

SHOULD YOU BE USING AI IN YOUR BUSINESS?

YAZMIN HOW, DIGITAL CONTENT MANAGER

WHITE PAPER

Artificial Intelligence (AI) is transforming every industry it touches from healthcare, to retail and advertising, finance, transport, education, agriculture and so many more. The purpose of AI? To take care of all the mundane tasks employees currently handle, freeing up their time to be more creative and perform the work that machines cannot do. Today, the rapidly advancing technology is used mostly by large enterprises through machine learning and predictive analytics. AI is not a technology of the future, it's happening now, and companies who fail to adopt it will get left behind.

This paper will explore the application of AI in business, delve into who should be employing these technologies, and hone in on the transformative impact AI is having on every industry with research contributions from leading minds in the field. Expert opinions from academics, industry leaders, researchers, CEOs, founders and many more are included to comment on the impact of artificial intelligence across multiple industries.

CONTENTS

1. Contributors.....p. 2 - 3
2. Introduction.....p. 4 - 5
3. The Landscape of AI.....p. 6 - 8
4. Progressions.....p. 9 - 11
5. Use Case & Benefits.....p. 12 - 15
6. Case Study 1: Daniel Golden, Arterys...p. 16 - 17
7. Case Study 2: Maggie Mhanna, Renault...p. 18 - 19
8. Where Are We Heading?.....p. 20
9. Learn more...p. 21
10. Glossary.....p.22
11. References & Additional Reading.....p. 23

1. CONTRIBUTORS

- **Jörg Bornschien**, Research Scientist, DeepMind
Specialist in unsupervised and semisupervised learning using deep architectures. Jörg Bornschien was chair and one of the founders of the german hackerspace "Das Labor" which was awarded in 2005 by the federal government for promoting STEM programs to prospective students.
- **Ben Chamberlain**, Senior Data Scientist, ASOS
Holds a Royal Commission for the Exhibition of 1851 Industrial Fellowship, which funds his PhD studies in statistical machine learning at Imperial College London. Ben Chamberlain has previously worked as a data scientist in the social media, defence and security industries.
- **Eli David**, Co-Founder and CTO, Deep Instinct
Specialist in deep learning and evolutionary computation with over thirty publications in leading AI journals. Developer of Falcon, a chess playing program based on genetic algorithms and Deep Learning. Eli David received Best Paper Award in 2008 Genetic & Evolutionary Computation Conference.
- **Daniel Golden**, Director of Machine Learning, Arterys
Uses machine learning to predict outcomes and disease characteristics in cancer patients. Daniel Golden founded a machine learning team at CellScope that used the field of deep learning to diagnose ear disease and streamline the process of recording ear exams at home.
- **Ian Goodfellow**, Senior Research Scientist, Google Brain
Inventor of generative adversarial networks, and lead author of the MIT Press textbook deep learning. Main focus of research includes generative models and security and privacy for machine learning. Ian Goodfellow has contributed to open source libraries including TensorFlow, Theano, and Pylearn2.
- **Ankur Handa**, Research Scientist, OpenAI
Dyson Research Fellow who focuses on using simulations to generate data for ML to do scene understanding with SceneNet RGB-D. Currently working on 3D scene understanding for robotics, reinforcement learning, and transfer learning.
- **Daphne Koller**, CCO, Calico
Co-Chair of the Board and Co-Founder of Coursera, the largest platform for massive open online courses (MOOCs). Previously, she was the Rajeev Motwani Professor of Computer Science at Stanford University, where she served on the faculty for 18 years. She is the author of over 200-refereed publications appearing in venues such as Science, Cell, and Nature Genetics.
- **Hugo Larochelle**, Research Scientist, Google Brain
Recipient of two Google Faculty Awards. Associate editor for the IEEE Transactions on Pattern Analysis and MI, on the editorial board of the Journal of Artificial Intelligence Research and program chair for the International Conference on Learning Representations.
- **Miao Lu**, Research Scientist, Yahoo Labs
Current work focuses on native/display/search ads recommendation and forecasting. Interdisciplinary background in statistics, machine learning and data mining, with wide applications in biomedical science and internet technology.

- Maggie Mhanna**, Data Scientist, Renault Digital
 Part-time University Professor at Leonardo da Vinci Engineering School. Maggie Mhanna's main focus is on the application of data science and ML with connected cars. Maggie earned a master of science in renewable energies from École Polytechnique, and an engineering degree in computing and communication from the Lebanese university.
- Ed Newton-Rex**, CEO, Jukedeck
 Learnt to code in order to start Jukedeck, which now comprises a team of 20 musicians and engineers. Named one of WIRED's Hottest European Startups and has won a number of competitions, including the Startup Battlefield at TechCrunch Disrupt and LeWeb in Paris, as well as a Cannes Innovation Lion.
- Kimberly Powell**, Senior Director of Business Development, NVIDIA
 Specialises in business development for Healthcare, she promotes GPU computing into the bio/life sciences and medical imaging fields. A graduate of Northeastern University, Powell holds a bachelor of science in electrical engineering with a concentration in computer engineering.
- Vijay Ramakrishnan**, Machine Learning Researcher, Cisco
 Specialises in developing AI and Natural Language Processing (NLP). Expert practitioner in developing NLP models and a leader in building deep domain conversational AI products. He has built AI assistants for fortune 500 companies at Mindmeld Inc before they were acquired by Cisco.
- Jasper Snoek**, Research Scientist, Google Brain
 Held postdoctoral fellowships at the University of Toronto, under Geoffrey Hinton and Ruslan Salakhutdinov, and at the Harvard Center for Research on Computation and Society, under Ryan Adams. Jasper co-founded the machine learning startup Whetlab, acquired by Twitter in 2015.
- Raquel Urtasun**, Head of Uber ATG Toronto, Uber
 Associate Professor at the University of Toronto, a Canada Research Chair in Machine Learning and Computer Vision and a co-founder of the Vector Institute for AI. She is a world leading expert in machine perception for self-driving cars. Research interests include ML, computer vision, robotics and remote sensing. Her lab was selected as an NVIDIA NVAIL lab.



2. INTRODUCTION

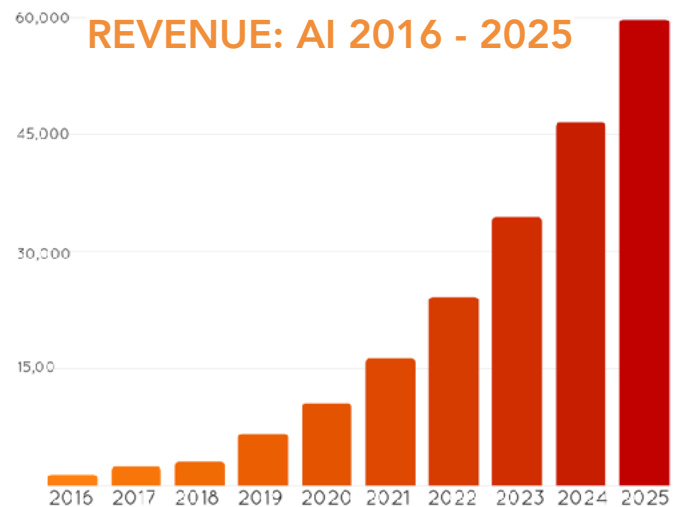
Artificial Intelligence is disrupting and transforming every industry it touches. From business operations and efficiency to innovative means of customer service, medical research breakthroughs, smarter transport systems and targeted advertising campaigns, it's an inescapable reality of today's world. Businesses unwilling to adopt AI will fall behind, and it's predicted that the revenue generated from both the direct and indirect application of AI software will grow from **\$1.38 billion** in 2016 to **\$59.75 billion** by 2025. (Tractica, 2017)

"As machines become smarter, consumers will expect flawless customer service around the clock, and by 2025 AI will drive 95% of all customer interactions, with consumers unable to differentiate bots from human workers via online chats as well as over the phone."

(Servion, 2017)

Thanks to the availability of huge amounts of data and increasingly intelligent algorithms, machines can **learn, speak, make informed decisions** and **carry out complex tasks** in an increasingly effective manner. Not only is this driving research breakthroughs, but implementation in industry is demonstrating the huge potential impact that real-world applications of AI can have on businesses across all industries from retail and advertising, to healthcare, sales and marketing, transport, travel and tourism amongst others.

PROJECTED WORLDWIDE REVENUE: AI 2016 - 2025



In the 80s when little progress was being made, three pioneers **Yoshua Bengio, Yann LeCun** and **Geoffrey Hinton** toiled away working on neural networks where other scientists had abandoned them due to lack of computational power. *'In the lean times when no one believed in neural nets, these [were] the people who really kept the torch burning and really inspired a lot of people.'* (Recode, 2015)

Since then AI has been applied in business, enhanced through research, and has had some astounding breakthroughs: DeepMind have mastered the Atari console and conquered the game of 'Go', and AI is now becoming superior to humans in several areas including object recognition and face detection, as well as working towards passing the Turing test. (DeepMind, 2013)

“

It is machine learning, and in particular neural networks, that seem right now to hold huge promise; but the history of artificial intelligence warns us not to assume we can accurately predict what will work, or when.

(Ed Newton-Rex, Jukedeck)

WHAT ACTUALLY IS AI?

Depending who you talk to, you usually find two definitions: one where AI aims to embed human intelligence into a machine, and another where AI aims at discovering possibly super-human levels of intelligence. If interpretability of the AI system is important, we might prefer an intelligence that's closer to the human's. But if we wish to design the best AI system that detects diseases in patients, we would be happy if it were better than a human doctor.

(Hugo Larochelle, Google Brain)

AI is the simulation of intelligence in computers: behaviour exhibited by non-biological systems that we would consider intelligent if exhibited by humans. A more recent approach, is 'machine learning' where the computer learns how to complete tasks by being exposed to large datasets.

(Ed Newton-Rex, Jukedeck)

What's allowing us to progress so quickly?

AI requires huge data sets and the coupling of 'really great science with amazing advances in technology has allowed us to collect data at that we've never had access to before', enabling models to learn more quickly.

(Jasper Snoek, Google Brain) The pace of current advancements wasn't foreseen, for example **Ankur Handa**, OpenAI didn't expect to see 'super-human performance on ImageNet within only three years of the first paper on CNNs (convolutional neural networks) from Geoffrey Hinton's group in Toronto.'

Whilst these progressions are rapid and impactful, consideration is necessary to identify whether your business in should be applying AI. Factors such as cost, available data, industry relevance and staffing are a consideration for businesses of all sizes as well as the likely ROI. Further chapters will provide solutions to these key points in order to identify the impact of AI in industry and whether you should be employing these technologies in your business.



Progress has been driven primarily by new ideas and insight rather than bigger datasets and faster computers.

(Jörg Bornschien, DeepMind)

Who is driving progressions in AI?

It's not just technology giants leading the AI race, but Universities, venture capitalists (VCs) and internal researchers. Research from institutions and industry experts opens doors for businesses to apply these models to their work, and AI specific VCs are assisting in breakthroughs through their funding.

VENTURE CAPITALISTS



khosla ventures



ACADEMIC INSTITUTIONS



INDUSTRY



3. LANDSCAPE



The current landscape of AI consists of four areas: models, training data, hardware and people. New models (such as Generative Adversarial Networks) are being used with great success; bigger datasets are available to train on; hardware improvements are speeding up training; and every success in the field leads to more people entering the industry.

(Ed Newton-Rex, Jukedeck)

Progress is consistent, yet to have a direct impact on society takes time and money from organisations implementing research progressions. There are limitations in the availability of data, the power of the computation, and the training period and intelligence of each model. Advancements in unsupervised learning are revolutionising business applications saving both time and money - between 2010 and 2014, global investments in AI tech grew from \$1.7bn to \$14.9bn (Merrill Lynch, Bank of America, 2016).

Although AI isn't a new concept, the journey to mainstream implementation has taken years, with recent decades seeing the most rapid advancements. This means that new advances become obsolete quickly thanks to our current culture of using arXiv and social media for dissemination. This 'hyperactivity' in AI can be disruptive and is forcing companies to reconsider the products they are designing (Hugo Larochelle, Google Brain).

Factors in the development of AI that have brought us today's models:



Researchers figured out how to use **graphics processing units (GPUs)** to speed up training of neural nets. These methods allowed us to scale up to larger

datasets & to iterate faster when doing research on smaller, academic-minded datasets.



The DL community developed much better open source libraries and tools for deep learning. For example, **Theano** was a breakthrough in how we can

implement DL models, which inspired the now widely popular **TensorFlow**.



ArXiv and **social media** have become the prime strategy for publicising research. This has made it possible to iterate and build on other researchers'

work at a much higher pace.

(Hugo Larochelle, Google Brain)

HOW CAN ADVANCEMENTS BE APPLIED IN BUSINESS?

ML is always held back by limitations in the amount of computation we can use. These advances are able to have very apparent impacts on real world problems, and **Ian Goodfellow** explained:

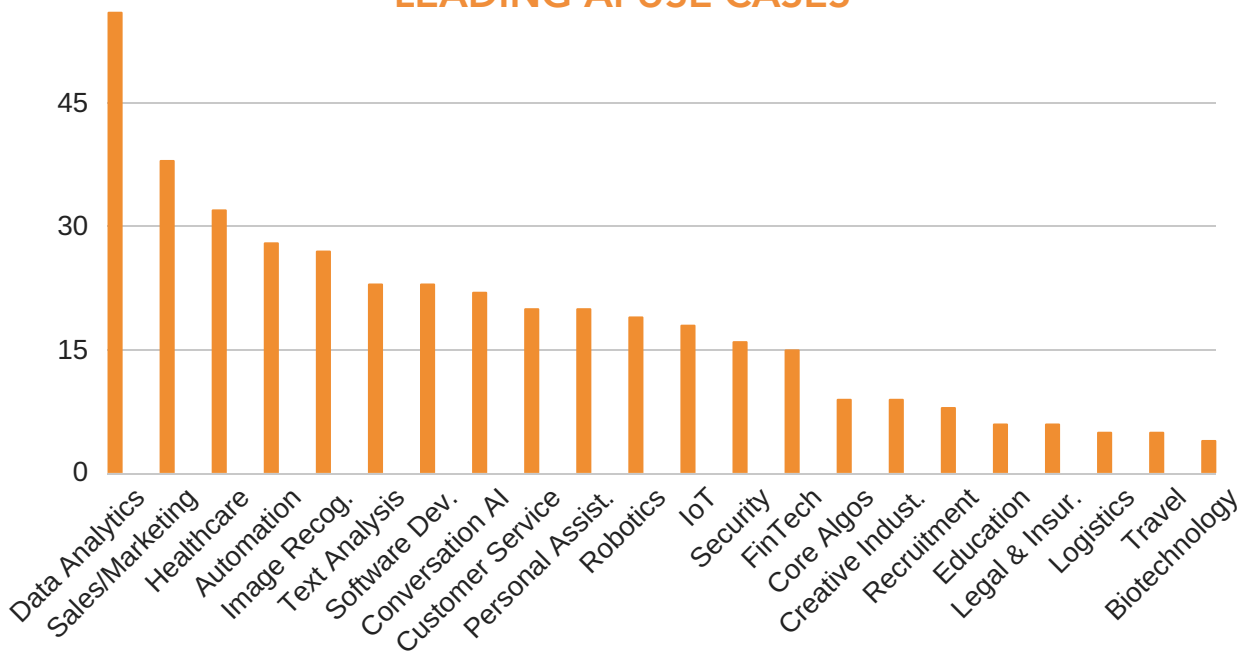
'the May 2017 announcement of the new generation of Google TPUs is huge. The new Google TPU helps bridge the gap between the amount of computation we can leverage in DL experiments and the amount of computation used in a biological nervous system. The previous generation was available only to Google engineers, but the new one will be available to Cloud customers, and researchers can apply to get access for free.'



Whilst we have seen a lot of progress in unsupervised learning, especially in generative modelling, we have yet to see it's full potential. It's impossible to attribute progress to a specific technique, but it's fair to say that progress has been driven by new ideas and insight rather than bigger datasets and faster computers.'

(Jörg Bornschien, DeepMind)

LEADING AI USE CASES




SOURCE: CITY.AI, 2017

As more companies are able to implement AI models in their businesses through the availability of cloud platforms (for instance from Google, Amazon, Microsoft), AI technology will continue to become more accessible to both industry and society at large. There is no need for any particular industry to get left behind, and industries across the board are being positively disrupted by AI advancements (**Hugo Larochelle**, Google Brain). Business intelligence tools are able to source, analyse, transform and report on data to provide valuable customer insights and allow businesses to invest their time and money in the right areas. Companies striving to build emotional connections with their users will see increased customer satisfaction through emotional AI, after all 'in real life situations, people are actually pretty bad at emotional

intelligence', causing us to end up in pointless fights, dismiss good arguments because they go against our biases, and judge people based on stereotypes. Once AI can be trained without bias it should be able to provide more rational responses than its human counterparts (Mikko Alasaarela, Inbot, Oct 2017)

Across multiple industries we see mind-boggling results due to deep learning in data preparation, speech recognition, text understanding, computer games, cybersecurity, etc. Deep learning has provided the greatest leap in performance in the history of AI (and arguably, in the history of computer science), and has rendered many traditional methods obsolete. As a result, within the next decade any company that will not heavily rely on deep learning will be left behind (**Eli David**, CTO, Deep Instinct).

AI ISN'T JUST FOR 'TECH' COMPANIES



20%

of business content will come from AI by 2018 ([Salesforce, 2017](#)).

57%

of consumers expect voice-activated smart assistants to have a major or moderate impact on their daily lives by 2020. ([Salesforce, 2017](#)).



60%

of AI platforms will be hosted by Amazon, Alphabet, IBM & Microsoft by the start of the next decade due to their status in cloud computing ([IDC FutureScape, 2017](#)).

75%

of developer teams will use AI technology in one or more business applications or services by 2018 ([IDC FutureScape, 2017](#))



80%



of customer relationship management will be done without human interaction by 2020 ([Salesforce, 2017](#))

DEEP DIVE: WHAT ARE GPUS?

Graphics processing units (GPUs) accelerate deep learning applications by allowing computers to offload intense portions of applications to the GPU whilst the remainder of the code runs on the CPU. Pioneered in 2007 by NVIDIA, GPU accelerators now power energy-efficient data centres in government labs, universities, enterprises, and small-and-medium businesses around the world. ([NVIDIA, 2017](#))

WHAT ARE TPUS?

Tensor processing units (TPUs) accelerate the neural network computations behind the scenes delivering 15–30X higher performance and 30–80X higher performance-per-watt than contemporary CPUs and GPUs. These advantages help many of Google's services run state-of-the-art neural networks at scale and at an affordable cost. ([Google, 2017](#))

4. PROGRESSIONS



"There's been such rapid progress in recent years. Since 2012 on the image net, all of that progress is essentially because of the advances we've made to the learning algorithms - it's not only because of the computational resources, but also the research progress"

(**Aaron Courville**, Assistant Professor, University of Montreal, Oct 2017)

There's a lot of discussion about AI '**stealing jobs**', but this isn't the first time. During the industrial revolution factory workers were put out of jobs by machines; fast forward to the internet and industries were disrupted. However, this wasn't the end for industry: factory workers learned to operate machines, journalists used the internet as a resource rather than a hindrance. **For every job that's destroyed, there will be countless new opportunities created.** (Forbes, 2017)

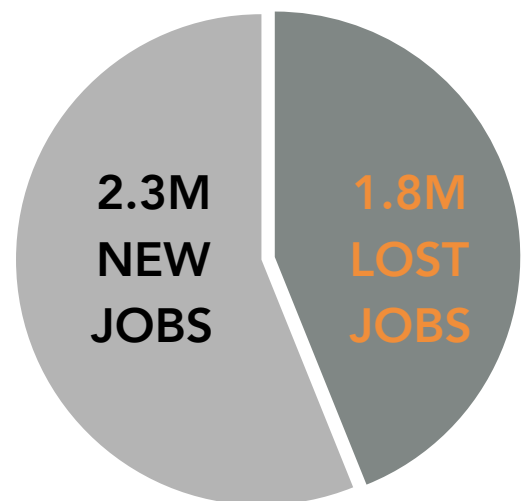


Hugo Larochelle explained, whilst we can't know for sure the future of the AI landscape and the economy, *historically there are examples of tech developments that aimed at "replacing people" yet growth in employment for that industry was observed*

(the advent of ATMs is an often cited one).

For some industries which rely heavily on labor, AI might be better than human beings. Since AI will never be tired and hardly make mistakes, this will create job openings in the development and automation of the robots and machines. Industries which require more human intelligence and emotion however, are unlikely to be completely transformed by AI, but will see the introduction of AI assistants to both save time and helps us do more intelligent work. (**Miao Lu**, Research Scientist, Yahoo Labs).

AI is creating new jobs in 4 out of 5 companies



SOURCE: CAPGEMINI SURVEY, 2017

AI IS HELPING IMPROVE EFFICIENCY, CUT COSTS & OPTIMISE BUSINESS

AI allows people to spend time on the more thoughtful and creative aspects of their jobs rather than wasting time on menial tasks. For instance, smart customer relationship managers (CRMs) can be used to automate data input, HR employees can use smart recruiting software to sort applicants more efficiently, and marketing and sales teams can send personalised and targeted campaigns. Businesses of all sizes can be reinvigorated by these software processes allowing employees and management to invest their time in growing the business. (Salesforce, 2016)



We made AI computing very fast, and we know speed means productivity.

***Productivity ultimately leads to performance.** Beyond speed, we must make sure it works well and is widely accessible.*

(Kimberly Powell, NVIDIA)

AI is now far more accessible, even to those outside the world of computer science and AI in general, and as Hugo Larochelle highlighted, the increase in AI cloud platforms is improving the accessibility. It's now important for Universities and Colleges to embrace the progress of AI and move away from treating it as a narrow topic in computer science meant for graduate study.

Historically, AI was developed by large companies since **a lot of training data was required to reap the rewards of AI and ML.** For the company to acquire this data meant that the business needed to either have a pre-existing and mature product with a lot of user traction, or for the company to have high financial backing to buy the data at an exchange. Today, the cost of acquiring training data is significantly reducing with platforms like scale.api and Amazon Mechanical Turk matching human annotators to startups at cost efficient prices. Moreover, a lot of Platform as a Service systems that directly serve machine learning models to input datasets like Amazon ML and Azure ML can help bootstrap startups in the AI/ML space. **(Vijay Ramakrishnan, Cisco)**




AI allows researchers to focus on the problems at hand instead of spending time creating new tools to approach new problems.


(Ian Goodfellow, Google Brain)




Feel that AI threatens their jobs today.



Worry they won't have time/money to retrain in AI.



Would be comfortable working with robots in the next 3 years.




Think younger generations need new skills to succeed as AI develops.

SOURCE: ECONOMIC TIMES SURVEY, 2017

As consumers demand more **personalisation** and **customisation** in their lives, personal assistants that cater for a particular brand and products catered specific to their likings will increasingly be expected from the consumer. (Kimberly Powell, NVIDIA). Businesses have to reach such expectations and manually doing this will stop being scalable after a certain point.

Consider the service your company provides and the challenges you're currently experiencing, whether they lie in ROI, organisation, efficiency, accuracy, customer service or any other aspect of the business. Consider the impact an AI could have on not only your day to day activity, but on your company's success as a whole.



SOURCE: BOOMTRAIN, 2017

RESEARCH PROGRESSION TESTIMONIAL (Eli David, Deep Instinct)

In 2004 I led a project that created a critical real-world system with immense complexity, entirely based on neural networks and evolutionary computation. Back then, merely mentioning neural networks resulted in raised eyebrows and bewildered looks. And there were good reasons for that: we could train rather shallow neural networks, the training was excruciatingly time consuming, and for most applications the results were inferior to other traditional machine learning methods which were much faster and more convenient to deploy.

During the past few years we have witnessed two revolutions: The first, much improved neural networks methods that allow us to train "deep" neural networks (and hence the name "deep learning"), and the second, the use of GPUs which allow us to train neural networks which are by orders of magnitude larger. And these two breakthroughs are responsible for the current ongoing deep learning revolution.

5. USE CASE & BENEFITS

“

Five years ago I was talking to a friend who was leading one of the leading computer vision groups in the world. When discussing deep learning he referred to it as "another transient hype". Today, his entire research group is focused only on deep learning, obtaining amazing breakthroughs, and barely using any of the traditional image processing methods which have been used in previous decades. I expect every leading business in every industry to heavily rely on deep learning in the upcoming years (or otherwise be completely left behind).

(Eli David, Deep Instinct)

Whilst the obvious contender for implementing AI is 'tech companies', cross-industry enterprises are seeing benefits incomparable to previously used models when implemented successfully. Smaller companies can be deterred from investing in an long term AI strategy due to the vast amounts of labelled data that's required. Many businesses don't have access to this, but that isn't to say that it isn't possible or advisable to build a successful strategy.

Companies currently benefitting from AI are the ones which already have an existing store of archived data, like hospitals or businesses where the cost of acquiring incremental data is low, like internet companies. However, the time investment required to collect training data or develop ML models is decreasing. Additionally, cloud infrastructure services like Amazon AWS and Google Cloud ML reduces the need to acquire expensive capital goods upfront, thereby reducing the friction of small/medium size companies entering into this space. (Vijay Ramakrishnan, Cisco)

"We are already seeing surprising success stories of unusual AI applications. One I particularly like is that of Makoto Koike, an engineer who used TensorFlow to design a cucumber sorting system for his parents' farm in Japan."
(Hugo Larochelle, Google Brain)



1.

BUSINESS OPTIMISATION

Tasks that have previously relied on manual labour from employees such as customer support, data management, supply chain management, and marketing strategy decisions and automation are areas predicted to be the most affected by AI in the next five years ([Forbes, 2017](#)). There are very few companies who don't rely on employees in any of these areas, and paying large teams is costly, so maximising their efforts by alleviating time consuming tasks that could be managed by an AI model is an obvious direction for businesses to take. Introducing a technology that can provide more insights into making an informed decision like a human can is sure to improve not only ROI, but optimise efficiency and allow staff to excel in areas that require human skill (**Ankur Handa**, OpenAI).



75% of over 200 businesses interviewed said AI will be "actively implemented" in their companies within the next three years

79% of executives surveyed around the globe believe AI will make their job easier and more efficient.

SOURCE: ECONOMIST INTELLIGENCE, 2017

2.

HEALTHCARE

The adoption of AI in healthcare is on the rise, helping not only with diagnoses but also in the production of new drugs and research advancements. **Deep learning methods are already enabling analysis that outperforms experts in applications such as the detection of diabetic eye disease and cancerous tumours.** These progressions will not replace doctors and medical experts, but will allow them to achieve higher accuracies with more time to focus on the hardest cases and with a strong second opinion to help minimise human error (**Jasper Snoek**, Google Brain). For example, in the U.S, there is one radiologist for every 10,000 people, and in India there is one for every 100,000 people. With AI, radiologists will be more effective and will therefore have more time to focus on those cases that are more difficult. (**Kimberly Powell**, NVIDIA). Jasper Snoek expects this to spread across medical domains to, for example, cardiology where algorithms could automatically analyse EKGs taken at home.



DeepMind is working and researching with the **NHS** to build a ML system to recognise impending blindness. ([The Guardian, 2016](#))



Researchers have created '**dermatologist level**' classification of skin cancer using **neural nets**. ([Springer Nature, 2017](#))



IBM Watson carried out cancer research where **99%** of the treatment recommendations were coherent with the physician decisions. ([SVN, 2017](#))

3.

TRANSPORT

Soon, autonomous vehicles won't just be delivering pizza. Their development will see reduced congestion, improved mobility for those who have difficulty getting around today, greener cities, and increased access to public transit.

Raquel Urtasun, Head of Uber Advanced Technology Group (ATG) is using ML techniques to make their self-drive cars more intelligent and safer than human drivers. Raquel explains that when she thinks about cities today she 'can't help but see all of the parking spaces and garages dedicated to our personal vehicles.' A future where we can begin to transform those parking structures into green space and new businesses is not too far away. Additionally, the huge popularity of Amazon has created a massive opportunity for autonomous vehicles, as the already prevalent shortage of last-mile drivers will continue to increase based on current consumer trends (**Kimberly Powell**, NVIDIA).

By 2035, the UK autonomous vehicle market is estimated to be worth **£28bn** capturing **3%** of the **£907bn** global market.



151K

Direct jobs



£71.6bn

Annual
turnover

1.56m

Cars built



1.2m

Exported



>800k

Automotive

SOURCE: UK GOV MARKET FORECAST, 2017

4.

RETAIL

“ We use ML to identify and distinguish between high and low value customers to optimise market spend and minimise exposure to unprofitable customers. A large percentage of customers churn and have a 0% CTLV, whilst some will spend millions each year on ASOS. It's important to identify which users are which to drive sales and provide the best UX.

(**Ben Chamberlain**, ASOS)

Retail sales are dropping as the cost of living increases. This is problematic for retailers who've had to turn to tech to encourage consumers to jump back on the spending wagon. Using machine learning, retailers can capitalise on market data to understand and anticipate consumer behaviours and trends (RTInsights, 2017). ASOS are using machine learning to predict customer lifetime value. In an ideal world, ASOS would 'know every action a customer will make for the rest of time, but [they] can't do that'. Distinguishing between high and low value customers allows companies to optimise market spend and minimise advertising to unprofitable customers. Retail has massive amounts of consumer data so is able to use AI to show real-time behaviours in the customer base. This further benefits the industry by allowing companies to implement better inventory management systems (Ben Chamberlain, ASOS).

The examples above demonstrate the positive transformations of applying AI in a variety of sectors, **but are there any negative ramifications?** Data scientists or an AI officer can be expensive hires, but there are an abundance of online training courses available such as Andrew Ng's Coursea as well as platforms like Kaggle. Hugo Larochelle explained that they have seen examples of people without formal training in AI who are able to produce high quality solutions to problems using AI. This suggests that the barrier of entry to using AI is lowering, and new AI-enabled systems in new domains should continue to be observed. There are some industries that will be required to tread more carefully than others in their implementation of AI, for example law enforcements agencies that have an enormous stake on public life that employ AI need to tread carefully since the outputs of the such systems could over-fit on past data, so racial and gender bias could be learned.

In this case interpretable models are preferred to black box models. Additionally, deep learning systems can be easily tricked using carefully crafted input data. MindMeld's Vijay Ramarkishnan explains that in these circumstances, systems will need ample supervision and common-sense before being the first responders to control important decisions like physical gate-keeping of sensitive areas. It's important to remember however, that fictional representations of AI tend to paint a dystopian picture. We need to be very careful with AI development, as it could certainly bring about negative scenarios; but we shouldn't assume that things will go badly, in any industry. AI has the potential to do enormous good, probably more than any technology we've ever created before (**Ed Newton-Rex**, Jukedeck).

TAKING AI FURTHER: APPLYING DEEP LEARNING ALGORITHMS

Whilst AI is now readily available for implementation cross-industry, there are still challenges in the commercial applications of deep learning. Despite the availability of multiple research oriented open-source deep learning frameworks, they are not efficient enough for application in many domains (e.g., cybersecurity); their incorporation in the final product is complicated; and even where applicable, highly skilled and scarce deep learning experts are required to employ them. As a result, today most of the real-world applications of deep learning in industry are in computer vision, because they mainly rely on convolutional neural networks only, for which the barrier of entry is lower.

Eli David from Deep Instinct shared his expertise on overcoming this obstacle: *I expect that in the next few years we will see deep learning being employed effectively in additional industries in real-world solutions. We will have more deep learning frameworks which focus on efficiency in real-world rather than being geared towards research, as is the case with currently available deep learning frameworks, and they will lower the barrier of entry for adoption in the industry.*

6. CASE STUDY:

DANIEL GOLDEN, DIRECTOR OF MACHINE LEARNING, ARTERYS

Artificial Intelligence has the capacity to transform all aspects of healthcare from the research and production of new drugs and treatments, to the way in which patients are cared for.

Cloud based company, Arterys, has created an **intelligent medical imaging platform** that supports advanced AI tools and ultra-fast image processing. Back in 2011 when Arterys was founded, most medical imaging happened on hospital premises with limited computing power, with measurements taken manually and automated tools minimal. Out of this, the vision to advance medical imaging via cloud computation and advanced analytics was born.

Arterys was the first clinical cloud-based deep learning in healthcare model to receive FDA approval to be used in a clinical setting, a huge step for AI in healthcare. Their aim of creating *'A world where clinical care is data-driven, intelligent and patient focused'* comes to life through their *'deep learning engine with unlimited computation to translate deep learning models into clinical products.'*

MEDICAL IMAGING PROBLEMS:

Medical images contain a wealth of information that isn't being leveraged to optimise patient care. Through images, it's difficult to precisely track disease and even more so to predict the course a disease will take or how it will respond to specific treatments. Several challenges in medical imaging are currently preventing data-driven diagnoses:



Increasing Productivity Demands

Complex scans with **more images** at a **higher resolution** are becoming the norm. Scanning volume has increased significantly in the past **10 years**, overburdening radiologists and causing burnout.



Challenges in Quantifying Diagnostic Imaging

Quantification of anatomy is tedious, and clinicians don't have time to measure precisely. While volumetric measurement is ideal, disease is often measured with few linear dimensions which can be **inconsistent** between readers.



High Error Rates

The retrospective error rate among radiologic examinations is around **30%** with real time errors in daily radiology practice averaging **3-5%**. These errors lead to **delays in treatment, unnecessary biopsies** and **mis-informed treatment** decisions.

These problems deliver real-world ramifications to healthcare professionals and patients alike. **The lack of efficiency means that delays and errors in screening, diagnosis and monitoring are common and lead to poor outcomes at higher costs.** Delays can come from either the time it takes for an image to be interpreted, or the time it takes for physicians to detect disease progression or assess the effectiveness of a treatment. **Arterys identified this issue and are leveraging their platform to accelerate these insights to provide more timely treatments as well as stopping and altering treatments that do not work.**

How can AI Help?

Whilst the system augments the radiologist in several aspects, **it does not aim to replace the human**, but enhance it in the following ways:



Expediting the tedious work surrounding setup, interpretation, and reporting on image findings.



Improving accuracy and consistency when measuring anatomy so the radiologist can better track changes in patients.



Helping the community standardise how disease is measured and categorised, so that there is more consistency across practices.

This will allow for real-world study of outcomes of similar patients and optimisation of treatment for individual patients.

Arterys have faced a good amount of skepticism surrounding AI in the clinical community, however **having allowed clinicians to try the system for themselves and compare the results has proven a convincing demonstration of the success and positive impact of AI**. The clinicians are then able to experience the benefits of the automation and assess its accuracy. The AI-enabled platform simultaneously enables the clinician to edit the measurements and assessments it automates, allowing them to remain in full control of the interpretation.

What's Stopping Mainstream Implementation, and What's Next?

AI requires very large datasets, and due to the physical recording of data that were popular until recently, these high-quality datasets can be hard to come by. Data found in a single institution can be biased to how clinicians practice there, so real leaps will occur when clinicians from different institutions collaborate in creating datasets that represent the population in a given disease.

As more and more clinicians and researchers embrace AI and see the emergence of new patterns and biomarkers thanks to their research, more people are going to be keen to employ the technologies.

What does Arterys implementation of AI teach us about cross-industry applications?

- Costs can be cut whilst improving efficiency of data collection
- Image analysis does not need to be time consuming, and can become more accurate
- Jobs do not need to be lost, the employee simply becomes more effective working alongside the intelligent system

7. CASE STUDY:

MAGGIE MHANNA, DATA SCIENTIST, RENAULT DIGITAL

Renault Digital is 100% Renault-owned subsidiary that started operations on January 1st, 2017 and **aims at digitalising Renault's core business for its employees, partners and clients worldwide**. Renault Digital's mission is to **deliver digital projects for all of the company's business lines**, whether for the upstream (in factories, for the design of vehicles) or for the downstream (points of sale, e-commerce, after-sales).

"We have based our approach on an inspiring start-up ecosystem. One of the most important core competence at Renault Digital is Data Science and AI. Driverless Cars are the main goal. In this project (AD/ADAS), the focus is on offering one the widest range of advanced driver assistance systems (ADAS) on the market, enhancing vehicle safety and in some cases operating without human intervention. These ADAS form the starting point for autonomous vehicles, capable of instigating decisions on behalf of the driver."


AI FOR A POSITIVE IMPACT

- **ML and data science for predictive maintenance:**



Predicting what will fail and when promises cost savings over routine or time-based preventive maintenance because tasks are performed only when warranted.

- **Accurately incorporating analysis of customer feedback:**



Helps in building vehicle and sub-systems performance for guiding future product design. This also helps in detecting failure patterns for establishing a relationship between the failure and causes of failure. Take an example of a cause of failure in several operations in the car which happens to be associated with region-specific issues such as inferior fuel quality, climatic conditions, road infrastructure, etc.

BRINGING CONNECTED CARS TO THE ROADS



The product lifecycle of a vehicle from inception through engineering design and manufacture to service and disposal is 5 years. New features, such as system upgrades and applications are provided regularly for smartphones; this can't be said for connected vehicles. AI models must be optimised to speed up this process and allow connected cars to run sophisticated applications.



As well as consuming power and battery, **connected cars also consume bandwidth**. Here, we need to talk about mobile operators: **passing data from the vehicle to the network and vice versa will represent a non-negligible amount of data**. Operators need to work with car manufacturers to make the transition seamless.

Obstacles in the Roll-Out of Connected Cars:

Data & Privacy

- The **privacy** of the individuals whose data is being collected and analyzed is **at risk daily**. Data is produced **via navigation, multimedia and other applications**.
- It provides information on the driver's habits, which can be transmitted to the car manufacturer, as well as to third parties. The **EU's General Data Protection Regulation (GDPR)** will come into force next year with tougher punishments covering a larger scope.

Obstacles in the Roll-Out of Autonomous Cars:

Reliability

- The vehicle must be able to perceive its environment perfectly, to make the right decision and carry out this action safely.
- In the real world, roads are a highly dynamic and varied environment, therefore there are millions of use cases the vehicle must learn.

Artificial Intuition

- Behaviours of other users can sometimes escape AI, humans are able to behave intuitively where computers are not.
- Humans take subtle signs for granted: the body language of a traffic controller, or a cyclist making eye contact. Machines will be able to learn these interactions with AI, but this may take infinite hours of driving and testing.

Legal Framework

- The **Vienna Convention on road traffic**, (the basis of European road codes), must evolve - it stipulates that the driver must keep control of the vehicle at all times.
- Most manufacturers are working with authorities on these changes, which can be assumed to be completed before 2020.

HOW RENAULT DIGITAL ARE USING AI IN OTHER PROJECTS



Demand forecasting



Optimising manufacturing



Distribution and the supply chain processes and networks



Finding correlations between incidents occurring at different phases of a project by identifying existing questions having the same context and linking them using NLP.



Optimising On-board diagnostics (OBD) by grouping diagnostic trouble codes and proposing them as labeled patterns to technicians with the corresponding solutions, **saving them both time and effort**.

How is Renault Digital's implementation of AI relevant?

The structure of their AI teams isn't that of a giant corporation, and can be replicated by businesses of varying sizes:


The projects development methodology used isn't waterfall but rather an agile approach relying on small teams. The teams self-organize and run on short cycles, with a strong link with the experts of the field and systematic feedback to users to validate the concept and quickly derive value.

8. WHERE ARE WE HEADING?



Contrary to what people think, scientific progress is slow and continuous, but social and economic impact can be disruptive - there's no doubt that we will reach human level AI but we don't know how long it's going to take. We need to make sure that AI and society will bloom for the benefit of all.

(Yoshua Bengio, Université De Montreal, RE•WORK Deep Learning Summit 2017)



Companies will dedicate up to **20%** of their workforce to neural nets, and **85%** of customer interactions will be managed by AI.

2020



57% of consumers expect voice activated smart assistants to have a major or moderate impact on their daily lives. (Salesforce, 2017)

2025



Revenue generated from the direct and indirect application of AI software will grow from **\$1.38 billion** in 2016 to **\$59.75 billion** by 2025. This represents a significant growth curve for the forecast period with a compound annual growth rate (CAGR) of **52%**. (Tractica, 2017)

2030



38% of U.S. jobs and **30%** of U.K. jobs could be automated with the industries leading being transportation & storage (56%), manufacturing (46%) wholesale and retail (44%) and healthcare (17%). (Tractica, 2017)

IN THE NEXT 5 YEARS: WHAT THEY SAY

Eli David, Deep Instinct:

Whilst we may face an initial wave of job loss, many new jobs will be created, which instead of competing with AI, will work together to achieve complex tasks not feasible today.

Raquel Urtasun, Uber ATG:

Thanks to AI, I see a future with reduced congestion, improved mobility for those who have difficulty getting around today, greener cities, and increased access to public transit.

Ankur Handa, OpenAI:

I see healthcare, public sector and government being affected positively as these are the places where most often important decisions are made.

Jasper Snoek, Google Brain:

New techniques for merging disparate data sources and NLP will allow us to analyze medical records, discover symptoms and predict medical outcomes .

Jörg Bornschien, DeepMind:

We will see progress related to rapid acquisition of knowledge: Few shot learning for example, where a generative or discriminative model has to generalize from only a few positive examples.

Kimberly Powell, NVIDIA:

More enterprises will implement AI in their own industry by crafting ideas to service each unique customer with AI. With an increased accessibility, everyone in every industry will be allowed to innovate with AI for themselves.

BEYOND THE NEXT FIVE YEARS...

“Research in *real* intelligence (humans and other animals) strongly suggests that there is a correlation between the number of neurons in the brain and the overall intelligence. This seems to hold true for artificial neural networks as well, as one could argue that our current state-of-the-art deep learning models are very similar to the neural networks from the 1990s, with the main difference that they contain about a million times more synapses (connections). Extrapolating from this trend, I predict that within the next several decades as the size of our DL models grows exponentially, and we would find better and better neural network architectures and parameters, we will get close to or surpass human-level. My guess is that this will happen in the lifetime of most of us.” (Eli David, Deep Instinct)

9. LEARN MORE

ABOUT

RE•WORK is an all-female run events organising company that brings together breakthrough technology, cutting-edge science and entrepreneurship shaping the future of business and society. We create and organise globally renowned summits, workshops and dinners, bringing together the brightest minds in AI from both industry and academia. At each RE•WORK event, we combine the latest technological innovation with real-world applications and practical case studies. Learn from global pioneers and industry experts, and network with CEOs, CTOs, data scientists, engineers and researchers disrupting their industries with AI. We also provide an analysis on current trends and innovations, through podcasts, white papers and video interviews.

Join us to learn more about applying AI in your business at our upcoming summits. [View our calendar of events here.](#)

ADDITIONAL READING FROM RE•WORK

- [Can AI Transform the Way We Buy and Sell?](#)
- [Yariv Adnan, Google Assistant: Are We Becoming Reliant on AI?](#)
- [Is a Chief AI Officer Needed to Drive an AI Strategy?](#)

RE•WORK PODCAST:

WOMEN IN AI

Women in AI is a biweekly podcast from RE•WORK, meeting with leading female minds in AI, Deep Learning and Machine Learning. We will speak to CEOs, CTOs, Data Scientists, Engineers, Researchers and Industry Professionals to learn about their cutting edge work and advances, as well as their impact on AI and their place in the industry. Guests include:

- **Antonia Creswell**, PhD Candidate, Imperial College
- **Elena Kochkina**, PhD Candidate, Warwick University
- **Ira Kemelmacher**, Research Scientist, Facebook

10. GLOSSARY

ArXiv - An open-source repository of electronic pre-prints of scientific papers

AD - Advanced driver

ADAS - Advanced driver-assistance systems

CPU - The central processing unit is the electronic circuitry within a computer

Deep Learning - Area of machine learning research which has been introduced with the objective of moving machine learning closer to one of its original goals of AI.

GAN - Generative adversarial networks are a class of AI algorithms used in unsupervised learning introduced by Ian Goodfellow in 2014

GPU - Graphics processing units accelerate deep learning applications by allowing computers to offload portions of the application to the GPU whilst the remainder runs on the CPU

Machine Learning - An application of AI that provides systems the ability to automatically learn and improve from experience without being explicitly programmed.

ROI - Return on investment

TPU - Tensor processing units are an application specific integrated circuit developed by Google especially for machine learning

Theano - Numerical computation library for Python

Tensorflow - Open-source software library for data flow programming across a range of tasks.

REFERENCES & ADDITIONAL READING

- 'The RIC Report', Bank of America, *Merrill Lynch*, ([12 Jan 2016](#) [View our calendar of events here.](#))
- 'Is Your Business Ready for Artificial Intelligence?', BCG Henderson Institute, *Philipp Gerbert, Martin Reeves, Sebastian Steinhäuser and Patrick Ruwolt*, ([6 Sept 2017](#))
- 'Reshaping Business With Artificial Intelligence', MITSloan, Sam Ransbotham, David Kiron, Philipp Gerbert and Martin Reeves, ([6 Sept 2017](#))
- 'Playing Atari with Deep Reinforcement Learning', DeepMind Technologies, Volodymyr Mnih, Koray Kavukcuoglu, David Silver & others, ([19 Dec 2013](#))
- 'Voice and Facial Recognition Could Help AI Surpass Humans in Emotional Intelligence', Venture Beat, *Mikko Alasaela, Inbot*, ([Oct 15 2017](#))
- 'How AI is Already Changing Your Life Backed By 6 Incredible Stats', Salesforce, *Kim Honjo*, ([1 Mar 2017](#))
- 'IDC FutureScape: Worldwide IT Industry 2017 Predictions' IDC FutureScape, *Frank Gens & others*, ([Nov 2016](#))
- 'What is GPU-Accelerated Computing?', NVIDIA, ([2017](#))
- 'An in-depth look at Google's first Tensor Processing Unit (TPU)', Google Cloud Platform, *Kaz Sato*, ([12 May 2017](#))
- Video: 'Visual Reasoning via Feature-wise Linear Modulation', RE•WORK Deep Learning Summit Montreal, *Aaron Courville*, University of Montreal, ([11 Oct 2017](#))
- 'The Paradox of Robots Taking All Our Jobs', Forbes, *Adam Oximek*, ([28 Oct 2017](#))
- 'Artificial Intelligence - Where and How to Invest', Capgemini, ([2017](#))
- 'Artificial Intelligence not to kill jobs yet: Survey', Economic Times, ([15 Nov 2017](#))
- 'Website Personalization Powered by Artificial Intelligence | Stats and Facts', Boomtrain, *Nitin Deshdeep*, ([14 Feb 2017](#))
- 'How Artificial Intelligence is Revolutionizing Business in 2017', Forbes, *Louis Columbus*, ([10 Sept 2017](#))
- 'Artificial Intelligence in Healthcare: Past, Present and Future', SVN, *Fei Jiang, Hui Zhi, Yi Dong & others*, ([21 June 2017](#))
- 'Google DeepMind pairs with NHS to use machine learning to fight blindness'. The Guardian, *Alex Hen*, ([5 July, 2016](#))
- 'Artificial Intelligence Market Forecasts', Tractica, ([Q2, 2017](#))