the importance of

CHALLENGES FOR VOICE TECHNOLOGY

Share of business professionals identifying the following as barriers to voice technology adoption, worldwide



indigenous language preservation in ecological sustainability. In that "many minority language communities possess highly developed bodies of knowledge about their physical environment and have elaborated technologies to adapt and make use of their environment that are transmitted through equally highly developed linguistic forms in their languages. With the loss of those languages and cultures entire areas of human knowledge are also at risk and at the same time biological diversity is also threatened."

Fortunately, the emerging field of language analytics – a combination of business intelligence, big data and linguistics – is used to inform developing AI and machine learning tools that can work to bridge the gap between established languages and those being left behind.


SIL is a non-profit organisation involved in 1,600 active language projects in 98 countries. SIL data scientist Daniel Whitenack says that although the rapid integration of AI into existing digital systems has "the potential to further marginalise already marginalised language communities, artificial intelligence creates amazing new possibilities for much of the world".

If you speak an 'official language', particularly one within Europe, you are much more likely to find digital options available in your mother tongue

Specifically, the advancement of natural language processing (NLP), which "has enabled researchers and engineers to decipher long-lost languages, translate speech in one language to speech in another language directly without converting to text, generate long-form text that adapts to the style and content of human prompts, and translate between language pairs never seen explicitly by computer systems". Because they relied heavily on "manually crafted linguistic resources", earlier NLP models largely ignored vulnerable indigenous languages. But this has changed, according to Whitenack: "Large tech companies, academic institutions, and grassroots organisations are beginning to realise that NLP methods could – and should – be extended to minority languages."

"For a technology like machine translation, this shift is evident from Facebook's recent request for proposals on neural machine translation for low resource languages. Google's focus on building translation systems for 'the next thousand languages', and the large, multi-institute Masakhane grassroots organisation's research efforts that focus on machine translation for African languages," he says.

The new attention placed on AI and NLP for indigenous languages will enable local community members, organisations, and non-profits to build momentum towards new, multilingual applications that give minority language speakers a chance to bring their voices into the global conversation and shifting economy. This will shift the attention, currently directed towards English and Mandarin, to a more democratised focus and participation across language communities. ●



WHAT ORGANISATION WITHIN YOUR COMPANY LEADS THE AI KPI PROGRAM?				
Chief data officer/data science organisation	CTO organisation	Other		
30.4%	30.4%	26.1%	8.7%	4.3%
	CIO/IT organisation		Product management	

ROI: how to define AI success

As AI adoption reaches critical mass, confidence around projects is rising and the pace of deployment is increasing. However, setting and measuring relevant targets to protect returns on investment remains a major challenge

It is perhaps the strongest indication of any technology's popularity when general hesitancy around its use begins to abate and projects proliferate. This is clearly the case with artificial intelligence, as nearly a fifth of companies now use AI and are pushing it into multiple business areas to drive change.

According to new surveys by the research, events and training company Informa Tech, business AI adoption is finally exiting the 'pioneer' stage and reaching an early maturity, a kind of 'adolescence' where the future is being established. Some 18% of businesses are now using AI across myriad business functions. "It is an explosive period for AI adoption, and in passing this tipping point we expect that the curve will soon rise rapidly, with a doubling or tripling of projects in just a few years," explains Mark Beccue, a principal analyst at the company.

The 'style and nature of businesses' AI deployment has changed significantly, he notes: "Adoption has moved from those early pioneers who progressed at a trial-and-error pace, to newcomers who have been able to build on those foundations and drive headlong into AI initiatives, armed with market-tested solutions and tools, as well as best practice playbooks."

Setting targets and leadership

Despite the growing momentum in projects, boardroom tensions around AI are rising as increasing expenses meet headlong with severe difficulties in quantifying the business benefits. Setting relevant key performance indicators (KPIs) has become a major sticking point here, with some 36% of businesses warning that the unclear goals for their initiatives are creating severe impediments to achieving success. Other businesses suffer from immature initiatives and a lack of proper responsibilities and leadership strategies, affecting their AI progress. Only 18% said they have proper goals in place for all their initiatives.

"After speaking to businesses working on important AI projects, it's evident that many are still struggling to set measurable goals for their efforts and are often working with unclear, ad hoc KPIs that shift around constantly as projects progress," notes Beccue. "With the growing examination of AI's return on investment, it is important that companies work methodically towards tying their deployments to clear, consistent, specific business impacts."

Among the challenges for organisations is to simultaneously draw out both the granular impacts of AI, such as specific improvements to a particular process, and broader business gains such as cost reductions or customer retention - and to derive relevant further actions to take. "It is essential that

granular-process KPIs fit tightly into overall strategic goals. This requires consistent alignment in the direction of travel of a business throughout all levels of management and teams," explains Jenaele Howell, vice-president of the applied intelligence group at Informa Tech. "The overall goals for the systems must consistently cascade right down to the technology, operational, and day-to-day levels."

Across sectors, companies are working to improve how they measure and demonstrate the value of AI. Examples include a retailer offering virtual digital assistant services via online channels and using AI to link customer satisfaction to different elements of those interactions so that processes and outcomes can be continually improved. Elsewhere, a major document technology company has invested extensively in AI to consistently assess the performance of its software-as-a-service solutions and make regular necessary refinements.

It is an explosive period for AI adoption and we expect that the curve will soon rise rapidly, with a tripling of projects in just a few years

Typically, the greatest success is being achieved by businesses first identifying clear strategic problems, designating who will bring about and measure progress with regard to them, and taking a pragmatic, methodical approach to the technology's deployment. As Beccue notes: "These organisations are practical in their approach and clear on who is responsible for measuring and execution." Among the technology's innovators and end users, along with academics and other experts to meet, share experiences, discuss best practice ideas and help businesses overcome key challenges. At the upcoming AI Summit in London, attendees will hear from sector counterparts and AI technology experts, and learn about current projects, the most effective new ways of working and what lies ahead.

With artificial intelligence projects beginning to reach critical mass, their prominence within enterprises is growing rapidly. But so too is demand for reliable measurement of material successes and clear reporting of returns on investment. Establishing the right KPIs and oversight structures will be fundamental to ensuring these projects consistently unlock the important strategic results that executives expect.

necessary outcomes. However, the risks they face with AI continue to arise, not least the threat of regulatory or legal problems. "People are definitely concerned about how to comply fully with regulations around AI and use of data in general, which began with the important GDPR rules in Europe. Further AI rules are on the way and concerns have spread to threats of major litigation, including around the idea of who specifically is responsible when AI makes decisions that cause negative consequences," explains Beccue.

While such concerns could drag on project progress, the pressure on businesses to stay ahead of competitors from within and beyond their industries remains unrelenting and is sure to motivate a continued acceleration of AI, as well as a sharper focus on how to measure the impacts.

Many organisations are working with Informa Tech to master their core capabilities around AI. The company's AI Business Toolkit provides best practice reports, peer and vendor benchmarking, market trackers and business performance metrics, helping optimise AI, improve practice, ease system selection and establish a new competitive edge. Meanwhile, Informa Tech's dedicated AI Summit Series enables the technology's innovators and end users, along with academics and other experts to meet, share experiences, discuss best practice ideas and help businesses overcome key challenges. At the upcoming AI Summit in London, attendees will hear from sector counterparts and AI technology experts, and learn about current projects, the most effective new ways of working and what lies ahead.

With artificial intelligence projects beginning to reach critical mass, their prominence within enterprises is growing rapidly. But so too is demand for reliable measurement of material successes and clear reporting of returns on investment. Establishing the right KPIs and oversight structures will be fundamental to ensuring these projects consistently unlock the important strategic results that executives expect.

Mark Beccue will be speaking about artificial intelligence KPIs at the AI Summit in London, which runs from 15 to 16 June. Get your complimentary pass at london.theaisummit.com



Commercial feature

Bolstering human expertise for more agile decision-making

Effective decision-making is critical to enterprise success. A new category of tech is helping firms grapple with complex data, siloed operations and macroeconomic challenges

Executives have a huge appetite for actionable information. Across sectors, most are constantly obtaining new data sources and attempting to assimilate and interpret them in order to react to current realities or effect industry disruption. Being able to fully harness analytics in real time can bring a strong competitive edge.

"Organisations are becoming ever more sophisticated in modes of operation, from agile development and lean manufacturing, to just-in-time and just-in-case supply chains," says Marc Warner, chief executive and co-founder of the technology company Faculty. "The problem is that they face major obstacles to delivering upon the promises of these ways of working."

Businesses are struggling to distill the sheer volume of information in order to make the best and most agile decisions, particularly in the current context of large, varied macroeconomic shifts. These difficulties are often accentuated by managers' tendency towards holding information in silos, which can lead companies to miss the broader picture entirely.

Using AI to empower humans is not only the best use of the technology, but the best way to make good decisions

"The lived experience for many decision-makers is getting into meetings where it is hard to work out what the problem really is and what can be done to fix it, because excessive and conflicting information is causing unnecessary and heightened confusion," says Warner. Ultimately, these challenges cause near-constant obstacles and end up hindering a business' competitiveness, leading to potential lost profits.

In increasing numbers, businesses are attempting to move away from the disjointed answers and improve the situation by using advanced AI to refine complex decision-making and optimise supply chains. A new category of AI is emerging that is reclaiming the entire decision-making process.

The right way to apply AI, Warner says, is to use it consistently alongside the expertise of humans. Executed well, decision intelligence can unlock understanding, drive relevant actions and connect all parts of an organisation. He explains: "Using AI to empower human experts is not only the best use of the technology, but it's the best way to make good decisions. AI can handle the heavy lifting such as data extraction and trend identification, surfacing the deep insights in the right formats for expert decision-makers to focus on pertinent information. This results in well-informed choices."

Organisations across sectors are using Faculty Frontier, the decision intelligence software, to unlock a better understanding of their operations and surrounding conditions. The technology draws from a multitude of data, presenting simple and actionable insights. This shows business leaders exactly what is happening and what will happen inside their business - and crucially why. All areas are covered, from supply and demand, to distribution and strategic planning.

Among the many businesses working with Faculty is one of the largest US food manufacturers. By accurately predicting how many people would work on production lines every day, it has been able to drive more effective resourcing and ensure production meets demand. Faculty's forecasts for the manufacturer were shown to be 99.5% accurate in a retrospective assessment, a notably high success rate for predictive processes.

In the UK, the NHS uses Frontier to meet critical staffing and supply chain challenges in highly complex and fast-changing clinical environments by anticipating demand up to three weeks ahead. Having successfully forecast admission rates, and the lifesaving equipment needed, for hospitals in 137 trusts during the first year of the Covid-19 pandemic, the system is currently being rolled out across many A&E departments to predict arrival levels and ensure managers are prepared for incoming demand surges.

Enterprises increasingly count on shrewd and future-ready decisions to drive simplicity, efficiency and profitability. As collating and understanding data becomes more important, AI adoption will be crucial for business leaders to support expert human decision-making and gain a clear competitive advantage.

To find out about Frontier, which combines human and machine intelligence for better decision-making, visit faculty.ai/frontier

faculty

INSIGHT

'The ultimate measure of effective AI governance is minimising the harm caused by AI deployment'

A Q&A with Sue Daley, director of tech and innovation, techUK, on the importance of AI governance

What is the importance of robust AI governance?

One of the incredible things about AI is its range of applications across tasks and industries, where different types of AI software can transform practices to increase efficiency and improve outcomes. This diversity, however, also translates into a variety of risks that may arise or be exacerbated by AI technologies. Robust AI governance is a mechanism for safe and responsible deployment of AI.

Recent years have seen many high-profile cases of AI gone wrong. To prove worthy of people's trust, we must be able to address the risk of harm throughout the AI lifecycle, and to take the task of governance seriously.

Good AI governance may not prevent all negative outcomes, but governance approaches within any business area cannot promise perfection. It will, however, allow organisations to spot potential and actual harms faster and therefore rectify them in a timely fashion.

What respective roles do legislatures and private enterprises play in the development of AI governance best practices?

AI governance is a collaborative endeavour, and private companies, civil society, governments and international bodies all have a major role to play. As AlgorithmWatch has documented, as of April 2020 around 170 AI ethics guidelines had been published, and the continually updated OECD AI Policy Observatory shows that both public and private initiatives are being launched.

It's encouraging to see that these initiatives are becoming concrete and

practical. Within the tech industry, governance initiatives are moving from the stage of outlining principles to implementing processes specifically intended to catch adverse consequences of AI applications.

Legislatures play a natural role in securing their inhabitants' wellbeing, and may for those purposes bring forward AI-related policy and regulatory action. But while good AI policy helps create baseline consistency and sets expectations among the public, seeking compliance with legislation is only one part of AI governance.

Many private companies have measures in place which go above and beyond minimum requirements – and it is not necessarily the case that more heavily regulated jurisdictions also warrant higher internal company AI governance standards. Rather, it depends on the company's own grounding in AI governance and how this aligns with their principles.

How can AI governance be managed and measured in private organisations?

AI governance practices vary between organisations, and it is too early to say conclusively which approaches are more effective. The 'right' approach is likely to be different depending on each organisation. If a project is potentially risky, some techUK members create groups spanning legal, engineering and policy departments to look further into how these may be mitigated. This illustrates how different skill sets are important in the process of AI governance. Some bring the technical background to interpret models behaviour, others will understand legal or political implications and others again may

provide perspectives on the higher-level ethical implications.

The ultimate measure of effective AI governance is the extent to which harm caused by AI deployment is minimised. This can be difficult to measure, however, as organisations cannot know exactly what would have happened if they had chosen a different approach.

It can therefore be useful for organisations to have more 'bite-sized' metrics to check the success of their AI governance approaches. It would be useful for organisations to register how often AI projects are flagged for governance purposes, how often concerns are considered to be substantial and what actions are taken as a result. Crucially, if issues occur with AI products or services which have been launched, teams should go back to check what might have been missed in the governance process. AI governance is not over after a one-off, tick-box exercise. ●



Sue Daley
Director of tech and innovation, techUK

REGULATION

Brussels responds to call for clarity

The EU is set to become the first major market to create a legislative framework specifically covering AI tech. There are hopes that others will follow its lead

Daniel Thomas

When Microsoft unleashed Tay, its AI-powered chatbot, on Twitter on 23 March 2016, the software giant’s hope was that it would “engage and entertain people... through casual and playful conversation”. An acronym for ‘thinking about you’, Tay was designed to mimic the language patterns of a 19-year-old American girl and learn by interacting with human users on the social network. Within hours, things had gone badly wrong. Trolls tweeted politically incorrect phrases at the bot in a bid to manipulate its behaviour. Sure enough, Tay started spewing out racist, sexist and other inflammatory messages to its following of more than 100,000 users. Microsoft was forced to lock the @TayandYou account indefinitely less than a day later, but not before its creation had tweeted more than 96,000 times. AI systems are, of course, adaptive, learning from cues in their environment and changing their behaviour autonomously. This means that, once they’re put in place, unforeseen ramifications may ensue, for which humans can quite easily avoid responsibility. They can also be an invisible hand that discreetly influences our choices to an extent that goes beyond our understanding and/or wishes. As these risks proliferate, legislators have stressed the need for clearer and more consistent regulatory frameworks to deal with them. The world’s major economies have yet to establish such measures

so far, relying instead on a suboptimal patchwork of old laws and standards to police the industry. But that will change when the European Commission’s proposed Artificial Intelligence Act becomes law. Heralded as the world’s first legal framework designed specifically to cover AI, this will seek to identify and regulate higher-risk forms of the technology – biometric identification, for instance. It will impose far-reaching obligations on developers, covering standards of governance, design, transparency and data security. Those that fail to comply may be subject to hefty fines. Once passed, the act will still need to work its way through the European Parliament for adoption. And after that, it will be at least a couple more years before it becomes enforceable. Other key markets have yet to design an AI-specific regulatory regime, although China issued guiding principles for its regulation in 2017 and, this March, a data privacy law called the Internet Information Service Algorithm Recommendation Management Provisions. The UK is awaiting the publication of a government white paper on AI regulation, while the White House started preliminary discussions on the need for what it called “a bill of rights for an AI-powered world” in November 2021. Dr Mariarosaria Taddeo is an associate professor at the Oxford Internet Institute and a faculty fellow at the Alan Turing Institute. She believes that a clear “Brus-



sels effect” is setting the pace for other administrations to follow. “Starting with the General Data Protection Regulation and continuing with the Digital Services Act, the Digital Markets Act and now the AI Act, the EU has created a framework for the coherent regulation of digital technologies,” she says. And, in her opinion, any tech provider seeking access to the single market will have to abide by it. “I suspect that we’ll be moving on to a

“We know that bias is an issue that crops up in numerous machine learning systems

‘transatlantic effect’, because the EU and the US have strengthened ties over the past year and sought shared points to align their regulations. Other markets are likely to follow,” she notes. Part of the problem with the current mishmash of regulations is that it creates inconsistency. AI-based businesses would be much better served by a set of shared international standards if they are to build a thriving global industry. Dr Cosmina Dorobantu, co-director of the Alan Turing Institute’s public policy research programme, notes that, in practical terms, AI is a “general-purpose technology” that touches all sectors. This means that some issues require a common approach, yet often they will be policed by separate regulators, each of which may take a different approach. “We know that bias is an issue that crops up in numerous machine learning systems,” Dorobantu says. “We will see it in credit scoring algorithms used by mort-

gage lenders, facial recognition technologies used by police forces and automated triage systems used by hospitals, to name but a few examples.” In the UK, these applications fall within the remit of different watchdogs, which increases the risk of regulatory arbitrage, whereby users seek the path of least resistance, she warns. Also, important problems could fall through the cracks if one regulator were to assume – wrongly – that another authority is handling them. The EU’s AI Act is not without its weaknesses. The European Commission’s consultation on the draft legislation attracted more than 300 feedback submissions – a far stronger response than its other tech bills have generated. Criticisms have focused on the act’s overly broad definition of AI and what it classes as high-risk uses of the tech that will be subject to more stringent controls. Human rights campaigners, meanwhile, are concerned that the act is not strict enough in con-

AS OF 2020, THE UK HAD:

£63bn

spent on AI technology and AI-related labour

1,300

AI companies – representing a 600% increase since 2010 – which attracted...

£1.8bn

of funding

1.3million

UK business are expected to be using AI by 2040

Gov.uk, 2022

trolling the use of AI in law-enforcement applications such as predicting criminals’ behaviour and conducting mass surveillance using facial-recognition systems. As the global AI market develops, we may see flashpoints when different jurisdictions’ approaches to regulation clash, Taddeo predicts. The EU’s approach to digital governance has historically been based on values such as human dignity. In the US, the focus has been much more on preserving freedom, of speech or of markets. But both value democratic values and basic human rights, even if they may interpret these differently. China’s attitude to such matters will be harder to reconcile. The EU has already vowed to ban AI systems such as Beijing’s so-called social credit system, which enables organisations and individuals to be tracked and evaluated for trustworthiness. Despite their differing approaches, democratic countries are likely to seek regulatory alignment wherever possible as they seek to unlock the opportunities created by AI, according to Taddeo. Even if they can’t achieve full alignment, she says, some measure of convergence will at least “create a playground that allows different actors to collaborate”. ●

Conversational AI isn’t going anywhere—it’s going everywhere.

“Robb has always been a pioneering mind in the world of experience design and he’s established a deeply compelling vision for the next phase of our relationship with technology. *Age of Invisible Machines* is surprisingly easy to follow—laying out a clear path for the inevitable work of leveraging conversational AI and making complex business operations autonomous. I can't recommend it enough!”

Rebecca Flavin
Global Experience Strategy Officer
Ogilvy

