



AI Business
eBook Series



Unlock **enterprise** **document intelligence** with AI

In collaboration with:



Infosys NIA

2

- 4** Reading like a human: How advanced AI pulls data from media
- 6** The business case: Five ways AI can help your company thrive
- 8** Can AI work for you?
- 12** A place for humans: Why Artificial Intelligence still needs the real thing
- 16** The need for privacy: Why some data must be kept private, and how AI can help
- 18** What does the future hold? How can things improve?

eBook contributors:

John Kuriakose
Principle Product Architect,
Infosys Nia, Infosys

Kajari Ghoshdastidar
Senior Product Manager,
EdgeVerve

Siddharth Pandya
Director and Head of Engineering for Nia
Contracts Analysis,
EdgeVerve

eBook in collaboration with:



You're sitting on data gold and AI can mine it

Business growth in the digital world is often hindered and slowed down by the effort required to understand document data. Document understanding is much more than OCR or handwriting recognition and requires the capability to detect and recognize various structural elements in enterprise documents like textual paragraphs, sections, table, fields, checkboxes, signatures, logos etc. Document extraction, processing and comprehension requires the ability to stitch together computer vision models, NLP models and information retrieval technology into a single data pipeline and provide a simple unified experience across all kinds of documents. The unified experience and document pipeline also require some automation and business process management to complete the overall story and integrate the digitized information extracted from documents into enterprise processes and workflow.

Artificial intelligence (AI) is the critical piece to enable comprehensive document extraction, processing and comprehension. This is because, only AI can handle this document variation tolerance with its innovative approach and can take contextual decision on the fly. This is where traditional rule-based approach fails, RPA fails. This doesn't just apply to text, either. With advances in AI and machine learning, any source of data can provide extra information: pictures, audio, infographics – even video or speech.

In people-hours alone, the benefits to a business simply can't be overstated. Legal teams that are used to spending hours sifting through documents for crucial information can locate relevant passages in minutes. Insurance and estate agents can access detailed historical asset and customer information collected in one place, neatly structured to be easily processed. And that's just scratching the surface of what is possible with AI.

In this eBook, sponsored by AI-based software company EdgeVerve, we take a look at the following artificial intelligence technologies, Document Digitization, Computer Vision, Natural Language Processing and Intelligent Search, as applied to document management. We investigate use cases, report market projections, and interview the experts. We hope this publication will equip you with the knowledge to succeed in today's volatile market conditions.

Alan Martin | Associate Editor | AI Business

AI in the workplace: An expanding market with room to grow



Artificial intelligence has been generating media hype for a while now. But having now proved its worth in specific applications, it is moving from an experimental phase towards becoming an essential part of many businesses, and their spending plans.

What does that mean in terms of market projections? According to a recent report from analyst firm Omdia, ([web link](#)), annual AI software revenue totaled \$16.4 billion in 2019, and will grow six-fold in the next six years, hitting nearly \$100 billion by 2025.

“The global AI market is entering a new phase in 2020 where the narrative is shifting from asking whether AI is viable to declaring that AI is now a requirement for most enterprises that are trying to compete on a global level,” Keith Kirkpatrick, principal analyst at Omdia, said.

The Natural Language Processing (NLP) technologies specifically are expected to generate \$9 billion per year by 2025, and a large slice of that comes from document AI: the structuring of unstructured data for

easy human consumption, and the resultant opportunities for computer analysis.

Which industries are doing the best job of capitalizing on the new tech? Omdia estimates this by total spend, and on this metric the biggest early adopter is the healthcare sector. “It’s electronic healthcare records and how you work through some of the biggest problems in healthcare,” Mark Beccue, another principal analyst at Omdia, specializing in AI, tells me. “You have this issue around the world where records are not indexed very easily. You have physicians writing notes and administrators trying to figure them out.”

All these medical notes – this unstructured data – can be interpreted via NLP. “That’s a long-term trend that we think is going to continue over time,” Beccue says. As for adoption elsewhere, “we’re watching some of the biggest enterprise application companies like Oracle, Salesforce, SAP, those kinds of guys,” he says – in other words, customer relationship management (CRM), and other


documents-based processes. Certainly not least, in the financial services, another industry with money to spend.

It’s not surprising that the largest industries are the ones that catch the eye with early investments, but there’s plenty of potential for Document AI outside of traditional enterprise space. Take a lawyer’s office: fully indexed and structured data would save the need to search through piles and piles of documents and case records. “The legal business model has forever been hourly rates, but I’ve talked to people in that profession and some clients won’t pay for research,” Beccue says. Data ingestion that utilizes AI systems has the potential to massively reduce research time, allowing clients to pay for the expertise.

Are there any industries that wouldn’t benefit from this type of automation? “I’m not sure there’s a whole lot,” Beccue says, before highlighting businesses without a need to keep big central databases: think plumbers, or home services. “The more scattered, the more decentralized an industry is, the less they would use these technologies,” he says.

For the moment, anyway. These projections are based on the current ground costs of introducing AI in the workplace, which remain high due to the expenses involved in training AI models using massive data sets. Some of the world’s brightest minds are working on reducing the barriers to adoption, and if they succeed, it may well be that Omdia’s eye-watering growth projections end up being on the conservative side.





Reading like a human: How advanced AI pulls data from media

Analyzing documents for contextual meaning is an essential skill for humans, who are accustomed to seeing data being presented on slides or posters, in advertising, infographics, and email. For AI systems, this presents its own set of challenges – something the technology simply wouldn't have been able to cope with just a few years ago.

Even a simple text was a problem at the turn of the century. Remember how CAPTCHA boxes on websites used to ask you to type in a couple of random words on screen? These weren't as random as they first appeared: as well as proving you were human, these little tests were helping Optical Character Recognition (OCR) software recognize words that it couldn't decode on its own. People typed in enough words that, by 2011, every tome in the Google Books

archive had been digitized, along with 13 million articles from the New York Times ([web link](#)).

We've come a long way since. In fact, with enough human-provided training, AI has proved suitable not just for digitizing books – a relatively simple task in the greater scheme of things – but for decoding and interpreting many different types of media. That's not limited to plain text; as a general rule of thumb, documents can be divided into four types:

- Continuous text (e.g: books)
- Form-based documents (e.g: receipts, invoices, forms, etc.)
- Visually-rich documents (e.g: brochures, magazine articles, infographics, etc.)
- Pure images (e.g: pictures, photographs, screenshots, etc.)

Let's take the third – visually rich documents – as an example. For AI to extract the data in a useful manner, it not only has to understand contextual clues from the text itself, but also be primed to handle and interpret anything else that might plausibly appear in a magazine.

Things like:

- Images
- Logos
- Word Art
- Symbols
- Charts
- Captions
- Font style and sizes
- Colors
- Backgrounds
- White space
- Positioning
- Orientation
- Symmetry



As you might imagine, OCR isn't cut out for this, as it relies on consistency of format – something which visually rich documents are not known for. Fortunately, object identification and image classification have reached 99% accuracy over the last decade, so other AI tech is ready to make up for OCR's deficiencies and decode documents designed with people in mind. Think pamphlets, posters, catalogs, digital adverts, invoices, receipts, marketing assets, and so on.

To do this, an AI engine not only has to accurately scan text; it has to have the basic contextual understanding of what things like font size and spacing actually signify to the human eye. This holistic interpretation of a document is known as “multi-modal image extraction” and combines OCR and computer vision to extract the pertinent information, and turn it into structured data.

Consider the following example, shared in the image below.

In January 2018 a hurricane caused severe damage to the roof, walls and the garage. I raised a claim for \$7850 with Geico Insurance on my policy: A248822.
I am interested to know what natural disasters are not covered under the new policy.

The AI system will be used to comprehend the disparate elements of the document to figure out the meaning.

While this might be simple for a human to understand, simple OCR technology would falter. This may feel like a difficult process, but it's actually not too far removed from how the human brain would interpret the same information. And while outlining the stages in a step-by-step fashion may make it seem like slow work, an AI platform actually processes these stages in the blink of an eye, making the consumption of hundreds of thousands of documents something that's eminently achievable.

Crucially, that speed doesn't sacrifice detail – indeed, AI in general has been known to pick up details that even the most specialized humans miss. Physicians trying to tell the difference between restrictive cardiomyopathy and constrictive pericarditis have a

success rate of around 56% – an AI looking at exactly the same echocardiograms managed to push the success rate up to 96% in just two months ([web link](#)).

Challenges ahead

The potential is clearly there, but mass adoption of Computer Vision is not without its challenges. Chief among these is the vast amount of data that an AI model needs for training in order to be effective – in reality, such datasets are hard to find, which is why a great deal of research is currently being undertaken on training models that require less detailed data. Active learning, semi-supervised learning, unsupervised learning with data reconstruction, noise injection and N-shot learning are all areas being explored to make this dream a reality.

Computer Vision is also intensive in terms of computational resources. There's no two ways about it: deep learning models are costly. Still, this is a comparatively simple problem that will likely take care of itself as we see advances in hardware and cloud solutions. And when we do, AI platforms will be appearing everywhere.

The business case: Five ways AI can help your company thrive

Humans are smart on an individual level, but generally don't deal well with scale. Artificial intelligence is the opposite



The business case for introducing AI in the workplace may feel a bit abstract, and a tough sell to people with hands on the pulse strings. Here are five tangible benefits that introducing NLP and document management AI can bring to your business.

1: Time saved

"Time is money," as the old adage goes, and by far the biggest selling point for structuring that unstructured data with AI is the sheer number of admin hours saved when your data is neatly indexed and searchable.

Whether you run an insurance company with a complex range of policies and exclusions, or a healthcare company with thousands of hand-written doctor's notes, AI will save time – and ultimately, money.

2: Making the most of your employees' skills

The time saved is good enough on its own, but it's the knock-on effect that's really exciting. With less time spent in the 'needle in a haystack' search for pertinent data, employees have space to use the skills that you actually hired them for in the first place.

Whether that's medical expertise, a sharp legal mind or the ability to upsell, employees who have more time to put their unique skills to good use are not only more productive – they're happier, too.

3: An improved customer experience

With structured data easily searchable, you'll be surprised at the improvements available in terms of customer service.

Taking the example of the insurance company again, setting up a customized

policy might once have taken hours to complete, while all the information about a customer's unique circumstances was collated and evaluated. With their credit history, demographics, policy options and possible exclusionary risks highlighted in a searchable database, setting up with a bespoke policy could be done in a single phone call.

4: Improved document security

With data breaches a daily occurrence, and with governments around the world upping punishments for those that let them happen, security has never been more important.

Document AI can help here, too. Properly trained artificial intelligence systems can scan for sensitive information and automatically redact it, as and when required. The same systems can be on the lookout for unusual activity, warning you of a possible data breach before it happens.

5: Unexpected business insights

Humans are smart on an individual level, but generally don't deal well with scale. Artificial intelligence is the opposite: the more data it has, the more unexpected insights it can produce.

You may have heard headlines about AI making correct medical diagnoses that trained doctors miss, and when fed on millions of business documents Artificial Intelligence may well have some unexpected insights for your company too. Who knows what AI could spot in years' worth of sales data, or by analysing thousands of resumes from successful job candidates?



Can AI work for you?



One of the companies at the forefront of the Document AI revolution is EdgeVerve – a subsidiary of Indian IT giants Infosys, founded in 2014. One of its products – Nia DocAI ([web link](#)) – a document extraction, processing and comprehension platform that unlocks business value from enterprise data by extracting intelligence from enterprise documents, regardless of complexity or domain specificity. The company has helped businesses across all industries in transforming digital journeys from fintech to the aerospace.

AI Business sat down with some of the people who have overseen the development of Nia DocAI: Principal Product Architect, John Kuriakose, Senior Product Manager Kajari Ghoshdastidar, and Director and Head of Engineering for Nia Contracts Analysis, Siddharth Pandya.



John Kuriakose
Principal Product Architect,
Infosys Nia, Infosys



Kajari Ghoshdastidar
Senior Product Manager,
EdgeVerve



Siddharth Pandya
Director and Head of Engineering
for Nia Contracts Analysis,
EdgeVerve

kind of analytics product on top of the data platform. That set in motion a plan to create an AI-based platform that gives enterprise customers everything that they need to execute AI projects.

From there, we found other market opportunities. For example, we found out that the space around scanned document digitization was opening up. The contract solution initially assumed that you had digital-native documents, which means documents that have digital text included. But that was not always true: when customers had scanned documents, typically they had to first digitize those into digital text, and then pass the result to the contract solution. So we found out that digitization is both required to improve the value of contracts, but also that there was a separate market for digitization.

AIB: Can you start by telling our readers how EdgeVerve came to be working in artificial intelligence?

JK: I think our first foray into the document space came in 2016, when we started with an offering around the legal documents, which eventually became what we call Nia Contracts Analysis. It was just an experiment in a focused, closed domain where we could apply some of the AI-based techniques to derive intelligence from documents.

We discovered that there was a huge market waiting for this kind of thing. We found that there are many adjacencies available – not just legal documents, but any types of contracts, everything from rental contracts to artist contracts; every time we had a new use case, we kept improving the product capabilities.

By 2016, [contract analysis] had done well and we knew we needed some

AIB: It sounds like DocAI was shaped by internal design, rather than by clients approaching and asking “would this work for our business?”

JK: It’s a mixture of both, actually. As we went to market with contracts, we discovered the clients’ specific needs, and then we also had our mechanism to learn from the market.

In hindsight, I would articulate three broad dimensions that we found our clients needed. First, the breadth of AI: to be able to work with structure, data and prediction problems. Secondly, they wanted to have useful, integratable, easy to consume AI. Thirdly, we discovered it is important to talk about IT readiness. So questions like, “How will you scale? How can I integrate my documents? How do I bring the documents in? How do I pull the results out? What kind of security do you support?” All of the questions that a typical enterprise architect or



an IT director would have around a product coming into the data center.

AIB: Can you take us step-by-step through the process of DocAI tackling client's raw data and making it manageable?

KG: There are lots of paper documents, significant documents, unstructured data, unprocessed data that is lying there but nobody can use. The first step is the digitization of data. You can think of a document as a collection of different elements, but what are those elements? There are visual objects in the data, graphs and charts and tables, logos and, most importantly, the text.

The next step is to make sense out of the data you have extracted. Take a form with the name "Alan" on it: just by looking at it, the AI doesn't know what "Alan" means, till it is paired with the keyword (field) "Name:", printed on the left of the word "Alan". So the DocAI platform will do a full field value pairing to give a meaning for the values in the document. Legal contracts,

insurance policies, annual report book... each keyword we pick up, we try to find a meaning for it whether it's a proper noun or a kind of business clause.

Then we try to understand the intention of the text blocks, whether it's a termination clause in a legal document, or explaining some medical procedure. After we've got a better sense of the information of the document, we'd also try to understand the context from where it is coming up so that we can really zone in and try to get the holistic information out of it.

We also do classification of documents. So if somebody uploads a bunch of documents, which is a collection of checks, or invoices, purchases or sales orders, we can automatically separate them.

AIB: That's how the data is extracted. Then what happens?

KG: Then there's the analytics. We can do document comparison. We can use other advanced techniques to understand the sentiment. We can

find the historical implications, and based on those historical implications, we can predict something. We can highlight risks.

After that, when we have extracted this actionable insight, we have to make sure that this data – this information – is available for consumption. Unless we allow the business user to consume it when needed, then what's the point of having so much information?

AIB: How would you sell the benefits of DocAI to a company?

KG: The crux of the problem is the sheer volume of data. It's not that humans aren't intelligent, it's that machines save time.

JK: Along with productivity, you also want some consistency in the execution of the process. Depending on the type of the document, you have to pull out certain information out of it, make some decisions, and then update some enterprise system. ➤

➤ I'm simplifying it, but if you are to do this with 100 documents, that's one thing. If you have to do this repeatedly, or work 90 days with 1,000 documents produced every day, that's another level of complexity.

AIB: You've already mentioned insurance, but what other industries do you feel really benefit from AI assisting with document management?

KG: If you look at the market sizing, it's banking, financial services, insurance and healthcare, they comprise 50% of the growth in document processing.

JK: There are types of documents that are applicable in any industry: invoices and account payable functions, where they have to triage between an invoice, a purchase, and a delivery receipt, for example. You can be dealing with telecoms, logistics, or banks, and you'll have that. In fact, the whole logistics space is something that we have out-of-the-box readiness to deal with.

Different industries have different kinds of requirements, though. If you look at an automobile manufacturer, the challenge that one of the customers

came to us with was they had a whole set of parts and suppliers, for each part across their models. And when a field engineer is talking to a service dealer, they have to quickly answer questions on part availability, part replacement alternatives, dimensions, and things like that. But the documents that are used to answer these questions are basically engineering drawings.

So that's an example of where you might need to first digitize and pull out all of that information from the engineering drawings – but that's not enough. You need to open it up on an interface where it becomes very easy for you to find the information you are looking for.

SP: At the time when the CARES program was launched in the US, we were helping one of the banks to process loans. The volume of loan applications was enormous, and we were helping the bank analyze various application forms and supporting documents to assess the loan eligibility using AI models.

Simultaneously, an orchestrated algorithm was run in the background

to make recommendations on which of these loan applications could be approved.

AIB: And are there any industries which wouldn't benefit from AI-based document management?

KG: If you want an answer in one word, it's a "no". Because we are talking about artificial intelligence, and that means we are trying to mimic the biological intelligence of humans – or at least, that is what AI aspires to do. So wherever human intervention is needed, AI is applicable.

Today, the degree of uptake is not due to a lack of need, but because of technology limitations. Say, for example, in healthcare, AI is doing wonders. It's helping doctors to find lots of insights into diseases, and helping them to diagnose patients properly. But still, AI systems are working as a guide here, they are not replacing the doctor.

I would say, wherever common sense is still saving the day, that's where, to this day, AI has a little bit of a backseat, because of the fuzziness in decision-making that a human has.

AIB: So essentially, everyone should benefit when the technology advances enough?

KG: Yes, absolutely. AI is pervasive, and there is a reason why it is pervasive.

JK: AI is definitely applicable across industries, but if you talk specifically with respect to DocAI, the simple question to consider is how critical are documents to the core business processes of that industry?

The ones that we already talked about – insurance, banking, manufacturing and logistics – obviously these heavily depend on documents. But some are less dependent – for example, retail. They're more dependent on structured data than on documents. So you find much less need for text analytics or document recommendations that come from analyzing unstructured data.



“AI is definitely applicable across industries, but if you talk specifically with respect to DocAI, the simple question to consider is how critical are documents to the core business processes of that industry?”





A place for humans: Why Artificial Intelligence still needs the real thing

“AI technology is good, because it handles the sheer volume, which is beyond the capacity of a single human, but at the same time, we can never ensure that AI can guarantee 100% accuracy all the time.”

Artificial intelligence systems have a number of advantages over humans, and when it comes to document management, the most important one is obvious: AI never gets bored. “Crunching a huge volume of data and comprehending it, is something that humans struggle with,” explains Kajari Ghoshdastidar, senior product manager at EdgeVerve.

“Scalability is something that doesn’t come easily to humans. We are telling machines to read hundreds and thousands of documents from different data sources to find out, directly or indirectly, which factors could impact a business decision. For humans, they need a big team to do it, and they need months and years.”

In other words, it’s simply not a fair fight, and you can’t get too far into a conversation about artificial intelligence without arriving to the question: if AI is several orders of magnitude more productive than humans, doesn’t that just spell job redundancies at best, and human redundancy at worst?

As with most things in the AI space, there are no definite answers. But to assume that wider use of AI means less need for humans would be to fundamentally misunderstand both how AI works, and its actual purpose: to augment and build upon the human experience, rather than to replace it.

Humans at the heart of AI

“For us, keeping humans in the loop is a very essential part of the process,” Ghoshdastidar says. “AI technology is good, because it handles the sheer volume, which is beyond the capacity of a single human, but at the same time, we can never ensure that AI can guarantee 100% accuracy all the time.”

In other words, for the foreseeable future, AI will need human input to decode unstructured data. Dr Riza Batista-Navarro, a lecturer in Computer Science at the University of Manchester, has a great case study in this, via one of her NLP PhD students, Gavin Abercrombie ([web link](#)), and his work analyzing over 100 years’ worth of speeches and debate ([web link](#)) in the UK’s Houses of Parliament.



On paper, that should be easy enough, as Batista-Navarro herself highlights: “there are many types of data sitting there.” But the strange conventions of Westminster would be lost on a conventionally-trained AI model. “The language that they use is ambiguous, or it’s indirect, with a lot of hedging,” she says. “By regulation, or according to guidelines, they have to be polite. They’re always subtle, but there’s an underlying message: they’re saying ‘I don’t like this,’ or ‘I like this.’”

In short, even in this environment with – in some respects – a perfect supply of data, algorithms need a human guiding hand. In this instance, Abercrombie used the voting data of British MPs to guide the model into understanding their actual sentiment – but even that’s imperfect. “MPs might say something in their speeches which are not necessarily consistent with how they voted,” she explains. “They have to stay loyal to their parties.”

Sentiment analysis, in other words, is a tricky field for AI to navigate. “Humans have a way of communicating that

machines will never really completely get,” explains Mark Beccue, a principal analyst at Omdia, who specializes in AI. Face-to-face communication is best for data density, he says, because of the visual cues, followed by voice-only, and finally, text-only. By that point, things like detecting mood and sarcasm come down to your previous knowledge of the person, rather than anything that can necessarily be gleaned from the text itself.

As Dr Batista-Navarro points out, even the richest companies on the planet occasionally make a misstep here – and through no fault of the software. Take Facebook, which has often been criticized for its AI-based moderation tools being either too aggressive, or too lax. “Part of the reason, I think, why the automatic content moderation is failing is because what could be hateful or offensive in one context or for one culture is fine for another,” she explains.

In other words, humans need to be a part of language analysis, and they

need to be there right from the start. “I certainly believe that humans still need to be in the loop, as we say, or involved in the process,” Batista-Navarro says, “and at several points.” Humans are needed to definite classifications and consider the purposes the data might be used for once structured.

“If you have more complicated types of questions, for example, those which require some sort of reasoning over temporal relationships and so on, and those types of relationships are not represented enough in the data set, then that means that the model won’t be able to learn those kinds of things,” she explains.

That’s not to say wider adoption of AI won’t lead to job losses, of course, but opinions on the size of the impact differ; the idea that humans will be made redundant by machines in some form or other is almost 100 years old. As Jared Robert Keller, a specialist in the history of automation and labor once told me ([web link](#)), there’s a danger in treating the job market as ➤

➤ a zero-sum game where one job taken by a robot means one less slot for a human.

“We know the job market is constantly changing and creating new occupations,” he said, highlighting cyber security – a field that simply didn’t exist until a few decades ago. “I think if we’re going to say that ‘we’re the generation that proves the rule, where mass unemployment hits,’ we have to be very clear about what makes us the exception.”

There’s also the question of where redundancies occur. And on that subject, Omdia’s Mark Beccue believes that, at least in the short term, it’s largely in jobs with low satisfaction ratings and high turnover – such as call centers. “What’s happening is the automated agents are taking the redundant work and they’re assigning the humans to better work. The kind of work that needs human emotional intelligence to deal with complex issues, beyond AI’s current skillset.

“The more linear the thinking involved in a job is, the more likely it will go away with AI,” Beccue adds. “So humans are going to do nonlinear thinking and emotionally intelligent work. That’s where we’re headed.”

Optimization, not replacement

Overall, the current trend for document management AI is about optimizing human workflows, rather than eliminating them. As John Kuriakose, principal product architect at EdgeVerve, puts it: “We are amplifying the human element, not replacing it.”

For him, AI offers consistency in an area where humans are typically inconsistent. “Depending on the type of the document, you have to pull out some information with it, make some decisions, and then update some enterprise system.” Doing that for 100 documents could be more or less the same, but scale that up to



The more linear the thinking involved in a job is, the more likely it will go away with AI

thousands every day and consistency will inevitably vanish. “There are a lot of variables that affect the overall performance.”

Then, of course, there’s the simple hourly rate. Spending hours scouring documents for information costs time, and time is money. Having AI to sift through the unstructured data frees up huge amounts of time that can then be used on things that clients are less reluctant to pay for.

There’s no shortage of industries where this can be applied. Doctors can access complete patient histories at the touch of a button, real estate agents can find area and property information in seconds, and so on.


The future of work?

Could we ever reach a stage where humans are no longer required? “I

don’t think so,” Beccue says. “Human communication being as difficult as it is, it’s just too difficult for machines to ever get to that point where they can fully understand what we’re trying to do.” While he can see natural language in analytics being good enough, the external interface feels like a stretch.

“Wherever common sense is still saving the day, that’s where AI takes a bit of backseat,” Ghoshdastidar says, adding that the “fuzziness” of human decision-making makes us indispensable. “The character recognition tools are very, very mature,” she says, but “replacing our codified human fuzziness” is something that might never happen.

And yet, never say never: the world of AI is advancing at a rate that futurist Ray Kurzweil calls “the law of accelerating returns ([web link](#)).”



His theory is that humans made equal technological progress between 2000 and 2014 to the entire previous century – and we’re likely to make the same progress again in an even shorter time span. “I started my PhD in 2010, right?,” Dr Batista-Novarro says, “and in ten years, the area has changed a lot.

“So that’s why I find it hard to predict, because in 2010, I wouldn’t have imagined that by 2020 we would have a machine that can write news on its own, and people won’t be able to tell if it’s written by a human author or not.

“At some point, it might accelerate again.” Enough for humans to become surplus to requirements? We’ll just have to wait and see.

The need for privacy: Why some data must be kept private, and how AI can help

In recent years, people really seem to be championing the privacy cause – something that was considered a niche subject as recently as a decade ago

Anyone following the advancement of artificial intelligence will be acutely aware of the privacy issues surrounding its use. At the moment, with so much of the world's data unstructured, if any of our private information is being stored against our will, there's some comfort in knowing that finding it would be like finding a needle in a haystack. Once AI has structured that data, suddenly all that information is readily indexed.

The issue is that artificial intelligence training requires data, and lots of it. So how can privacy be integrated into the process from the ground up? Patricia Thaine is CEO and co-founder of Private AI ([web link](#)). Its software product, Galatea, can redact key information from text, images and video – think names, locations, license plate numbers, faces and other personal identifiers. So it should be pretty unsurprising to hear that she believes privacy is absolutely essential to Natural Language Processing.

“One of the reasons why we're doing this is because there are two major problems that we're dealing with in the next ten years,” she tells me. “Climate change is number one, but number two is privacy, because it's really the basis of all our freedoms.

“I was speaking to a data protection officer yesterday, who said that a few years ago, if you wanted to know what my bank balance was, you'd have to basically go to the branch where I bank.

Whereas now, you can be in China and I could be in Canada, and you can find out.”

She is reassured that, in recent years, people really seem to be championing the privacy cause – something that was considered a niche subject as recently as a decade ago. “It's so nice to see that people are caring,” she says. “It's nice to see that there's actually a competitive advantage to integrating privacy.”

She cites the explosive growth in popularity for software that puts user privacy front and center: think of the likes of Telegram, Signal, DuckDuckGo, and Brave. “Whenever there's some sort of scandal or some sort of need for privacy, you see a spike in downloads in these apps, browser options and search engines.”

But while some turn to DuckDuckGo because of a deep-seated fear of corporate intentions, Thaine is far less pessimistic as to the root causes of privacy breaches. “I think it's the lack of education on the topic and the lack of time and resources, more than deliberate desire to do bad things with the data,” she says.

“There's also that level of abstraction. You don't necessarily understand what can go wrong in a person's life if their data gets leaked. It's hard to think about all the falling dominoes.” She thinks the privacy community needs to help people understand the implications more clearly, on a macro level: “This is what happened to Bob, because you didn't take care of his data. Are you going to be more careful next time?”

\$9bn

Annual revenue generated by NLP applications by 2025 (Omdia)

A long way to go

Thaine is reassured that governing bodies seem to be taking privacy concerns seriously. “We talk about the law not catching up quickly enough to technology, but I don’t think it’s necessarily doing a bad job considering the speed of change that’s happening – how quickly communication and AI is changing how data is spread, stored, and shared.

“There are more regulations that are popping up, and a lot of what companies do is very regulation-driven,” she says, highlighting privacy-positive statements being made at the highest levels of government in Brazil, India, China, Canada and Nigeria. Not necessarily the countries you would consider being the most concerned about the sanctity of citizen data.

But an understanding that privacy should be an essential building block of all software takes a while to filter down. Thaine says she’s met a surprising number of American lawyers who are unfamiliar with the EU’s GDPR laws, for example. “If even lawyers aren’t aware of it, there’s no reason for developers to really be tuned into that. It’s really about education, and about also a general understanding of the ethics of using data that isn’t necessarily obvious.

“There’s often a mix-up between what privacy is versus what security is,” she adds. “One is a requirement for the other, but you can’t necessarily use the same frameworks of security and translate it to privacy.” Nonetheless they are, she says, “partners in good.”

For AI-based software vendor EdgeVerve and its document product Nia, balancing the need for privacy against the usefulness of the data is an ongoing struggle. AI systems are extremely good at cross-linking data, meaning that even redacted information can be revealed via logical deduction on non-redacted bits and pieces.



“Machines can connect the dots, and end up revealing sensitive data that they are not supposed to know,” product manager Kajari Ghoshdastidar explains. The way to fix this is to “aggressively mask all the data,” but that might impinge on the ability to perform useful analytical tasks. The answer to the problem is “differential privacy” – where patterns of groups are revealed without sharing individuals’ information.

This hints at a way AI can help deal with broader data privacy issues: if models are properly trained to understand what information is and isn’t private, then it is possible for systems to redact resources proactively, depending on who’s accessing them at a given time.

Certain information could be made available only to specific departments that have explicit permission from the user, for example, while business-sensitive data could be limited to those who have signed non-disclosure agreements.

Whether pushed by legislation or simply the desire to do the right thing, a strong regard for user privacy could be the most important weapon in your AI arsenal. After all, all it takes is one big data breach to push trust in the technology right back to square one – and when you consider the advances that AI, still in its infancy, has already offered humanity, that would be a great loss.



What does the future hold? How can things improve?

“The unsupervised ability to learn from general content, and then quickly re-target that to specific use cases in [different] industries is going to be significant

In just a few decades, the state-of-the-art in artificial intelligence has evolved beyond what many of us would have thought possible. And in the world of document management – thanks to a head-start with Optical Character Recognition (OCR) – AI-based systems are pretty mature. There remains, of course, plenty of room for improvement, so what exactly should we hope for, and what can we expect?

By far the biggest problem facing AI-based document management today is the amount of data AI models need to be trained on in order to be effective. Big enough data sets are quite hard to come by.

For this particular problem, transfer learning could be a game changer. It’s exactly what it sounds like: the ability to apply knowledge from one domain to another.

“The problem is you have to keep re-training the models for every specific industry, and within the industry, to specific use cases,” explains John Kuriakose, principal product architect at EdgeVerve. “The unsupervised ability

to learn from general content, and then quickly re-target that to specific use cases in [different] industries is going to be significant. It’s already emerging in the research space, but it will take a couple of years to get industrialized.”

Dr Riza Batista-Novarro, a lecturer in Computer Science at the University of Manchester, agrees this area of research is hugely significant – especially for document AI projects. “Even if you have very specialized domains, which probably means that you don’t have a lot of labeled data, transfer learning will enable you to fine-tune in that specific domain and produce a model that is effective enough.

“For the past 10-20 years, very few domains have been able to make use of NLP and machine learning because of the lack of large amounts of training data.” If that changes, niche areas previously immune to insights generated with AI may be able to benefit in the way that areas like fintech and healthcare have.

Then, there’s summarization. Although AI models are good at picking out the



pertinent bits of a document, turning it into a brief digest is something that AI models still struggle with. It's an area of research that Siddharth Pandya, EdgeVerve's director and head of engineering, is keeping an eye on. "Whether it's text, speech or video summarization, these are hard problems to solve and some thing that I am following closely," he explains. "Because the moment we crack this problem, gigs [gigabytes] and gigs of data can be very easily summarized, and this would expedite decision making."

Finally, there's the toughest nut to crack: sentiment analysis. While you can teach an algorithm the definitions of words, the way they're employed by humans doesn't always directly correspond to what's written in the Oxford English Dictionary. Think sarcasm, or words with a meaning that has flipped over generations, like "wicked", "sick", "cool", and so on. This may not seem like a big deal, but if you're feeding an algorithm customer service emails or call transcripts, then an angry customer can be mislabeled as satisfied, or vice versa.

"When it's in the case of ERP, CRM or healthcare, there's likely not a lot of nuance to your data," explains Mark

Beccue, a principal analyst at Omdia, who is specializing in AI. "But when you get into the intricacies of customer service transactions, that's when it gets really hard.

"As humans, we look at the way things are said in terms of tone. Is he or she talking with urgency? Are they loud?" Much of this is lost when voice is turned to text – as even humans sometimes find, when trying to decode if someone is angry with them by the curtness of a text message. Of course, AI can be trained on customer service calls and video to understand sound and visual cues, but it's far more complex than 'simple' textual analysis.

Powering up

The potential solutions to all of these problems may be varied, but they all have one thing in common: computational power.

"On the hardware side, edge computing and custom hardware, FPGAs (field-programmable gate arrays), is something I find really fascinating," Ghoshdastidar says. "Today, customers come to us and say 'I want you to do AI magic with the minimal hardware that's available to me' which is not easy. AI is resource-hungry."

The promise of FPGAs is that they allow a chip to be customized for specific business needs. "With this light hardware, I can do lots of heavier work," she explains. "The way we are advancing, in leaps and bounds, from the Von Neumann model of computing to neuromorphic computing, we are really seeing more and more miracles every day."

The investment in hardware, in other words, is something that Ghoshdastidar is keeping a close eye on. "The kind of investments going on, and the development of collaborative systems, where humans and machines can work together in a seamless and effective way – that's something that I find very promising. It should give humans wings, right?"

We can but hope. The problems outlined above – and others facing the wider artificial intelligence research community – may seem insurmountable today, but if you want to put things in perspective, just take a step back and see how far we've come. With the progress we've made in the 70 years since the Turing Test was first proposed, who would really dare to suggest that AI is anywhere near its peak?



AI Business eBook Series



EdgeVerve Systems, a wholly owned subsidiary of Infosys, develops innovative software products and offers them on premise or as cloud-hosted business platforms. Our products help businesses develop deeper connections with stakeholders, power continuous innovation and accelerate growth in the digital world. We power our clients' growth in rapidly evolving areas like banking, interactive commerce, distributive trade, customer service and enterprise buying and help them navigate their digital journey with our AI enabled Automation and Business Solutions. At EdgeVerve, we are making constant strides towards transforming enterprises by providing AI enabled Business Applications, leveraging the Infosys Nia™ Platform with capabilities across the Automation continuum. Today EdgeVerve products are used by global corporations across financial services, insurance, retail and CPG, life sciences, manufacturing and telecom.

To know more, visit www.edgeverve.com
or write to contact@edgeverve.com

Infosys NIA

Infosys Nia is an enterprise grade AI platform that simplifies the AI journey to enable the enterprise's digital transformation. It digitizes and harvests intelligence from documents and data across the polycloud. Infosys Nia industrializes AI deployments, accelerates business outcomes and assures fastest time-to-value for AI implementation. Simplifying the consumption experience in the AI journey, it integrates with existing AI infrastructure and protects investments. Infosys Nia enhances the consumption of AI workloads through the end-to-end lifecycle. The platform truly speeds up the enterprise AI journey by empowering the citizen data scientists to build complex use-cases easily.

Nia DocAI is a document extraction, processing and comprehension platform that unlocks business value from enterprise data by extracting intelligence from enterprise documents, regardless of complexity or domain specificity.

Book a Demo

