60 LEADERS ON

ARTIFICIAL INTELLIGENCE

60 THOUGHT LEADERS ANSWER 17 QUESTIONS ABOUT ARTIFICIAL INTELLIGENCE

60LEADERS.COM

2022 EDITION

George Krasadakis
Dublin, Ireland

Contributing Authors: Abhishek Gupta, Agnis Stibe, Alex Guilbault, Alex Wang, Alf Rehn, Anand Ranganathan, Angeliki Dedopoulou, Anthony Mills, Arathi Sethumadhavan, Aruna Kolluru, Aruna Pattam, Badr Boussabat, Boyka Simeonova, Christian Guttmann, Cynthia Rudin, David Pereira, Dima Turchyn, Edward Pyzer-Knapp, Eleanor Watson, Emma Duckworth, Enrico Panai, Eva Agapaki, George Panou, Golestan Radwan, Gertjan Boulet, Harry Mamangakis, Himanshi Allahabadi, Jair Ribeiro, Jaroslav Bláha, Jennifer Victoria Scurrell, João Azevedo Abreu, John Thompson, Jon Skirnir Agustsson, Jordi Guitart, Joseph Yun, Jyotirmay Gadewadikar, Kashyap Kompella, Luca Sambucci, Marily Nika, Mathew Hughes, Mayte Hidalgo, Michael Wu, Mike Tamir, Nazar Zaki, Netanel Eliav, Prakash Natarajan, Professor Andy Pardoe, Professor Eleni Mangina, Raj Madan, Richard Benjamins, Samiran Ghosh, Sherin Mathew, Siddha Ganju, Simon Greenman, Steve Shwartz, Sunil Kumar, Therése Svensson, Tirthajyoti Sarkar, Veena Calambur, Yannis Kalfoglou, Yasen Dimitrov

Concept, Production and Editing: George Krasadakis
Production assistance: Lucy Woodhead
Original Visual Design by Coy Chen

First Edition 2022.
Contact: https://www.60leaders.com/contact - g.krasadakis@gmail.com
https://www.60leaders.com/

PREFACE

GEORGE KRASADAKIS

PRODUCER & EDITOR OF THE '60 LEADERS' SERIES • IRELAND

Artificial Intelligence is probably one of the most popular and hyped technology terms of our time. But, beyond the media hype, there is a real revolution happening. To understand the potential, just consider the recent advances in domains like Computer Vision or Natural Language Processing and their applications.

The ability of a computer to 'see' is an astonishing achievement. AI-powered systems can 'understand' the context of an image or a video at an impressive level of detail. They can identify a growing set of entities — such as persons, named individuals, cars, buildings, and eventually any object — with increasing levels of accuracy. For example, given an image or video, algorithms can estimate properties such as the number of persons in the picture, their gender, age, their emotional state, or even the occasion and the dynamics. Computer Vision is making huge steps, with massive applications in autonomous cars, navigation, robotics, medical diagnosis, and more.

A chat with one of the popular digital assistants like Alexa or Siri is sufficient to demonstrate the advances of Natural Language Processing technologies. Digital assistants are becoming more intelligent, contextual, and proactive – they continuously learn by leveraging every single interaction with the user. At some point soon, digital assistants will become autonomous by seamlessly leveraging deep knowledge about the user, signals from the user's environment, and global trends and dynamics. Your digital assistant will respond naturally, in a conversational mode, and possibly with a style, attitude, and humor that matches your personality and your current mood.

AI is changing our socioeconomic system and redefines the way we work, live, collaborate, decide, and act as individuals, businesses, or societies. We have entered a phase of a drastic transformation of markets, businesses, education, government, social welfare systems, companies, employment models, and social structures. The transportation sector, for

example, is going through a radical transformation, as fully autonomous cars will soon be a reality - and they will be far safer, more efficient, and effective. Autonomous trucks, smart containers, driver-less taxis, and smart cities are just some examples of the new era. Consumer habits will possibly shift from owning a car to consuming smart car services, on demand. Entire transportation networks consisting of fleets of autonomous cars will be orchestrated by AI algorithms and adapt in real-time to demand, traffic, and other conditions. This will transform the way people commute along with the way cities operate and grow – leading to a more efficient, sustainable, and safe transportation system. As another example, AI is changing how we build products and services: physical product manufacturing processes already benefit from AI-powered production lines, automated quality control systems, and continuous improvement processes.

I believe that under certain assumptions, this technological revolution will lead to a new era of prosperity, creativeness, and well-being. While there will be disruption in employment patterns and a certain level of technological unemployment, AI will create numerous new roles and specializations — with a focus on technology and science. In most cases, AI will play a supportive role to humans — empowering people to perform better in handling complex and critical situations that require judgment and creative thinking. The general adoption of AI will allow people to free up time from monotonous, low-value work, and engage with more creative and strategic activities. The workforce and the underlying employment models will move from long-term, full-time commitments, to flexible, adaptive models of services and value exchange. Moreover, AI will empower people to discover and learn more. Smart content discovery and collaboration tools will give everyone access to the world's digitized knowledge, ideas, and talent. The education system will be improved by AI-powered experiences on top of world-scale digitized content and data, along with scientific and general knowledge. For example, intelligent 'education agents' will be able to capture the needs of each student and synthesize optimal, personalized educational programs — matching students' intent, setting the right level and pace, and utilizing optimal synthesis of content and experiences. As a result, I do believe that there will be streams of new business opportunities and a strong basis for entrepreneurship, creativeness, and innovation. In the health space, AI will bring more accurate medical diagnoses, personalized medicine, and shorter drug development cycles that will significantly improve the overall effectiveness, level of service, and general access to health services.

But of course, there are risks and concerns associated with the application of AI. We need to find ways to ensure that AI is used in the interest of our society - and this is not only a responsibility of the policy makers. People need to achieve a general awareness and sufficient understanding of the technology, its potential, benefits, and associated risks. Societies need to adapt to the new technology landscape and embrace AI as a

‘smart tool’ that helps people to achieve more. People must switch to a life-learning mode — become open and agile in acquiring new skills and exploring new talents and career paths that are more relevant to the new order of things. But governments need to adapt as well - by modernizing laws, frameworks, social programs, and the education systems. New strategies are needed — for education, businesses, and social systems – and leaders must rethink how markets, companies, and employment agreements should work in the new era of AI. Thought leaders need to propose effective new rules, and frameworks to mitigate the risk of centralization of power and control of data and technology.

But these are just my thoughts. This book brings together unique insights and answers on the above topics from a diverse group of global thought leaders. It is organized into 17 chapters, each presenting one question on AI and multiple responses reflecting a variety of backgrounds and standpoints. I am extremely grateful to the 60 leaders – the amazing group of academics, business leaders, technologists, authors, and researchers who made this happen. I am also grateful to Lucy Woodhead who helped in the various stages of the production process and to Coy Chen who created aspects of the original visual design. A project made of pure passion for AI from 60+ people across the globe.

This book is ‘connected’. You are invited to join the 60 Leaders community on LinkedIn where you may share your thoughts and ask any questions regarding Data and AI. At the end of each chapter, there is a link that takes you to the corresponding LinkedIn thread where you can join the discussion and connect with the 60 Leaders community.

May 2022
Dublin, Ireland

George Krasadakis

CONTENTS

CONTENTS

CONTENTS

1 HOW WOULD YOU EXPLAIN AI TO A 5-YEAR-OLD?

Artificial Intelligence tends to become hype and, in some cases, even perceived as magic. Explaining AI with simplicity and clarity is both essential and challenging.

We asked various leaders to present AI in simple words and attempt to demystify it by explaining how AI technology works, defining the different types of AI such as Narrow AI, General AI, Super AI, and providing characteristic real-world examples of Artificial Intelligence applications.

Mike Tamir, Tirthajyoti Sarkar, George Panou, Alex Wang, and Siddha Ganju share their insights.

MIKE TAMIR

Chief ML Scientist, Head of ML/AI
SUSQUEHANNA INTERNATIONAL GROUP • USA

AI is all about *getting a machine to complete complex tasks*. It focuses, in particular, on getting the machine to solve tasks that are not easily completed with a fixed set of instructions. These days this is accomplished with Machine Learning (ML) where a machine is taught to solve more complex problems through what (for the most part) can be described as a 'trial and error' process, making mistakes and then correcting for those mistakes based on feedback until it gets things right.

For a formal definition of the scientific discipline of AI (as for my definition for a 5-year-old), AI can be applied to any case where complex tasks can be completed by some artificially created system (machine). Again, in modern times, this has predominantly focused on the area of Machine Learning. Most prominently, Deep Learning (DL), a subdiscipline of ML where artificial Neural Networks containing many layers of artificial 'neurons' are trained to complete these tasks, has become central to the large advances we have made over the past decade. Neural network applications have dramatically advanced our ability to solve complex tasks involving computer vision, understanding and working with language, representing and organizing complex data structures, commerce, recommendations, and much more.

Mike Tamir, PhD is a data science leader, specializing in deep learning, NLP, and distributed scalable machine learning. Mike is experienced in delivering data products for use cases including text comprehension, image recognition, recommender systems, targeted advertising, forecasting, user understanding, and customer analytics. He is a pioneer in developing training programs in industry-focused machine learning and data science techniques.

TIRTHAJYOTI SARKAR

Sr. Director
RHOMBUS POWER • USA

Artificial Intelligence, or AI as it is called in the short form, is a powerful and transformative technology of the modern world. It is being said that the majority of the progress that human society will experience in the 21st century, will be linked to how AI technologies are developed,

applied, and governed. If we want to draw some parallels with similar developments in the 20th century, we can say that AI is as transformative as the discoveries of nuclear power, vaccine, or the internet!

AI is supposed to be a human being's 'intelligent companion.'

- Dr. Tirthajyoti Sarkar

However, the idea of AI is not a 21st-century one. Highly intelligent and visionary men and women have been thinking about it since at least the 1950s. The core idea started at the same time computers were first being developed to aid human beings with routine tasks like business calculations or spreadsheet preparation. AI pioneers (such as Alan Turing and John McCarthy) started envisioning a computer system (combination of hardware and software) that does not just carry out repetitive instruction but can develop a capability of thinking for itself. They thought, that if such a machine can be developed, then it would be a faithful companion to human beings for all kinds of intellectual tasks – helping us in discovering new scientific facts, designing better products and services, and solving great challenges that the world faces, making everybody's lives a lot better. In that sense, AI is supposed to be a human being's 'intelligent companion' who may look and feel different (a computer program or a robot) but think, converse, and perceive the world just like us and provide us all the help to make it better. There has been a lot of research and development in the AI field since those early days. Today, the largest organizations (both private and Governmental) around the world are investing huge amounts of money and human resources in this field. Some of the brightest minds in the academia and industry are working on solving tough problems related to AI.

However, the current focus is on something called 'Narrow AI' that does not aim to become the general-purpose intelligent companion of human beings that we talked about but focuses on solving a narrow, specific task at hand. This 'Narrow AI' is a specialized tool, and there will be many such AI tools for all kinds of tasks, such as making a movie recommendation, playing video games, helping with automatic driving, helping industrial operations in factories, assisting doctors with healthcare delivery, even discovering new medicines and vaccines.

In contrast, Artificial General Intelligence (AGI) is supposed to have the true thinking and analysis faculties of a human mind, sufficient for

handling all general-purpose everyday tasks. Such an AGI system may not be the master of a single task but can handle any task at least at the level that a human can do. Naturally, we also expect an AGI system to understand the world around us - through vision, audio, and language communication - as well as a human being is expected to. Unfortunately, there is no such AGI system in sight today. Understanding visual clues and natural languages (e.g., English, not some computer programming language like Python) has turned out to be the toughest task for AI tools to handle. A lot of progress is being made and researchers are hoping to get to some level of AGI within this century.

Modern AI systems mostly work by learning from examples. Careful programming allows humans to feed a huge amount of data into these systems autonomously. Generally, this data is in the form of scenarios and their responses. AI programs analyze that data, find the hidden patterns, and construct a mathematical model of how that data could have been generated and also what response goes with what scenario. Next time, a new data point or scenario is presented to the AI, the model can predict the correct response without explicit programming or instruction. In that sense, it becomes a prediction machine and helps humans solve many challenging problems.

For example, given medical test results, an AI system can predict what probable ailments a patient may have, thereby aiding a doctor. Or, it can predict what movie you may like given the preference data about your friends and what movies they have liked in the past. Please note, however, that the AI tool that makes the movie recommendation is completely different (in inner workings and programming) than the one that helps a doctor assess medical ailments. That's why they all are individual, Narrow-AI systems.

Dr. Tirthajyoti Sarkar is a Data Science and Artificial Intelligence leader. He has worked in multiple industries including semiconductor, edge computing, and SaaS. He has published multiple academic papers, Data Science books, and holds many patents.

GEORGE PANOU

Head of Innovation Centre
EUROBANK • GREECE

So, you are probably wondering what the term AI (Artificial Intelligence) really means, right? I am sure all these words like algorithms, computers, and machine learning fill your mind with questions and curiosity. Well, think of AI as computers having a small brain. AI, in simple

words, is like enabling a computer to have human-like intelligence, not in terms of simple repetitive tasks but of human cognitive functions. For example, in the future, AI will enable a computer to demonstrate empathy, and act in such ways that make it difficult to tell whether you are talking to a human being or an AI – thus passing the famous Turing test challenge.

Any device that includes a modern computer inside, can perform simple tasks that require a type of intelligence. Imagine that you want to type in the word 'Robot'; while you are trying to type it, the AI can help you autocomplete the word and even autocorrect it in case you have misspelled it. Like humans, AI is capable of learning. Using Machine Learning, if you feed a computer with data, after a while it will be able to distinguish a chair from a table or a bird from a tree. It would be also able to group different objects or predict whether it is going to rain tomorrow or even if you are happy or sad while you are talking to the AI. Sometimes this AI can even learn on its own and perform tasks or even self-correct its own program.

Sounds strange? I would rather say it is complex rather than strange. Let me give you some more familiar examples. You may have used Netflix to watch your favourite tv series or film. Have you ever wondered how Netflix displays first your preferences and how it recommends you watch similar tv shows? Well, that is possible based on the technology behind Netflix that knows what you like because you have been feeding it with data – e.g. how long you watched a film, what was the theme, or the genre. And with that data, the AI finds and promotes the most suitable content for you. In other words, *you teach the AI system what you like and it responds back with recommendations*.

There are many definitions of AI, but for now, you can think of it as 'a near-to-human-like intelligence that has cognitive functions like humans do, and reacts intelligently to external stimuli. There are also many types of AIs; some perform simple tasks while others perform more complex ones.

The term 'Narrow AI' describes algorithms that focus on performing a simple task, like playing a game of Connect 4, a game of tic-tac-toe, or even predicting the weather based on specific data. Narrow AI does not interact outside its predefined limits. Even though we might like a futuristic world like in the film 'Minority Report' or the film 'I Robot', the truth is that we are mostly surrounded by Narrow AI that performs simplistic tasks without having human-like cognitive abilities. You shouldn't underestimate Narrow AI though, since it makes our lives so much more effortless, like performing repetitive tasks faster and more accurately than humans - and, most importantly, tirelessly.

On the other hand, when an AI system has reached a high level of cognitive abilities that imitate human thinking, we will have reached the

'General' or 'Super AI'. Computers of this class can 'think' independently and take into consideration several parameters, most of the time unexpected ones, and make decisions based on not only *hard facts* but also *emotions and consciousness*. In the future, General AI will be able to create its own music, play a musical instrument and even create art like a painting or even a book.

One example where applied AI brings a significant positive change to our lives is the self-driving car. By using Deep Learning technologies, self-driving cars can eliminate transportation problems and minimize traffic accidents.

George Panou is the Head of the Innovation Centre at Eurobank, leading Digital Innovation and Digital Transformation for more than 20 years working with Financial Institutions-Banks, Large Enterprises, and Public Sector in Greece & EMEA.

ALEX WANG

Data Science Consultant
DELOITTE • AUSTRALIA

Human intelligence can be defined as the general mental ability that allows us to perform crucial activities like reasoning, problem-solving, and learning. Unlike human intelligence, Artificial Intelligence is built to solve specific problems. With AI techniques, we can perform various cognitive tasks, such as facial recognition, predictions, or even more complex such as music composition, etc. Moreover, AI enables a machine to make its own decisions, for example, a self-driving car is making decisions in real-time, while driving without human intervention.

AI systems are trained on specific data to solve a particular problem. But, whenever there is a need to address another type of problem, we have to obtain new data, specific to the new problem to be solved and train a new AI model. This is known as Narrow AI, according to which we have to follow this approach and train a specific AI model for every single problem we want to solve.

In the future, we want to move to General AI – a technology that can deal with many different problems. When we reach this point, we will be feeding a single AI system with a variety of data. This single system will be able to learn how to deal with multiple different problems. To some extent, General AI will have similar capabilities to what we humans have. When we reach the point where technology exceeds the capabilities of human intelligence, we will be entering the state of Super AI – a class of

systems able to build new forms of AI. Currently, we don't know when this will be feasible or even if we want this to happen in the future.

What we have today is Narrow AI. A characteristic technology is the Neural Networks - computer algorithms that imitate certain functions of the human brain. Neural networks are made up of interconnected layers of virtual neurons. The neurons are the core units responsible to pass on the information and perform calculations. The connections between the neurons have values associated with them, called weights. A bigger value means that the information that is carried by this connection is more important to the next layer. The values of neurons and the weights of the connections are learned from data through model training - the process during which the Neural Network learns from the input data via an optimization process based on backpropagation that attempts to find optimal values of the parameters that minimize the 'loss function'. This is how Neural Networks learn from failure[123].

Alex Wang has extensive experience as a data science consultant. She works at UTS as a casual academic, and a member of AI4Diversity, which is a non-profit initiative that engages and educates diverse communities about AI.

SIDDHA GANJU

Senior Artificial Intelligence Lead

NVIDIA, NASA FRONTIER DEVELOPMENT LAB, SPACEML • USA

A five-year-old probably knows the difference between a living being and a non-living thing. A living being exhibits some key characteristics like intelligence and emotions. Intelligence tells us not to touch fire and what a cat or a dog looks like. But how does a non-living thing like a computer know what a cat looks like? Using Artificial Intelligence we can teach a computer to think and reason what makes a cat look like a cat. This is the process of recognition - recognizing people (Taylor Swift), objects (chair), animals (cat), or places (Eiffel Tower in Paris).

If you are 5 years old today then you may have used a computer in some form - either a mobile phone, a tablet, a smart home device, or even the television. You may also have experienced using applications like Snapchat. How does Snapchat figure out where to put the dog ears on your face. When we are young we are taught parts of the face like where are my ears or nose or eyes are. So who teaches a computer that? A computer can also be taught to learn locations of things using Artificial

[1] Human intelligence and brain networks, Roberto Colom 1, Sherif Karama, Rex E Jung, Richard J Haier
[2] Chalfen, Mike, 'The Challenges Of Building AI Apps'. TechCrunch. Retrieved 27 November 2021
[3] The Cambridge handbook of artificial intelligence. Frankish, Keith., Ramsey, William M., 1960-. Cambridge

intelligence - so not just the location of eyes or ears, but also relative locations like left, right, bottom, etc. This is the process of detection - recognizing not just the object but also where it is located.

Artificial Intelligence is the same science that runs Alexa or Siri or Google Home. If you have talked to Alexa and given it instructions like which music to play or to place a call to your grandparents Alexa is using AI to understand and interpret what exactly you want it to do. AI helps a computer like a mobile phone understand language.

Siddha Ganju is an AI researcher who Forbes featured in their 30 under 30 list, is a Self-Driving Architect at Nvidia. As an AI Advisor to NASA FDL, she helped build an automated meteor detection pipeline for the CAMS project at NASA, which ended up discovering a comet. Previously at Deep Vision, she developed deep learning models for resource constraint edge devices.

JOIN THE DISCUSSION **SPREAD THE WORD**

2

WHAT IS THE ACADEMIC CURRICULUM FOR A SUCCESSFUL CAREER IN AI?

The demand for AI talent is growing fast and there is great potential for impactful careers in the AI space. We asked various leaders to provide their advice to a hypothetical high-school student who considers an AI career.

What are the key specializations in the AI space? What are the essential skills required for a successful career in the AI space? What should be the ideal academic background for a successful AI professional?

Michael Wu PhD, Aruna Kolluru, Mike Tamir, Cynthia Rudin, Tirthajyoti Sarkar, Anand Ranganathan, Jon Skirnir Agustsson, and Siddha Ganju share their insights.

What is the academic curriculum for a successful career in AI?

MICHAEL WU PHD

Chief AI Strategist
PROS INC. • USA

AI is already one of the hottest professions in the years to come. This trend is further accelerated by the pandemic, driving the rapid adoption of contactless digital technologies, many of which require AI to provide a good user experience. So how can you get involved and be successful in this exciting field?

Clearly, a solid foundation in STEM[1] education is crucial. AI is essentially applied Machine Learning to automate everyday human decisions and actions, with the ability to learn and improve through feedback. And machine learning (ML) is nothing more than applied Statistics to large-scale problems with high dimensionality. So you are going to need the Math (mainly linear algebra and numerical optimization) and Statistics background to think and talk like an AI professional. And you are going to need the coding and Computer Science (CS) skills to turn your brilliant ideas into efficient and scalable algorithms.

AI is inherently an interdisciplinary subject.

- Dr. Michael Wu

This foundation is a stepping stone to a successful career in AI, but it's insufficient because AI is evolving so rapidly. Although many universities are offering programs in Data Science, they can't keep pace with the latest technologies in the industry. This is where boot camps and conference workshops come in handy, but you do need the Math, Stats, and CS background to get the most out of these extracurricular training.

Finally, to be successful in AI, you will need to enrich the fundamentals you learned from school (which tends to be very generic) with domain-specific knowledge (which is rarely taught in school) through internships. This will also provide you with hands-on experience with large-scale cloud-based storage and computing technologies that are crucial for developing, deploying, and maintaining real-world, AI-based solutions.

[1] Science, technology, engineering, and mathematics - Wikiwand

What is the academic curriculum for a successful career in AI?

Because the field of AI is growing so fast, it's already impossible for any one person to keep up-to-date with all the latest technologies and innovations. AI professionals have no choice but to specialize. In industries, there are already specializations based on different stages of the data lifecycle. There are Data Engineers responsible for capturing and storing the data necessary to train the ML. Data Scientists are generally focused on using these data to train ML models to address business problems. Then there are engineers working in ML Ops[2], and they are responsible for deploying the trained ML model to the cloud, so they can be used anywhere the cloud is accessible via the internet. Moreover, there are technical specializations in AI that depend on the domains (e.g. autonomous vehicles, natural language processing, etc.), data types (e.g. image, video, text, etc.), coding language (e.g. Python, Java, R, etc.), and even the growing number of tools and packages (e.g. Tensor Flow, PyTorch, SciKit-Learn, etc.).

There are even specializations based on the development lifecycle of the AI solution. AI researchers tend to focus on the development of new ML algorithms, which leverage new data sources or use different ML paradigms. Although many industrial applications of ML use the Supervised Learning paradigm (which includes deep learning), there are more and more applications that leverage Ensemble Learning, Reinforcement Learning, Adversarial Learning, One-shot Learning, and many others (e.g. Semi-supervised, Unsupervised, Active, Online, Transfer or Transductive learning, etc.). AI researchers may sometimes need to develop new learning paradigms to optimize the utilization of the available training data.

Once a new ML algorithm is developed and tested, they are ready for use in AI applications. This is where AI Solution Architects and AI Application Engineers come in. Due to the rapidly evolving field around AI, these titles will likely change over time. Because they aim to build AI solutions to address business problems, they are sometimes referred to as AI Data Scientists or AI Scientists in general. These AI Scientists will design and build AI applications using the new ML algorithm as a mechanism to learn through feedback data. This will enable the AI applications to improve upon the decision it automates, with usage. Although the specializations in AI arose out of practicality, it's perfectly fine to be a jack of all trades in AI, but you do need to have expertise in at least one area to be successful.

As you can see, it's not easy to be at the center of the AI revolution. If you can check all the boxes mentioned above, you are likely a very good AI scientist already. However, AI is a very versatile tool, and it can be applied to such a wide range of problems that it's been said that AI will catalyze the Fourth Industrial Revolution[3]. So in practice, AI is inherently an

[2] MLOps - Wikiwand
[3] Fourth Industrial Revolution - Wikiwand

interdisciplinary subject. This often poses challenges for those working in a technical area of AI, because they must frequently collaborate with people from all walks of life. For example, you may need to work with lawyers to understand the legality of using certain data attributes. You may need to work with user experience designers to ensure the AI tool you are building will drive adoption. You may even need to work with anthropologists and sociologists to understand the societal impact of deploying your AI tool.

A great AI scientist can comprehend interrelated disciplines.

- Dr. Michael Wu

What distinguishes a very good AI scientist from a great one is the ability to comprehend these interrelated disciplines around AI and the ability to communicate and work with a highly diverse group of experts. Having the working knowledge in these tangential disciplines would make you a highly desirable AI job candidate. The good news in all this is that you don't need to be technical to work in AI. We need designers, ethicists, lawyers, doctors, teachers, entrepreneurs, and many others to help AI scientists build AI solutions that are more user-friendly to the layman, more compliant with ethical standards, and use data more responsibly.

So my advice to the future generation of AI enthusiasts is this, "stay curious and follow your passion." It is likely that in whatever you do or wherever your passion leads you, there will be some domain-specific AI tools that you can use, enhance, or maintain. And if you are curious enough, you will find the resources to learn what's necessary to excel in ML and AI. Aside from proprietary algorithms and domain-specific knowledge, most of the resources you need to get started with ML and AI are freely accessible on the Internet as open-source content. In the future, anyone can be an AI professional.

Dr. Michael Wu is the Chief AI Strategist at PROS (NYSE: PRO). He's been appointed as a Senior Research Fellow at the Ecole des Ponts Business School for his work in Data Science, and he serves as an advisor and lecturer for UC Berkeley Extension's AI programs. Prior to PROS, Michael was the Chief Scientist at Lithium for a decade. His R&D won him recognition as an Influential Leader by CRM Magazine. Michael has served as a DOE fellow at the Los Alamos National Lab. Prior to industry, Michael received his triple major undergraduate degree in Applied Math, Physics, and Molecular & Cell Biology; and his Ph.D. from UC Berkeley's Biophysics program.

What is the academic curriculum for a successful career in AI?

ARUNA KOLLURU

Chief Technologist, AI, APJ
DELL TECHNOLOGIES • AUSTRALIA

Artificial Intelligence is not a profession, it's a technology that enables the extraction of insights from data. AI requires varied skills and different roles throughout the data science life cycle. There are many specialised roles in AI, the most characteristic ones being:

- **AI Architects** determine what business problems to solve. They choose the right technology, frameworks, and architect solutions to solve business challenges.
- **Data engineers** manage the data sources and prepare the data. Well, we can argue that the Data Engineer is not part of AI, but I have included it as it's an important piece of the puzzle for AI to be successful. A data engineer is a comprehensive data specialist who prepares data and makes it available within your organization. By extracting information from different systems, transforming, cleaning data, and combining different sources into a working database, data engineers become an important piece of the puzzle in the AI process. They analyze and cleanse the collected data from the various systems and tools, used in an ecosystem. To uncover the hidden treasures in a company's data, these individuals need to deeply understand business processes.
- **Data scientists** train the models using the prepared data. They use machine learning and predictive analytics to collect, analyze, and interpret large amounts of data to gain insights that go beyond statistical analysis. A data scientist builds machine learning or deep learning models to help organisations improve efficiencies.
- **AI or Machine Learning Engineers** deploy the models. Once the models are built, the process is very similar to the software development lifecycle. The role of an AI/ML engineer is to develop deployable models from models developed by data scientists and integrate the models into the final product. AI engineers also develop secure APIs for delivering models on demand.
- **AI Strategists** work in research or real-world solutions. They should have hybrid skills in both technology and business. A good AI strategist understands the power of both technology and business; they must be able to explore the art of the possible as the means of transforming the business. They also identify the important AI projects for the organisation and strategize the AI journey for the organization. In my experience, the biggest gap we have in the industry at the moment is the AI Strategists.

What is the academic curriculum for a successful career in AI?

Most often the perception is that people working in AI need a Ph.D. It really depends on what exactly the person is interested in. Skills required for different roles in AI vary but they typically include mathematics (statistics, linear algebra, and probability), programming languages (python, java, R, C++), neural network architectures or frameworks, and most importantly domain or industry knowledge.

For a high school student considering a career in the AI space, I would suggest focusing on and learning Mathematics and Programming. It's important to get your hands dirty so, do some Python projects. Get yourself familiar with some Machine Learning or Deep Learning frameworks like Sci-kit learn, Pytorch, or TensorFlow and try to apply them in solving real-world challenges.

It is important to have a degree in Computer Science which helps understand the basics and provides a strong foundation. A Master's in Data Science will add to your knowledge and give you the tools and skills required for developing a career in Data Science. Also, it is essential to explore an industry you are interested in and look at how AI is transforming that industry. There are a lot of massive open online courses (MOOCs) and university programs with a focus on AI in certain verticals. And lastly, If you are interested in research and improving and building algorithms then you will definitely need a Ph.D. All these if you would like to have a formal education. I have also seen a lot of people without formal education excelling in the AI space by learning from MOOCs.

Aruna Kolluru is passionate about business, technology, and emerging technologies. She helps organizations to visualize the art of possible, architect, and build solutions with emerging technologies. Before joining EMC & Dell Technologies, Aruna held various senior roles for IBM & Oracle. She has over 21 years of experience in the industry on varying technologies including AI, Blockchain, Deep Learning, Computer Vision, IOT Big Data & Analytics, and Enterprise Architecture.

MIKE TAMIR

Chief ML Scientist, Head of ML/AI
SUSQUEHANNA INTERNATIONAL GROUP • USA

Defining the multitude of 'data professional' roles has been a challenge over the last decade. For instance, the definition of Data Scientist has mutated terribly; many data scientists tend to use ML in order to best understand the data with which they work, but some major tech companies have also simply rebranded their traditional analyst roles as 'data science' roles in order to attract talent.

What is the academic curriculum for a successful career in AI?

Regardless of job titles, a high-functioning data team, focused on complex problems, needs several contributors with a diverse set of skills to be successful. This includes:

- **Data Engineering.** An AI/data team will need data, of course, and professionals focused on managing large volumes of data for processing and preparation. These professionals, often called Data Engineers, need a moderate (though not extensive) background as software developers and must have a strong background in the principles of parallel processing with map-reduce concepts and the latest technologies for implementing them. Data Engineers will also have a strong background in mathematical literacy and must be fluent in the kinds of mathematical transformations applied to their data.
- **Data/ML Science.** Scientists/applied researchers work with the data to draw conclusions (especially predictions). The highest functioning scientists in this role will use the best tools at our disposal and so often (but not essentially) use ML. Ideally, an ML scientist will focus most of their time hypothesizing what kind of techniques (ML models) will detect desired patterns or generate successful predictive results, run experiments to test these hypotheses, and then conduct extensive analysis to understand if their hypotheses were validated (and in what ways). The best Data/ML scientists will have extensive (graduate level) ML training in addition to mathematical literacy, especially in the area of linear algebra, and ideally functional analysis. They will need decent coding ability but need not be software engineers. They will also require a background in statistical modeling, at minimum having a strong background in probability theory and hypothesis testing.
- **ML Engineering.** To make the team most efficient, I often like to have a dedicated group of ML engineers. An ML engineer's job is to develop centralized code for the scientists to use to eliminate practical software challenges reproducing boilerplate code so that they can focus on understanding their experiments rather than implementing, debugging, and ensuring that their results have parity with prior work. ML engineers will fundamentally be software engineers but will have a working understanding of (though perhaps not graduate-level training in) the machine learning and mathematical literacy expected of scientists.

Mike Tamir, PhD is a data science leader, specializing in deep learning, NLP, and distributed scalable machine learning. Mike is experienced in delivering data products for use cases including text comprehension, image recognition, recommender systems, targeted advertising, forecasting, user understanding, and customer analytics. He is a pioneer in developing training programs in industry-focused machine learning and data science techniques.

What is the academic curriculum for a successful career in AI?

CYNTHIA RUDIN

Professor of Computer Science
DUKE UNIVERSITY • USA

I generally recommend that before students take a machine learning course, they should know *how to program*, and they should have taken somewhat advanced courses in *probability and linear algebra*. Many professors post their course notes and lecture videos online (as I do), so I would recommend trying to watch those to gain more knowledge.

The best way (in my view) to learn Machine Learning after taking a course is to work on a real problem. There is a lot to learn because real data science is messy: it takes effort to specify the goals correctly (based on what the domain experts say) and to collect and clean the data. It is common to find that the problem has been formulated in a problematic way and needs to be fixed. Important pieces of advice are:

1. Keep in mind that Neural Networks are excellent for computer vision, but they don't tend to give an advantage on many other kinds of data. Thus, it is worthwhile to learn other techniques besides Neural Networks, as they can be difficult to train, and are often black boxes.
2. Interpretable Machine Learning is critical for fairness and trust, but it is a topic that is rarely taught in courses.
3. Ethics are important. Data scientists should evaluate the ethics of the project themselves, not just do what their manager tells them! Ethics is another topic rarely taught in data science courses.

Cynthia Rudin is a professor at Duke University. Her goal is to design predictive models that are understandable to humans. She is the recipient of the 2022 Squirrel AI Award for Artificial Intelligence for the Benefit of Humanity (the "Nobel Prize of AI").

TIRTHAJYOTI SARKAR

Sr. Director
RHOMBUS POWER • USA

To understand the needs of the academic curriculum for a successful AI career, it will be prudent to get a good understanding of the various roles that AI professionals play in a typical AI team/organization. It is already a widespread domain of expertise and knowledge and it is getting wider by the day. A few core specializations within the space can be thought of as follows,

What is the academic curriculum for a successful career in AI?

- **Data scientists**. The generalists among the specialists. These folks, depending on the organization set up and task demand, are expected to know the entire breadth of AI/ML and participate in all kinds of tasks. However, in a somewhat narrower sense the term 'Data Scientist' denotes somebody who analyses the raw data through visualizations and statistical tests, performs feature engineering, applies and optimizes a number of ML algorithms, and creates insights and reports. Above all, they build prototype solutions that are not necessarily production-grade but help the AI organization make decisions on further development. Often, Data Scientists come from a wide variety of backgrounds and educational training – Physics, Economics, Biology, Marketing – that are different from a Software Engineering professional.
- **Data engineers**. These folks are also called 'Data plumbers' as they literally build the essential plumbing infrastructure through which datasets flow. If you can imagine the AI system as a pipeline (raw data as input and intelligent insights as output), then data engineers build, maintain, and improve these pipes. They are mostly concerned about how the raw data will come in from all kinds of sources – standalone databases, web APIs, gRPC calls, industrial protocol services, sensor networks, etc. They also ensure the smooth flow of that data (through all the transformations) to various parts of the AI/ML system and subcomponents as data scientists and ML engineers need it.
- **Machine Learning engineers**. Even a few years back, there was no real distinction between Data Scientists and ML engineers. However, with the advent and rapid rise of MLOps (think of it as a fusion between Machine Learning and Software Engineering DevOps), the role of ML engineers is becoming more specialized. In a modern AI-focused organization, they are, generally, in charge of the creation, deployment, maintenance, and tracking of production-grade ML models for customer-facing and enterprise-scale operations. The term 'production-grade' distinguishes them somehow from Data Scientists who are mostly building prototypes and may not be directly doing the deployment and long-term maintenance. ML engineers should have solid Software Engineering and DevOps backgrounds as many principles from those disciplines are used in this role. Training and deploying an ML model is often not the end of their job but just the starting point. They also have to log and track how the models are performing in real-world scenarios and apply fixes regularly. Some organizations have started placing ML models at the centre of all operations.
- **Machine Learning/AI scientists, algorithm researchers**. These roles go to academically-trained research professionals (for example having a Ph.D. degree and research experience as a post-doc) whose

responsibility is to expand the frontier of AI/ML knowledge by coming up with novel algorithms and concepts altogether. Note that the immediate business impact or application of these research concepts may not be apparent, but they should open up new areas of exploration or solve (completely or partially) a significant existing problem in some AI sub-domain (even if it is a theoretical solution for the moment).

Having looked at these general descriptions, it should be apparent to the readers that the academic background required for a successful AI professional can be widely varying. Statistically speaking, most professionals come from some sort of STEM background with Computer Science, Electrical/ Electronic Engineering, and Physics backgrounds being the majority. A lot of Data Scientists come from Applied Mathematics and Statistics backgrounds.

Programming skills are highly desired and useful for any AI role but not everybody comes with years of experience with coding and software development. It is nothing to be worried about and can be picked up with diligent practice and self-learning (even on-the-job training).

A general inclination towards mathematical thinking and an analytical mindset is almost a necessary prerequisite. Experimental mentality and curiosity to solve problems using logical thinking and data-driven methodologies are what matter most.

Some roles are greatly benefited from specialized experience. For example, Data Engineers will benefit from previous database management experience, Data Scientists from Statistical model building experience, and ML engineers from Software deployment experience.

These days, new, exciting roles are getting created at all levels of organizations for making AI a truly transformative force. For example, AI application scientist or AI product manager will be in charge of pitching the great features of an AI-enabled product or platform in front of potential customers. A background in marketing, product engineering, or sales can help them. AI Ethics is a new and expanding field that looks at the ethical and moral implications and boundaries of AI applications and how complex dilemmas can be solved. People with social science backgrounds (e.g., philosophy) are filling many such roles in top organizations.

Overall, AI is a vast and complex field with a place for professionals with all sorts of backgrounds. With proper motivation and due diligence, anyone can hope to get an entry and prosper.

Dr. Tirthajyoti Sarkar is a Data Science and Artificial Intelligence leader. He has worked in multiple industries including semiconductor, edge computing, and SaaS. He has published multiple academic papers, Data Science books, and holds many technology patents.

What is the academic curriculum for a successful career in AI?

ANAND RANGANATHAN

Chief AI Officer
UNSCRAMBL • USA

In my mind, an essential academic curriculum would have a good balance of four different competing sets of prerogatives, namely: Symbolic & Sub-symbolic techniques, Algorithms & Ops, Theory & Applications, Formal Rigor & Story-telling.

Data storytelling is critical to conveying the results to business or other stakeholders.

- Anand Ranganathan

Let's look at each of these elements in detail.

- **Symbolic & Sub-symbolic AI**. AI can be broadly broken into logical (or symbolic) AI and Statistical (or subsymbolic AI). Symbolic AI includes topics like Knowledge Representation, Logical Reasoning, and Planning. The key idea is that knowledge is represented and reasoned upon explicitly. This branch of AI was popular till the 90s with the rise of expert systems, rule formalisms, and logical programming languages like Prolog. Sub-symbolic AI includes topics like Machine Learning, including Deep Learning. The key idea is that knowledge can be learned with sufficient data. This branch of AI dramatically increased in popularity since the 2000s, and today is probably the dominant branch. Most AI and Data Science courses focus on subsymbolic AI, but I feel that a true AI system will need a combination of Symbolic and Sub-symbolic techniques. Hence, I believe that AI researchers must focus on Symbolic AI and its co-existence with Subsymbolic AI.
- **Algorithms & Ops**. A lot of emphasis today is placed on different algorithms and techniques for processing data, building models and evaluating or tuning them, but not on the operationalization of these algorithms for production settings. I believe that an AI curriculum must cover different aspects of AI or ML Ops like robustness, scalability, and monitoring. Furthermore, it also needs to cover the data

architecture – how is data sourced, transformed, stored, and made available for training and scoring. This is especially important for AI application architects.

- **Theory & Applications**. I believe the AI curriculum must have a good mixture of theory and applications. For example, for various Machine Learning models, AI professionals must understand both the internals of the training and scoring algorithms, as well as various applications of the model (for solving different problems in different industries). A good understanding of the theory of the algorithms is essential to understand its behavior and pros and cons. And the use-cases are important so that a Data Scientist or a Researcher quickly knows what algorithms to apply for different kinds of problems they might encounter.
- **Formal Rigor & Story-telling**. Todays' AI curriculum focuses on formal rigor – i.e. making sure that the algorithms and techniques perform well in a formal sense (e.g., they have high accuracy, precision, recall, or other metrics). Not as much emphasis is put on story-telling, i.e. creating a story about the data. The story about the data is critical to convey the results to business or other stakeholders in an easy-to-understand manner. The story might include different elements like an overview of the key takeaways, a description of the source and properties of the data, a diagnosis of key findings, predictions of what might happen next, and prescriptions of what actions might be taken to optimize some outcomes. The stories help weave different elements of the data and the models to communicate insights effectively to the listener. This is especially important for Data Scientists.

For a high-school student considering a career in the AI space, I would recommend looking for courses that cover a gamut of the topics I mentioned above. This will help them acquire skills that are immediately useful in the marketplace, as well as skills that will be useful even in the future as the state of AI evolves. Furthermore, I'd recommend them to take inter-disciplinary courses to get more ideas of interesting ways of applying AI in different disciplines.

Anand Ranganathan is a co-founder and the Chief AI Officer at Unscrambl. He is leading Unscrambl's product development in several cutting-edge areas, including natural language processing, automated insights, data story-telling, real-time optimization and decision-making, and marketing optimization. Before joining Unscrambl, he was a Global Technical Ambassador, Master Inventor, and Research Scientist at IBM. He received his PhD in Computer Science from UIUC, and his BTech from the IIT-Madras. He also has over 70 academic journal and conference publications and 30 patent filings in his name.

What is the academic curriculum for a successful career in AI?

JON SKIRNIR AGUSTSSON

VP Artificial Intelligence and Data Research
NOX MEDICAL • ICELAND

AI is used to describe a broad set of tools used for pattern recognition, automation of cognitive tasks, augmentation of data, and generation of new outputs or content. Since AI encompasses a broad set of tools and methods which are applied to almost every area of our lives one can approach the field from multiple different angles. In my job, I lead an interdisciplinary team to develop Artificial Neural Networks and other AI tools to be used in clinical practice. In our world, developing and training the models encompasses only a very small part of the whole development and release journey of AI products. To be successful, we need to understand how the AI products are used and by whom; we must understand how the data we use is generated, this includes understanding physiology, sensors, and instrumentation; we must understand how the users and patients interact with our AI products and are impacted by our outputs; and we have to do all of this within the relevant ethical, regulatory, and legal boundaries.

In our world, there is no one way of having a successful career in AI, but overall my impression is that the people who outshine others are those who have a strong foundation in a specific field to build on and have an affinity for learning and collaborating with people with very different backgrounds. In our case, we may have engineers having very deep dialogs with medical doctors where they need to build a mutual understanding of how a disease is diagnosed, what are the edge cases, how AI is best developed for each use case, and how a wrong diagnosis can be mitigated. It takes strong communication skills to explain and learn from people who approach the same challenges from very different directions.

For young people who are interested in pursuing a career in AI, my advice would be to specialize in a field that highly interests you. If you are inclined to specialize in Computer Science, Engineering, Math, Physics, or another technical domain you would learn Programming, Statistics, and other relevant hard skills. When you have established your hard skills, it becomes increasingly important to understand the ecosystem your AI products will be used in and to learn to communicate with key stakeholders in that ecosystem. For others who are more interested in areas where AI might be a useful tool, I would advise them to get their education in the field they are interested in. Studying languages, history, medicine, or something that is not directly related to AI does not mean you will not have a shot at building a career in AI. If you learn to understand how AI can be leveraged in your field and learn to communicate with technical

people you can be at a great place to introduce or enhance AI in your area of speciality.

AI is just a tool and for it to have a real impact it should be applied to the right challenges, at the right time, and in the right manner.

Jon S. Agustsson is an experienced AI and Research leader working in the medical device and medical research industry-leading an interdisciplinary team in Data Science, Physics, Electrical Engineering, and Research. Passionate engineering, inventing, and building new things, with multiple patents and scientific publications.

SIDDHA GANJU

Senior Artificial Intelligence Lead

NVIDIA, NASA FRONTIER DEVELOPMENT LAB, SPACEML • USA

Having interviewed many people who have multiple specializations from deeplearning.ai, Coursera, Udacity, DataCamp, Udemy, etc, the candidates who tend to stand out are those who demonstrate a working knowledge of AI/DL. How do you do that? With a portfolio of projects that involve different data domains (text, image, multimodal) and tasks (classification, generation). Assuming that you have a good grasp of the theory and you're ready to take up a small project that allows you to practice your skills, it is a good idea to find a project that resembles something in the industry and work on it to improve the breadth of your knowledge and try things into production. Let's say for example that you decided to develop an image similarity search system that finds the closest image to a given query image. The following list of questions can help you, not only to prepare the project itself but also to get a good sense of the required skills and knowledge required:

- What would your complete end-to-end pipeline look like?
- How would you build a cloud API to serve the web frontend?
- How do you scale from hundreds of images vs millions to billions?
- What would be the cost involved in scaling this up?
- How would you evaluate performance metrics, eg latency, and accuracy for model drift?
- How would you index the new incoming data? Would you rebuild or incrementally update?
- While scaling up how do you make your network and pipeline efficient? How do you reduce the floating-point computations in your network? How would you reduce the size of the embeddings while still having the same representative power?

What is the academic curriculum for a successful career in AI?

- What are the potential sources of bias?
- In retrospect, consider the effort you put into this. What would you have done differently if you had one day, one week, or one month to develop a solution?

Follow the trajectory of these questions in your projects and you'll find yourself approaching problems from different perspectives, acquiring knowledge, developing multiple solutions, and evaluating the merits and demerits of each. Beyond technical skills, you should be looking for opportunities to elevate your profile with ideas such as:

- Communicating your work and its impact effectively, to both specific stakeholders and the general audience.
- Develop and engage in leadership roles be it for organizing the project or organizing people and teams.
- Maintain professional relationships.
- Create an identity for yourselves in the field, and market yourself.

Siddha Ganju, an AI researcher who Forbes featured in their 30 under 30 list, is a Self-Driving Architect at Nvidia. As an AI Advisor to NASA FDL, she helped build an automated meteor detection pipeline for the CAMS project at NASA, which ended up discovering a comet. Previously at Deep Vision, she developed deep learning models for resource constraint edge devices.

SHARE YOUR THOUGHTS SPREAD THE WORD

3

WHAT IS THE MOST IMPRESSIVE TASK THAT AI CAN ACCOMPLISH TODAY?

Artificial Intelligence has already impressive achievements to present, for example, the ability of machines to 'understand' natural language, to 'see' and detect patterns, and to make sophisticated, real-time decisions.

But what would be the most impressive achievements of Artificial Intelligence so far? How is AI being leveraged across industries, from manufacturing to healthcare? How is AI helping the world become more connected?

Cynthia Rudin, Jordi Guitart, PhD, Jyotirmay Gadewadikar, Aruna Kolluru, and George Panou share their insights.

What is the most impressive task that AI can accomplish today?

CYNTHIA RUDIN

Professor of Computer Science
DUKE UNIVERSITY • USA

AI is helping society in numerous ways, though as we know, it is also being weaponized. Let's consider some of the positive examples first.

IBM was a 2020 Edelman finalist competitor for their work on using Machine Learning to predict computer server failures. These servers form the backbone of the Internet, so it is important that they don't fail. IBM used natural language processing to analyse trouble tickets and put them into a predictive model for server failures. It was an astonishing data science project that was quite technically difficult. As a second example, a charity called 'Give Direct' uses Machine Learning on satellite images to estimate which villages were the poorest. Give Direct can then target those villages to give out money.

My lab is working on automatically generated poetry and music.

- Cynthia Rudin

There are a lot of wonderful, innovative, and beautiful examples like the two above. In my own projects, I've worked with the power company in New York City to predict power failures (namely fires and explosions on the underground electrical grid) using data from the time of Thomas Edison. I've also worked with neurologists to design models that can be used in the intensive care unit to predict whether a patient might have a brain-damaging seizure in the near future. As another example, my algorithms are used by police departments to determine which crimes might be part of a crime series, which is a set of crimes committed by the same people.

As I mentioned, I'm quite concerned lately about the weaponization of AI. AI does help us be more connected since it underpins the recommendation systems used in social media and search engines, but it also

What is the most impressive task that AI can accomplish today?

allows bad actors to propagate misinformation, leading to genocide, bullying, and violence in the real world. I am quite honestly concerned about what AI is doing to the future of the world and I do not see any clear way to stop it. There are other serious concerns about the interaction of AI with personal privacy. We now have facial recognition systems that are extremely accurate, even across different races. But to build them, a company scraped all our personal photos from the internet! Of course, we did not intend our biometric information to be used for this purpose. Unfortunately, there are almost no regulations on the collection of such data and the use of AI on it, so these problems will continue to occur.

In terms of the characteristic examples of AI, I think the most basic one is automated handwriting recognition. Whenever you put a check into an ATM to be deposited at the bank, or when you send a letter through the mail, an AI system reads your handwriting so that the amount does not need to be typed up manually. In the future (not quite yet) we will hopefully have reliable vision systems for self-driving cars, which could be another characteristic example of AI - after we get it to work! But we should be careful not to launch AI systems too early, in cases when it is dangerous to do so. Another characteristic example of AI is the algorithms that are used for playing games, such as freestyle chess and Jeopardy.

In the future, I hope that AI will help humans with creative tasks too. My lab is working on automatically generated poetry and music. It's not as good as a human, but we are getting there. I find it easier to edit a computer-generated poem than to write a poem myself!

Cynthia Rudin is a professor at Duke University. Her goal is to design predictive models that are understandable to humans. She is the recipient of the 2022 Squirrel AI Award for Artificial Intelligence for the Benefit of Humanity (the "Nobel Prize of AI").

JORDI GUITART, PHD

VP of Artificial Intelligence
AIZON • USA

AI is here to stay, and we must agree that the hype is now over according to the steady growth of investment in AI, spanning virtually all industries and sectors along with their digital transformation. The tasks AI can accomplish today are almost infinite as per our human unbounded capacity to think and dream. Still, the financial payoff has not achieved the spectacular promised results due to a mix of failures and successes almost all enterprises and verticals have experienced by adopting

What is the most impressive task that AI can accomplish today?

Artificial Intelligence. Regrettably, many companies have discovered that having large amounts of data is not enough to secure successful outcomes from AI. And many more have paid high bills aiming at addressing wrongly stated problem statements, often too broadly defined. Without a delimited and well-defined purpose, no ROI should be possible and AI technology, whether Machine Learning or Deep Learning, is not an exception. But the resilience of first AI adopters will end up rewarding the big majority of them thanks to their unique learnings gathered across their AI transition journey.

The most impressive task that AI can accomplish is yet to come and it will deal with Human health.

- Jordi Guitart, PhD

Currently, Machine Learning and Deep learning are being intensively used in research with promising results, particularly Neural Networks as they dramatically accelerate computationally intensive processes, e.g., drug discovery[1]. Research benefits from a high degree of liberty that is hardly seen in other areas, thus being the ideal candidate for investing in powerful, though complex, Neural Networks architectures – only limited by budget allocation capacity. On the other end, there are all those industries that are subject to heavy regulations for which interpretability of Deep Learning models is a problem. Thus, following the example of drug discovery, the adoption of Neural Networks for the subsequent manufacture of these new drugs is a no go right now, according to Good Manufacturing Practices regulation in Pharma and Biotechnology.

Moving away from highly regulated industries, AI has already demonstrated its feasibility to generate potential value from the use of pattern recognition (people's faces, sign language, etc.) to fully autonomous equipment (self-driven cars, home vacuum cleaners, etc.) thanks to Computer Vision using Convolutional Neural Networks; and from voice-led instructions to activate appliances, select and execute tasks, to text mining

[1] "Deep Learning Driven Drug Discovery: Tackling Severe Acute Respiratory Syndrome Coronavirus 2". Front. Microbiol., 28 October 2021

or simultaneous translation using Natural Language Processing technology. No doubt Deep Learning is the field of AI that delivers the closest Human-to-Machine and Machine-to-Machine interactions without simply mimicking human behaviour as Robotics Process Automation does from a collection of pre-set rules.

For sure, we can all look astonished at the performance of a fully autonomous car, but AI still has a lot to do to make the self-driving experience fully secure, and again regulation has a lot to say. Then maybe simple yet irrelevant things like telling Siri that you feel happy, and it selects one of your most-loved songs to play becomes impressive as we are absolutely unconscious of the swift connection we have built over Artificial Intelligence.

Still, to me, the most impressive task that AI can accomplish is yet to come and it will deal with Human health, and more precisely, the diagnosis and prognosis of patients' life-threatening conditions. The challenge is paramount as there are no two equal persons in the world – yes, even identical twins have genetic differences[2]. Additionally, any intervention, such as any initiated treatment, severely biases patients' health data. Hence, a predictive model aiming at diagnosing, for instance, the risk of metastases in breast cancer patients, is subject to an excessive number of factors (variables aka *features*) that dramatically reduces any subset of available patients to an insufficient number of samples that could be decently used for model training.

And here the word 'decently' has all connotations: *legal, medical, social, cultural, ethical, and economical*. How to deal with data privacy and data protection is just an example of issues that easily prompts our minds regarding patients and healthcare institutions, data controllers, and data processors. But these issues are all inherent to AI and they must be addressed in full alignment of all healthcare stakeholders, a common will, and a titanic effort.

Jordi Guitart, PhD has been recently appointed CEO of the Barcelona-based healthtech startup Science4Tech Solutions, coming from Aizon where he served as VP of AI. He is concurrently Adjunct Professor of Strategic Management at ESIC Business & Marketing School.

[2] H. Jonsson et al., "Differences between germline genomes of monozygotic twins"

What is the most impressive task that AI can accomplish today?

JYOTIRMAY GADEWADIKAR

Chief Engineer (Enterprise AI – Systems and Mission Analysis)
MITRE • USA

We all interact with AI products, for example, through an in-home voice-activated personal assistant or receiving a product recommendation on your favorite e-commerce website. It is abundantly clear that AI is omnipresent, driving innovation and influencing how businesses run and compete. AI has been known for nearly six decades. However, AI was thrust to the forefront of the new industrial revolution only recently. AI is the dominant driving force for modernizing a seemingly never-ending list of industries and functions such as communications, healthcare, media, education, audit, taxation, and operations. In short, the current and future use of AI in innovation can change how people live, work, play, and even think. Below are a few examples of what AI can accomplish today.

Natural Language Understanding and Generation. AI can process natural language, and it can read and comprehend what the intent is in documents, news articles, blogs, books, and emails. Furthermore, it can generate summary information from processing the text and label the text with topics. The technology can detect a question in text and group similar questions together. Note that humans have a lot of variation in speech, as someone can ask the same question differently. Similarly, AI can use existing unstructured data to find answers to these questions. The ability to process unstructured text and respond to the text has further been helped by the technology to hear using state-of-the-art pattern recognition speech-to-text (STT) techniques. STT application programming interfaces can convert spoken words by humans into text in real-time. Generating human-like speech is an excellent accomplishment of speech synthesis techniques. Speech synthesis is the artificial creation of human speech where a text to speech (TTS) system converts natural language text into speech. This synthesized speech can be created by concatenating pieces of recorded speech that are stored in a database, and the quality of these speech synthesizers has been becoming more and more similar to the human voice every year. Text to speech technologies leverage Neural Network models to deliver a human-like, engaging, and personalized user experience. This capability of using existing unstructured databases and grouping questions and answers facilitates responding to those questions being asked in real-time. So not only is AI understanding, but it also generates natural language and appropriately customizes responses into a specific answer based on the context by identifying the sentiment of the question. These capabilities result in creating an engaging user experience in achieving a great degree of containment in virtual voice assistants with interactions that mimic a human

conversational style. The ability to be empathetic is critical. For example, "I would like to reschedule my flight" and "I just broke with my girlfriend, I'm rebooking my flight to leave early" would be different because the contexts in both these situations are different.

Advanced Inference Engines. AI facilitates incorporating rapid and advanced inferences, and inference engines are part of the decision systems that apply logical rules to the knowledge base to deduce new information. For example, inference engines based on Bayesian Belief Networks are used for automated breast cancer detection support tools. AI is a viable option for computer-aided detection by representing the relationships between diagnoses, physical findings, laboratory test results, and imaging study findings. In the example above, the work brings essential roles such as Radiologists, Image Processing Scientists, Database Specialists, and Applied Mathematicians on a common platform[3]. By exploiting conditional independencies entailed by influence chains, it is possible to represent extensive cause-effect relationships using little space. It is often possible to perform probabilistic inference among the features in an acceptable amount of time. The inference engines are critical in recommendation systems, such as suggesting which movie to watch next and deciding which price point will maximize the revenue.

Jyotirmay Gadewadikar received the Scientific Leadership award from the US Department of Homeland Security and is an Artificial Intelligence, Decision Science professional engaged in Strategy, Business, People Development, and Thought Leadership. He was previously Chief Product Officer of Deloitte's Conversational AI Practice and System Design and Management Fellow at MIT.

ARUNA KOLLURU

Chief Technologist, AI, APJ
DELL TECHNOLOGIES • AUSTRALIA

Artificial Intelligence is really transforming the way we live and work. From chatbots to recommendation engines we are surrounded by AI systems. The most impressive task that AI can accomplish is Computer Vision - a branch of Deep Learning which helps machines see and perceive the way humans do.

Since the first breakthrough with AlexNet in 2012, the accuracy of the models increased from 50% to 99%. Computer Vision applications are used in every industry making our world a smarter and safer place to live. Computer Vision is really augmenting humans. At Dell Technologies we

[3] Gadewadikar et al., "Exploring Bayesian networks for automated breast cancer detection," IEEE Southeastcon 2009, 2009, pp. 153-157, doi: 10.1109/SECON.2009.5174067.

have been working with customers across a variety of industries, leveraging AI to innovate and grow.

In healthcare, where 90% of all medical data is images, Computer Vision is playing a pivotal role in developing new-gen healthcare systems for diagnosing 2D and 3D images. AI algorithms are used from diagnosing brain tumours to determining which cancer treatment will work best for a patient, enabling doctors and hospitals to better analyze genetics, lifestyle, and past medical history along with the real-time data coming from wearables to provide personalized care to the patients.

In the automotive sector, autonomous vehicles leverage some of the most sophisticated AI technologies. An autonomous car must deal with complex analysis and decision-making for navigation, automatic braking, collision avoidance, lane change, parking assistance, and many other decisions. Self-driving cars capture data from their environment and feed it back to the AI - a loop called the *perception action cycle*. Then, AI makes decisions that enable self-driving cars to perform specific actions in the same environment. This is an iterative process and the more cycles the more accurate the AI model. Moreover, accumulated data from such cycles from multiple vehicles help produce intelligence that finally empowers every single vehicle.

Every industry has numerous use cases where AI can be applied. In the Retailer industry, companies are using Computer Vision for checkout, seamless transactions, and theft detection. Cities are using Computer Vision for surveillance to keep the population safe. In agriculture, Computer Vision is used to detect weeds and monitor the crop's health and growth. Construction and manufacturing industries use AI to run operations more safely and efficiently, for instance, they apply computer vision for defect detection and preventive maintenance. Research also benefits from AI as it enables the optimization of costs and increases the productivity of complex research projects: Considering that the vast majority of research activities are led by data, Machine Learning models can simplify the whole process, allowing researchers to go through the cycles faster and improve their outcomes.

Aruna Kolluru is passionate about business, technology, and emerging technologies. She helps organizations to visualize the art of possible, architect, and build solutions with emerging technologies. Before joining EMC & Dell Technologies, Aruna held various senior roles for IBM & Oracle. She has over 21 years of experience in the industry on varying technologies including AI, Blockchain, Deep Learning, Computer Vision, IOT Big Data & Analytics, and Enterprise Architecture.

What is the most impressive task that AI can accomplish today?

GEORGE PANOU

Head of Innovation Centre
EUROBANK • GREECE

Artificial Intelligence has numerous applications in almost every discipline and every industry. Technologies like Artificial Neural Networks (ANNs) have contributed significantly to the advancement of AI and its major adoption by industries worldwide. However, Neural Networks are not new; they were first introduced back in 1944 by Warren McCullough and Walter Pitts, two University of Chicago researchers. The technology took off only around 2012 with the creation of 'AlexNet' developed by Alex Krizhevsky, Geoffrey Hinton, and Ilya Sutskever which won the ImageNet Large Scale Visual Recognition Challenge. Actually, before that, the scientific community was ostentatiously ignoring ANNs and the whole concept.

But what led to the wide adoption and exponential development of ANNs and why it is so important to the domain of AI? That is mainly because of the exponential growth of computational power with GPUs, faster computer networks, and more efficient and reusable code libraries. Another reason for the growth of ANNs adoption is the huge amount of data generated by users and electronic devices (IoT, smartphones) giving the possibility to train ANNs to perform various tasks. Cisco reported that in 2021 there are 27.1 Billion devices 3.5 per capita creating 33 Zettabytes, predicted to reach a mind-boggling 175ZB by 2025.

The fully autonomous car is a classic example of advanced AI.

- George Panou

ANNs have helped in many different industries. In healthcare, for example, ANNs are used to predict Alzheimer's or Parkinson's disease by predicting the disorders with accurate diagnosis at early stages. IBM has used its AI Watson to detect cancer from MRIs and correlate clinical trials and medicines specific for the type of cancer through a huge dataset helping save the lives of patients. Actually, by digitizing the neurons of the brain and then training the ANNs properly with the right data you may end up finding correlations between dependent and independent variables that humans may fail to correlate. In the food industry, ANNs can help identify poisonous ingredients or test the quality of food products

What is the most impressive task that AI can accomplish today?

through a 'digital nose' – even in cases where microbes could develop and go unnoticed by conventional tools and technologies.

In the automotive industry, the fully autonomous car is a classic example of advanced AI. Data are continuously gathered through sensors that capture the behavior of the driver, the road conditions, and the objects in the environment in order to provide an effortless and safer driving experience. Although the technology has advanced, you cannot simply program a car to drive itself from point A to B with no risk as there are too many unpredicted parameters during the ride. You must have an evolving unsupervised deep learning system that can autocorrect itself and adapt to unpredictable conditions. Tesla cars for example, by using ANNs and sensor devices, might be able to 'see' an accident before it happens; they even update their firmware on the fly to resolve urgent problems. Of course, there is still work to be done as AI is lacking ethical decision-making. For example, in a critical situation, would the car select to hit the pedestrians or try to avoid them by crashing on a wall and putting the lives of its passengers at risk? Google's Waymo is also disrupting the automobile industry by creating new ecosystems and platforms for self-driving cars.

George Panou is the Head of the Innovation Centre at Eurobank, leading Digital Innovation and Digital Transformation for more than 20 years working with Financial Institutions-Banks, Large Enterprises, and Public Sector in Greece & EMEA.

ASK A QUESTION

SPREAD THE WORD

4

HOW IS AI IMPACTING THE WAY BUSINESSES OPERATE?

To survive, businesses need to adopt Artificial Intelligence. They need to realize the potential of AI and transform their key processes, services, and products.

We asked various leaders to name the functions of the organization that can be optimized with Artificial Intelligence, provide industry examples, and discuss the risks associated with AI in a business environment.

Edward Pyzer-Knapp, Professor Andy Pardoe, Simon Greenman, Rajvir (Raj) Madan, and John Thompson share their insights.

EDWARD PYZER-KNAPP

WW Research Lead, AI Enriched Modelling and Simulation
IBM RESEARCH • UK

Before the COVID-19 pandemic, many businesses had begun their digital transformation and were beginning to use analytics and Artificial Intelligence (AI) technologies to understand and improve their business processes. This move towards a more general position of digitisation seeded a focus on the understanding and exploitation of existing data assets. The pandemic has fundamentally altered the timescale in which this digitisation must occur to ensure the long-term productivity of most businesses.

AI is key in the evolution of business from a data-driven to a discovery-driven enterprise.

- Edward Pyzer-Knapp

It has also forced a paradigm shift in the complexity of the processes, which now must also consider a wealth of external information including global health and climate data to guide their strategy decisions. I believe that this transformation from a passive data-driven enterprise to an active discovery-driven enterprise, fuelled by AI and hybrid cloud, is the next major paradigm in the way businesses operate. We call this new paradigm Accelerated Discovery[1].

The information-driven industry will be the early proponents of the discovery-driven enterprise. Leveraging developments in NLP from Deep Learning such as transformers will enable these companies to ingest and structure the constant flow of information faster and more accurately than ever before. From this information, we will see an eruption in the development of ever more accurate digital twins of processes and structures, driven by our ability to use high-performance infrastructure to train and deploy complex AI models at scale. This combination of the 'real'

[1] IBMResearch_STO_2021_Whitepaper.pdf

world and its digital twin will also enable the rapid identification of new 'white space' in business processes and practices. I believe that generative models will begin to help us to collaborate with our digital worlds and drive advances in computational creativity to fill in this white space with high-performing novel solutions, which can then be linked back to the real world through automated deployment by combining AI and RPA. One example we are already seeing is materials discovery. Typically, it takes 10 years and $100M to bring a material to market, but through the combination of data ingestion, AI-accelerated digital twinning, generative models, and automated experimentation we have already seen those figures fall significantly[2].

As AI drives our push for accelerated discovery technologies further onwards, we must remain aware of the potential to accelerate all outcomes, both good and bad. This will mean establishing guardrails and protocols to ensure that responsible advances which provide strong societal benefits are enhanced, whilst alternate applications which do not follow ethical codes of conduct are side-lined. This will require investment into the application of technologies such as explainable AI and bias detection as well as the formalisation of ethical principles for AI development[3].

To conclude, AI is a key component in the evolution of business from a data-driven paradigm to a discovery-driven enterprise, but we must ensure that at all stages of this transformation, we are cognizant of the potential implications of our advances, and pledge to do so responsibly.

Edward Pyzer-Knapp is the Worldwide Lead for AI Enriched Modelling and Simulation at IBM, Visiting Professor of Industrially Applied AI at the University of Liverpool, and the Editor in Chief for Applied AI Letters. He specializes in the interface between AI, HPC, and Scientific Discovery, and has extensive experience in developing best-in-class solutions and deploying them into real-world environments.

PROFESSOR ANDY PARDOE

Founder & Group CEO
AGI WISDOM GROUP LTD • UK

Soon we will have a fully autonomous, AI-driven corporation, with no human intervention or oversight. Maybe this will be the future of Amazon. This could also be the future of smaller businesses, acting as a profit generation for the owners (providing passive income), with little or no need for day-to-day involvement. Imagine a world, where all commercial

[2] For more information on this please read our article (DOI : 10.1038/s41524-022-00765-z)
[3] DARPA's Explainable AI (XAI) Program: Applied AI Letters: Vol 2, No 4 (wiley.com)

activities are fully-automated and autonomous. From product development and innovation to supply-chain and customer support. This vision is not fantasy, it will become reality in the near future.

Other technologies are also helping to make this possible. The blockchain, smart contracts and DAO (Distributed Automated Organisations) are creating the foundation for the governance and administration of these types of automated organisations.

Relying on AI to look after us, might sound idyllic, however, this might be our downfall.

- Professor Andy Pardoe

Imagine a world where the entire process, from product design to manufacturing, is fully controlled by AI and actioned by robots. Even the maintenance of the robots uses predictive algorithms to determine when machinery needs to be taken out of the production line to be serviced. The servicing would be done by other robots too.

Once the products roll off the production line, they enter a fully automated supply chain that has taken the order from the customer. It will then not only deliver the product to the front door of that customer, but also interface with a home help robot to receive the parcel, unpack it, and put it in the correct location in the home. All these are done with zero interaction with a human. Moreover, the customer's need for the product could be predicted by AI so the order is placed automatically to ensure that the product is available when the customer wants it.

While all of this sounds like a perfect utopian world, it has some significant implications. The majority of humanity will no longer need to work for a living. They can dedicate themselves to self-improvement and indulgent activities. Relying on machines and AI to look after us, might sound idyllic, however, ultimately this might be our downfall. Many futurists are worried that this particular future for humanity could cause our extinction, as depicted in a number of films on the topic.

We will need to find balance here so that we gain the benefits of having the machinery and applications in place to make both our businesses and individual lives run more efficiently, but without sacrificing our knowledge and capabilities to control and manage these systems.

Ultimately these technologies will help businesses become more automated and efficient, however, there should always be a layer of humans in the loop for these organisations, making the ultimate decisions, guiding the evolution of new products and services designed, as well as ensuring we have the appropriate level of human governance as the business evolves.

The challenge here is what AI does for the human workplace. What skills and knowledge will employees in the 22nd century need? Will we move from a 5-day working week to a 4 and 3-day working week? Remote working is most likely going to be the norm. Humans will provide more oversight and approval to suggestions coming from the data analytics of AI and machine learning.

Professor Andy Pardoe is an experienced CTO and CEO working across multiple industry sectors including Financial Services, Retail, and Media. Over 25 years expertise of global change programmes delivering complex technical transformations including Enterprise Scaled AI delivery and commercialisation. Also focused on AI Innovation.

SIMON GREENMAN

Partner, Best Practice AI & Member of the WEF's Global AI Council
BEST PRACTICE AI • UK

AI is a general-purpose technology, and as such, the number of use cases where it can add value is only limited by our imagination. The full spectrum of sectors such as agriculture, financial services, healthcare, and transportation already leverage advanced AI. Similarly, a broad array of organizational functions such as HR, marketing, sales, and finance already benefit from AI. In a few short years, it will be hard to conceive of any area of a business that is untouched by AI. Artificial Intelligence will be ubiquitous.

A good way to demonstrate the wide range of use cases for AI is to look at the many tasks an organization does on a daily basis. Any task with some of the following characteristics is well-suited for automation, or partial automation, with AI.

1. Discrete, stand-alone
2. Repeated frequently
3. Similar and repetitive in nature
4. Human intensive
5. Focuses on prediction, optimization, and pattern recognition
6. Has clear inputs and outputs

7. Has lots of data to learn from

The more of these characteristics a task has, the more likely AI can help. This is why AI is already used widely in responding to routine customer enquires via service chatbots, predicting customer churn in marketing, screening resume and CVs of candidates, identifying quality defects of parts on a manufacturing line, predicting supply chain demand, and identifying the signature of abnormal behavior on a computer network.

AI today is not only embedded in the activities of a company but also embedded in the products and services provided to its customers. Automobiles are full of AI, such as automated parking assistance (if not full parallel parking ability). Online services from Google and Facebook are driven by AI algorithms that dictate what search results or advertisements are delivered to a particular consumer. Amazon's Alexa smart speakers use AI to translate consumer commands into understandable text. Apple iPhones use facial recognition to identify users and provide secure access to banking apps. The examples from everyday life are too numerous to list.

AI is not just for Big Tech. The key challenge for organizations not "doing AI" is to put in place the foundations necessary for AI innovation and implementation. This involves six steps from setting out the plan to managing risks:

1. **Plan**. Have a business plan for AI that includes a portfolio of use cases, a roadmap for development and deployment, and metrics of success.
2. **People**. Mobilize leadership, cross-functional teams, and appropriate engineering, data science, project management, and product management skills to do AI.
3. **Data**. Invest in quality, labelled, unbiased and inclusive data to drive AI models and systems – this can often be 80% of the effort on an AI project.
4. **Technology**. Invest in a technology infrastructure that allows for rapid iteration and deployment of AI models and systems
5. **Operations**. Ensure AI is operationalized in a production environment with workforce involvement, customer awareness, and constant monitoring and improvement of the performance of the AI system.
6. **Risks**. Manage the risks of AI.

For many companies, the most topical step is the last one: managing risks. The risks of AI in an organization can be considered in three areas:

1. **Strategic**. If a company does not adopt AI, it may lose out to competitors who have embraced AI technologies. Consumers have seen AI

and data-first companies, such as Amazon, encroach upon many traditional retail sectors. Many new challenger consumer banks are built from the ground up with an AI and data platform that can deliver a more personalized, robust, and satisfying customer experience than banks that are encumbered by legacy technologies.

2. **Operational**. There are myriad operational risks of running AI on a day to day. These range from safety issues if AI fails through to customer services challenges if the AI does not work well, for example, in customer service chatbots.
3. **Legal and Ethical**. These risks include ethical concerns over embodying bias in algorithms, risks over IP ownership, AI supplier contract risks, and the risks of not adhering to existing and new AI regulations.

Ultimately AI is a technology that requires Board of Director oversight as it can lead to brand reputational damage along with significant financial damage if it goes wrong. Organizations need to proactively put in place AI governance frameworks to help identify, manage, and mitigate the risks of AI. Existing enterprise risk management frameworks should be complemented by having clear principles for using AI, having a Head of AI Risk and Compliance, ensuring fair, unbiased, and safe AI by design and default in engineering practices, and providing training across the organisation on the risks of AI.

Simon Greenman is a technology innovator. He is partner at AI management and governance consultancy Best Practice AI. He sits on the World Economic Forum's Global AI Council. He co-founded the early internet brand and AI company MapQuest.com. He holds a BA in Computing & Artificial Intelligence from Sussex University and a MBA from Harvard Business School.

RAJVIR (RAJ) MADAN

Chief Digital and Information Technology Officer
ARCUTIS BIOTHERAPEUTICS • USA

Daily, I am reminded of how complex the healthcare industry is and how nascent we still are in terms of transforming our industry. It still takes the industry several years to discover/develop most drugs (although COVID vaccine development was one positive exception) and patient experiences are sub-optimal (often a slow path to diagnostics for several diseases, mismatch of treatment pathways to patient's determinant of health, etc.).

On the other hand, patient expectations have changed – the consumerization of healthcare is no longer a thing of the past and patients want

to take more control of their healthcare needs. The industry is also collecting more and more data from various sources (think smartwatches or remote patient monitoring devices) that, if leveraged effectively, can disrupt the industry and create huge opportunities to improve patient outcomes.

AI can help identify potential drug targets with a high likelihood of providing positive patient outcomes.

- Raj Madan

So how can AI bring value to the healthcare industry? While the application of AI in Healthcare continues to grow, here are the top use cases:

- **Diagnostics.** AI-based image recognition and matching technology has evolved rapidly over the last few years and has created opportunities for specialists (think radiologists) to validate their findings. In parallel, this technology is already starting to put diagnostics in the hands of patients, for instance, to empower them to initially diagnose skin and eye conditions. As an example of where this is being used, as per Google[4], *"2 billion people around the world have skin, hair, and nail conditions"* and in response to this need, they recently launched the DermAssist tool which allows you to *"find personalized information about your skin concerns after a few questions and three quick photos"* and can *"identify 288 skin, nail and hair conditions"*.
- **Drug Discovery and Development.** Historically, less than 10% of drug candidates commercialize. The entire process can take about 10-15 years and the costs are normally in the millions (if not billions) of dollars. This is where AI can help reduce lead times and costs by (a) identifying potential drug targets (a molecule in the body, usually a protein, that is intrinsically associated with a particular disease process and that could be addressed by a drug to produce a desired therapeutic effect[5]) that have a high likelihood of providing positive

[4] Identify Skin Conditions with DermAssist - Google Health
[5] Drug Target - an overview | ScienceDirect Topics

patient outcomes for a specific disease or therapeutic area and then (b) understanding how patient cohorts will interact and react to that drug target. How do AI models do this? AI models generally parse and search through a vast amount of scientific data sets, look for matches (often by also considering existing drug targets in other therapeutic areas) and analyze/predict their safety and efficacy profiles.

Now, the use of AI doesn't come with a lack of challenges and risks in the healthcare industry. For starters, the healthcare ecosystem consists of several players – think providers (healthcare systems, doctors, etc.), payers (insurance companies), and then pharma and biotech companies (that often perform drug discovery/development activities and commercialize these products). This means that for AI to deliver value across the entire healthcare value chain, frictionless coordination and integration is necessary, which has often been difficult to date.

Additionally, the healthcare industry is highly regulated, and the use of AI to provide and replace human-based diagnosis, although not impossible to achieve, will continue to be a steep climb and one that will always continue to require human intervention in my opinion - until AI models are 100% error-proof. Lastly, the promise of AI lies in its adoption, and with an increased focus and concern about data privacy and security, organizations deploying AI (especially directly to patients) need to ensure that patients have complete transparency to how their data is used, stored, and controlled.

Raj Madan has over 20 years of digital transformation and innovation experience and is currently the CDIO at Arcutis Biotherapeutics. Prior to joining Arcutis, Raj served as the Vice President of Digital, Data/Analytics, and Innovation for Consumer Healthcare Technology at GlaxoSmithKline (GSK), where, in partnership with the Chief Digital Officer, he was responsible for driving significant elements of their Digital and Tech transformation globally.

JOHN THOMPSON

Author, Principal Consultant, CEO
MARKETING SCIENCES, LLC • USA

All functions within a business can and should be impacted by advanced analytics, data, and Artificial Intelligence (AI). Businesses move quickly and conditions change rapidly. AI models and data pipelines need to be built and directly linked into operational systems. I am not saying that AI will make all the decisions and that AI should decide without human intervention, but in the majority of the cases, an AI-driven system can, and in many cases, will be making most of the decisions and will not

require human intervention. The best way forward is Augmented Intelligence where humans and machines work together. We have a long way to go, but we are on the right road today. AI can improve manufacturing efficiency, the retail experience, and many more business functions. Of course, we need monitoring systems and the ability to override and intervene when needed, but we can build AI systems that handle much of the drudgery that people do not want to be involved in on a day-to-day basis.

Explainable AI (XAI) will take our ability to drive change with AI to a completely new level.

- John Thompson

One area where there are exciting breakthroughs is Explainable AI (XAI). The innovations in XAI will enable us to use our most powerful models and modelling technologies in all industries and application areas. Today, in regulated industries like financial services and pharmaceuticals we cannot use Neural Networks in a number of applications due to our inability to examine and explain how the predictions and decisions were made.

Soon, with XAI, that will no longer be the case. We will be able to deploy our most accurate tools and technologies on the most difficult problems in all areas and be able to explain in detail how the models made decisions and predictions. AI with XAI is coming sooner than many people believe and that change will take our ability to make and drive change with AI to a completely new level.

John Thompson is an international technology executive with over 35 years of experience in AI and advanced analytics. John is the author of, Analytics Teams: Leveraging analytics and artificial intelligence for business improvement.

SHARE YOUR PERSPECTIVE SPREAD THE WORD

HOW TO TRANSFORM A COMPANY INTO AN AI-POWERED ORGANIZATION

Transforming an organization requires strong leadership, expertise, and planning – a strategic roadmap and various initiatives that inject AI into its core processes, services, and products.

But which business functions should get the initial focus – in becoming data-driven and AI-powered? Do companies need AI expertise in-house and how long would it take to develop an internal Data/AI function? Could AI/ML models become a source of differentiation for companies? Or they will be provided as commoditized AI services developed and operated by big tech companies?

Jyotirmay Gadewadikar, Richard Benjamins, Harry Mamangakis, Yasen G. Dimitrov, Aruna Pattam, Jordi Guitart, PhD, Dr. Sunil Kumar Vuppala, Dr. Christian Guttmann, and David Pereira share their insights.

How to transform a company into an AI-Powered organization

JYOTIRMAY GADEWADIKAR

Chief Engineer (Enterprise AI – Systems and Mission Analysis)
MITRE • USA

Today it is no longer a question of whether AI-fueled innovation can help drive business value; but how an organization can incorporate change into business operations to drive value, still is. Where do I start, what do I do, how much will it cost, and how long will it take are common questions leaders frequently ask, since a universally applicable AI transformation strategy does not exist. Enterprises are reluctant to begin a significant change journey driven by AI innovation. Enterprises must be aware of the competitors' use of AI and internal AI innovation and operationalization maturity across business functions. Internally, it is essential for enterprises first to measure AI innovation capabilities and potential and then address capability gaps, if any, to become an AI-driven disruptor.

It is naive to think that enterprises' efficiency and efficacy can be improved just by adopting AI tools.

- Jyotirmay Gadewadikar

It is vital to analyze, formulate, initiate, scale, and optimize an organization's AI capabilities and subsequent execution through a well-defined value-maximizing strategic framework and operating model that best fits the organization. It is also naive to think that Enterprises' efficiency and efficacy with using AI can be improved just with the tools and radical approaches without considering sociotechnical complexity and a 'system of systems' perspective. Achieving the transformation, though necessary but not sufficient for advanced technology insertion, is a system of systems (SoS) challenge because this SoS is a collection of independent stakeholders and systems external to organizations. These systems and

stakeholders are external technology vendors, technology integrators, regulatory agencies, internal organizational stakeholders, and systems such as data silos, human resources, information technology, and leadership. Each constituent system's goals are separate and may conflict with the purpose of the transformation itself.

For example, technology vendors are driven to satisfy the shareholders and regulatory agencies to improve compliance. This SoS challenge has stakeholders with competing interests, contradictory objectives between participating systems, unsynchronized lifecycles, desperate management, no clear accountability between the separate constituent systems, and no clear escalation routes. The transformation journey can start with assessing and strengthening the internal capabilities described below first[1]:

- **Strategy and Leadership**. Although cutting-edge technology and talent are needed to drive AI transformation, it is equally if not more important for leadership to align the organization's culture and strategy to support AI. Alignment of corporate strategy with AI strategy and definition of measurable goals and objectives are necessary to prevent disjointed programs
- **Ethics and Governance**. Ethics and governance capabilities are needed to define transparency, explainability, appropriate use of data sources, fairness appraisal, and compliance with regulatory and legal requirements. Technical processes for testing the behavior of algorithms through a quality assurance process play a critical role here.
- **Data**. Data is the seed across all enterprises that allow for AI to scale, and hence, data must be accurate to the extent possible and unbiased to train systems continuously. The more data available, the more the system's learning can advance. Although data availability and quality may seem like a straightforward concept across organizations, even organizations within the same industry still have different maturity levels.
- **AI and Data Science**. The maturity of the data drives better outcomes for an enterprise, but it is the methods and the science applied to the data that help draw meaningful insights and make intelligent decisions. This assessment identifies the use of mathematical modeling techniques and the maturity of the methods and models.
- **Technology Infrastructure**. A solid technical foundation is a critical component for AI transformation and contributes to developing a supporting ecosystem.

[1] D. Simchi-Levi, J. Gadewadikar, B. McCarthy and L. LaFiandra, "Winning with analytics," Accenture, 2015

How to transform a company into an AI-Powered organization

- **Talent**. It is essential to identify how the talent is recruited, developed, and retained and if a skill competency model exists. Enterprises should be capable of adapting and aligning to the new realm that is tying talent operating models to focus on reactive as well proactive recruiting and hiring.
- **Organizational Structure**. A one-size-fits-all organizational structure will not work because of social and technical variations. Instead, AI adoption and transformation can be driven and governed in several ways - from a centralized center of excellence where C-level executives lead a central group to a decentralized one where the structure is entirely independent of various business units.
- **Decisions, Feedback, and Learning**. AI-driven transformation can deliver astonishing results only if AI informs decisions. Suppose the relationships between insights and data are nonlinear, complex, and stochastic. In that case, executives may hesitate to make decisions, so it is essential to ensure appropriate decision-making mechanisms exist. Enterprises have only recently started establishing the necessary infrastructure to collect feedback and incorporate learning mechanisms. It is necessary to include multiple measurements to establish causality and identify the best strategies.

Many organizations plan to operate, upgrade, and transform themselves into data and AI-driven organizations. Still, the challenges may be the complexity within that enterprise, both technical and social. Model-based systems engineering approaches can help manage that complexity. Organizations have leveraged a model-based approach to capture, organize, analyze, and synthesize needs and develop a strategy and road map for execution. The intention is to create a repeatable process to optimize and synchronize AI initiatives and investments through the annual planning and budgeting process.

The interdependence of AI initiatives and the required underlying infrastructure in the investment management process is often ignored. To address it, a business process model that connects the to-be business architecture with technical architecture, can help coordinate the evolution of the necessary infrastructure to successfully transform using AI. The overall approach to developing an AI roadmap is based on gathering requirements, getting some validation, making sure they are aligned, and then communicating them across stakeholders – to ensure alignment. The typical process involves gathering information from organizations' strategic plans and other documents and then engaging with enterprise business lines and service organizations to confirm future strategic business needs and capabilities.

Once the information is captured across the organization, the next step is to synthesize the collected data to develop an organization-wide

strategy which would include a logical grouping of earlier gathered requirements into strategic initiatives and to-be-developed capabilities. For example, increasing employee efficiency can be a strategic initiative, including using AI-driven digital assistants and workforce management. Each strategic initiative can be further executed by developing a multi-year investment plan and a road map.

Your AI adoption strategy, must also consider socio-technical complexities and take a 'system of systems' approach.

- Jyotirmay Gadewadikar

Given that strategic initiatives may cut across multiple businesses lines and operating units, it is important to validate those initiatives across the organization. The next critical step is to align with funding providers and concur on the investment roadmap. Finally, communication with organizational leadership and other stakeholder completes this cycle. This cycle has to be repeated so that the primary inputs from strategic goals are converted into measurable outcomes, business needs are synthesized into strategic initiatives and then executed through a multi-year investment strategy.

Jyotirmay Gadewadikar received the Scientific Leadership award from the US Department of Homeland Security and is an Artificial Intelligence, Decision Science professional engaged in Strategy, Business, People Development, and Thought Leadership. He was previously Chief Product Officer of Deloitte's Conversational AI Practice and System Design and Management Fellow at MIT.

How to transform a company into an AI-Powered organization

RICHARD BENJAMINS

Chief AI & Data Strategist, Telefónica • Co-founder and VP of OdiseIA
BOARD MEMBER OF CDP EUROPE • SPAIN

If you are a conventional company that wants to become a data-driven and AI-powered organization, you will need to embark on a journey: the data & AI journey that forms part of your Digital Transformation. Conventional companies may generate much data as an exhaust of their operation, but are - by design - not ready to collect, store and exploit this data for better decision-making and value creation. The data & AI journey typically consists of different phases including:

- **Exploration**. In this phase, you will explore a few quick & dirty (business) opportunities to assess whether you want to embark on the journey. This phase usually takes from several months up to a year and requires little investment. Once you are convinced that there is value, you will move to the next phase.
- **Transformation**. In this phase, you start to organize yourself to become more data-driven. You will perform an analysis to find the most interesting use cases (applications) to start with, considering *value* as well as *feasibility*. You need to break silos (technological, departmental, vendors) and collect data across the enterprise into a coordinated platform. You will need to hire a CDO and set up a data team, closely collaborating with IT and HR. You will need to work on a 'single version of the truth' so that the whole enterprise understands data in the same way without ambiguity. This phase will require significant investment and may take between three and five years. Be prepared to have some patience as an organization.
- **Data-driven**. In this phase, you will start to enjoy the results of your endurance in the first years. You will be able to use data in a consistent way to inform the big decisions you have to make related to your core business. Moreover, new data-driven products and services will see the light. Depending on your sector, you might be able to externally monetize your data and insights to other sectors in a B2B business model.
- **AI-empowered**. In this phase, you will use Machine Learning and other AI technologies to scale the value of data throughout your organization. Given the massive scale of use, mastering data privacy and AI ethics become essential to create and maintain trust with your stakeholders.

In each of those phases, you will have to make many decisions that will determine how fast or slow you will progress on your journey. Those decisions are different in nature and relate to various aspects of the

organization such as business & finance, technology, people, and responsibility.

- On the organisational side, you have to think about where to place the Chief Data Officer, how to measure data maturity and what will the relation be between the data and the IT department.
- On the business and finance side, you need to make decisions about how to select the best use cases, how to measure economic impact, and how to finance the whole data journey, which can take years. There is no single right answer for the question of what business function to start with. In the beginning, it is important to choose a business function that matters for core business but that is also feasible from a practical perspective. Otherwise, it will take too long to provide the first results. Having said this, many organisations start with marketing.
- On the technology side, important decisions include whether you want to work in the cloud or on-premise, whether you need a unified data model, and you need to define a data collection strategy including planning and budgeting.
- On the people side, it is important to create a team with the right skills and expertise. Many organisations choose a mix between hiring expert personnel and training existing personnel. Sometimes they outsource the first initiatives to a third party to kickstart the activity with the objective to later internalise the knowledge. Other decisions related to people include how to democratise all the data initiatives, how to win over skeptics, and how to make people enthusiastic about data through appropriate communication.
- Finally, on the responsibility side, you need to understand the social and ethical challenges of AI and Big Data; to define AI principles and implement them in your organisation; and work out how you can use data as a force for good, to improve society and fight its challenges.

Dr. Richard Benjamins is among the 100 most influential people in data-driven business (DataIQ 100, 2018). He advises the European Commission and Parliament as well as companies and start-ups. He has a passion for AI for Good and authored three books on AI.

How to transform a company into an AI-Powered organization

HARRY MAMANGAKIS

Chief Technology & Operations Officer
VOICEWEB • GREECE

Transforming into a data-driven, AI-first organization, will always be an ongoing journey since it is part of the overall Digital Transformation Initiative which calls for continuous improvement and adaptation (being agile as an organization). In a typical transformation scenario, which is well documented in the book 'A Data Driven Company' by Richard Benjamins, we can consider 4 general states in this journey:

1. Exploration
2. Standardization & Transformation
3. Data democratization
4. AI-first

What happens in each state, depends on each organization and there is no golden rule for all. This means, that each organization will need to have the tools and metrics to measure its fluency in becoming a data-driven and AI-first company. Measuring its fluency helps understanding where the organization is and what needs to happen next.

Transforming into a data-driven, AI-first organization, is an ongoing journey.

- Harry Mamangakis

In the 'Exploration' state the organization usually attempts to run 'Proof of Concept' (PoC) for specific business (use) cases. Usual candidates for such PoCs are marketing campaigns with the goal to improve their effectiveness or use cases like predicting and reducing churn (where applicable) or increasing average order value by cross-selling related/recommended products. From an organizational perspective, an *owner* is assigned, sometimes also called a *champion*. In general, this is a bottom-up approach, which means that those involved could be some managers that favor experimentation and data enthusiasts – typically data engineers and/or scientists. There is no formal organizational structure in this state – as this is seen as a specific initiative. This process

can repeat over and over with the same or different teams; however, the organization is still considered to be in the "Explorer" state as this happens.

An organization has reached the 'Standardization & Transformation' state when there has been an executive decision that *data is to be treated as a strategic asset*, and that customer value creation must use data in a systematic way. Having reached this point, the organization already has a backlog of use cases to be addressed, and this backlog has been prioritized from the top and communicated throughout the entire organization. Two major initiatives take place during this phase: (a) Data Standardization and (b) Organizational Transformation. To deliver on these initiatives, the following must be there:

- **Data Sourcing Strategy**. This strategy is about what data to use, and where to find it but it also aims to address organizational difficulties where certain functions consider data "their property". Thus, the Data Sourcing Strategy also includes *how data is to be shared* across the organization.
- **Discrete Budget**. In this state, each department needs to explicitly state their data, analytical or BI requirements, and this is to be approved in the annual budgeting process.
- **Formation of a 'Data Team'**. In this state, the organization has realized that a special Data Team is required, with a Chief Data Officer heading it. The positioning of this function within the organization is a whole other discussion, but in general, if we look at successful initiatives, we will see that the CDO is usually placed in organizational structures that are horizontal and apply to the entire business, such as IT or the Digital Transformation Team or under the COO. In some cases, the CDO can also be under the CEO but again, this is a topic for another discussion. The Data Team will consist of data engineers and data scientists and will have worked with IT on the technological choices required to set up the tools to perform the required work.

The next state is 'Data democratization'. An organization has reached this state, *when the use of data is included in the normal decision-making process*, in addition to intuition, experience, or expertise. In this state, the Data Team has matured and has delivered the 'Data Architecture' for the organization. This includes:

- The inventory of data sources which is always kept up to date.
- A data dictionary to ensure a common 'language' and understanding across the organization. In addition, the data dictionary ensures data is traceable to its source, is granular enough and there is only one version of it.

- Processes for data ownership and stewardship to ensure data is of the required quality, always up to date, and available to all those required.

In this state, the organization has already acknowledged that a lot of processes need to be adapted. Hiring needs to change to include skills such as 'use of insights for decision-making'. Training programs need to be introduced to allow existing staff to acquire the new skills required. Training is very important because, in this state, employees are empowered through 'self-service' (processes, tools, knowledge of data, etc) to use data and insights in their daily tasks. Finally, for the Data Team, this state flags the transition from the focus on *data engineers* to *data scientists*, meaning that the number of data scientists at the end of this state, must surpass that of the data engineers.

'Data democratization' has been achieved, when data is used in regular decision-making.

- Harry Mamangakis

The final state is 'AI first'. In this state, the organization uses Machine Learning (ML) and other AI technologies to create value. This is the state where the organization will also reflect back and decide if adjustments are to be made to the data strategy, organizational structure, and so on. Having reached a level of data maturity, the organization will not hesitate to use ML for direct interactions with customers, for example for personalization, product recommendations, or chatbots (NLP); such decisions will be Business As Usual (BAU). The Data Team will be enhanced by AI engineers (e.g., ML engineers, ML researchers) and transformed into a Data & AI Team. Organizational structure is critical in this state. AI talent and resources should not be working in a silo but should be part of agile end-to-end teams, besides the Product Owner and the other roles required to deliver a product or service. This will enable them to contribute their expertise in real-world situations and not operate in a vacuum. This *agile structure* is what makes an organization an AI-first organization.

So how does an organization measure its data fluency in order to assess where it is positioned in the data transformation journey and what

needs to be done next? There are four dimensions required to be measured, to assess the organization's fluency:

1. **The Technology dimension**. What technology choices are made for tools and platforms, how they are used, and by whom. Budgets will need to exist for all these tools, platforms, and so on.
2. **The Data Management & Governance dimension**. This includes data protection, legal compliance, data security (encryption, anonymization, access control, etc). Data management also includes the functions of Data Architecture as previously explained.
3. **The Organizational Dimension**. Changes the organization is making to adopt the use of data and AI.
4. **The 'Business Dimension'**. This measures the adoption of using data and AI throughout the organization to help in decision-making and optimization of processes but also includes efforts (such as R&D efforts) to use insights for new business opportunities.

Measuring each of the above dimensions will help the organization understand where it needs to adapt.[2][3][4]

Harry Mamangakis is a Technology Executive for over two decades, balancing between fluency in technological breakthroughs and having a business mindset. He has led and participated in several transformation engagements for leading brands in industries such as Telcos and Retail.

YASEN G. DIMITROV

Co-founder & Chief Analytics Officer
INTELLIGENCE NODE • UAE

As businesses gear toward a post-pandemic world, AI adoption across industries has rocketed - with the pandemic pointing out glaring inefficiencies in conventional operations and the evolving economy demanding more automation, efficiency, and data-driven decision making. With the accelerating number of AI use cases (from simple image recognition to AlphaFold[5]), a lot of sophisticated AI solutions have become available in the market and continue to evolve. While a lot of companies are successfully leveraging AI at an organizational level and have seen tangible improvements across bottom-line revenue, productivity, and cost savings, the majority have failed at implementing AI at scale. These

[2] The ins and outs of becoming a data-driven organization - Telefónica (telefonica.com)
[3] The AI first Company, Ash Fontana, ISBN : 0593423089
[4] A Data Driven Company, Richard Benjamins, ISBN: 1912555883
[5] AlphaFold - Wikipedia

cases tell us that while transitioning to an AI-led organization seems like the obvious next step, it comes with its own challenges and needs a strategic roadmap to see any success.

Conventional organizations need to first analyze where they need AI applications and if they do need them in the first place. Companies need to understand in what capacity to leverage AI and what the short and long-term goals for AI implementation need to be. The first step to moving toward a data-driven, AI-powered organization is to take stock of the internal data capabilities and talent and start with building a 'Data and Analytics' team.

Artificial Intelligence and Big Data bring a paradigm shift in how we do business.

- Yasen Dimitrov

To do so, companies need to create an ownership function, starting with the Chief Analytics/Data Officer, who will report to the CEO (this is very important because otherwise activities can be steamrolled either by the CFO or COO). As a second step, a *small* team needs to be hired, with the following skillsets:

- Senior Analyst (to collect and analyze company use cases and create internal POCs)
- Data Engineer (to put all data sources in one place - internal, eg. Point of Sale data, competitive data, etc)
- Data Scientist (to start using external services like Google Cloud etc. and build some models to complete the POCs)

Depending on the POC outcomes and the expected ROI, the company may hire more of the above, eg., if more custom models are required, hire more data scientists.

However, finding the right talent and training them can be expensive, let alone building an internal AI infrastructure from scratch. Hence, if you are not a 'heavy data-handling company', using external APIs to achieve your goals is key. Partnering with vendors that specialize in AI solutions that you are seeking can be faster (through API access), less risky, and more economic. But if the cost of the external APIs is too high and there

is no ROI, the probability for you to succeed by internalizing the project is very low as you won't have the right talent. It is therefore important to strike a balance between partnering with the right Artificial Intelligence vendors that have a rapid ROI potential and building your own internal capabilities.

The next step is to identify which business area or function to target first. Rather than trying to introduce AI across organizational processes in one go, it is much more judicious to analyze the simple challenges first and identify the low-hanging fruits where AI can be easily integrated and create significant value.

For example, an AI-driven pricing solution will always start with a simple regression model tested on a small set of SKUs. Once this process is complete, review the ROI and analyze the business impact it can create at scale. It is important to understand that it is an iterative process that requires frequent reviews.

It is important to bear in mind that transitioning to an AI-driven organization takes time. Building AI applications is a continuous process and works better if you follow a test-and-learn approach to identify and resolve problems early on. Setting up and implementing AI processes cannot happen overnight. Finding the right talent to build a focused, in-house team with domain knowledge and then forging relationships with vendors to partner with, can easily take up to a year or even more in some cases.

AI and big data analytics are bringing a paradigm shift in how we do business - helping improve revenue, productivity, market share, and processes across departments and making organizations future-ready. To take full advantage of data analytics and AI, companies need to incorporate these technologies into their vision and core business processes. Companies need to be adaptable, and flexible; they need to realign their culture and conventional processes to make place for AI-driven decision-making.

Yasen Dimitrov is the Co-founder and CAO of Intelligence Node. Yasen has a proven track record of building BI and Predictive Analytics solutions across various industries and has experience in converting the 'value' of data into real, tangible business opportunities. In his current role, he manages operations, analytics architecture, and category expansion; including building and maintaining the largest database (1.2 billion unique retail products) in the industry.

How to transform a company into an AI-Powered organization

ARUNA PATTAM

Head of AI & Data Science - Asia Pacific, Middle East, and Japan
A GLOBAL TECHNOLOGY COMPANY • AUSTRALIA

Artificial Intelligence can help a business adopt a more data-driven approach to decision-making. By harnessing the power of big data and analytics, businesses can gain a better understanding of customer behaviours and preferences, as well as market trends. This, in turn, can help them make more informed strategic decisions about where to allocate their resources and how to best serve their customers. AI can also help a company automate many tasks and processes, thereby freeing up employees' time for more value-added activities. This would not only improve efficiency and reduce costs, but it would also enable employees to focus on more strategic tasks that require human interaction and judgement.

So, how can a company go about making the transition toward an AI-powered organization? There is no one-size-fits-all answer to this question, as the nature of this transition varies depending on the specific needs and circumstances of each business. However, there are a few key steps that all companies can take to get started:

- First, it's important to assess your current business environment and identify areas where AI could be most beneficial.
- Second, once you have a clear idea of the opportunities that AI brings, you can develop a plan for how to integrate it into your operations.
- Third, the business would need to create a data-driven infrastructure, possibly on top of a data lake, a data warehouse, and so on.
- Fourth, the business would need to invest in AI technologies, and find ways to use these tools to improve its operations.
- Fifth, the business should set up a dedicated team to manage and oversee AI initiatives, or develop partnerships with external service providers.
- Sixth, the business would need to adopt a data-driven culture. This means that all employees would need to be data-literate and understand how to use data to make better decisions.
- Finally, it's important to remember that a successful AI transformation requires a commitment from all levels of the organization. Employees across the hierarchy need to be on board with the new strategy and be willing to work together to make it a success.

If a company can successfully make these changes, it will be able to harness the power of big data and become a data-driven, AI-powered organization.

How to transform a company into an AI-Powered organization

But how could a company find the right talent and form a powerful Data & AI team? One option is to hire external service providers who have the necessary skills and experience. This can be a cost-effective way to get up and running quickly, as many providers offer a range of services that cover everything from data management to AI development. Another option is to develop an in-house team of AI experts. This can be a more expensive option, but it gives businesses more control over the development and implementation of AI initiatives. A company could also choose a hybrid approach, which involves hiring external service providers for specific tasks while developing an in-house team for other tasks.

No matter which approaches a company chooses, it's important to make sure that the team has the necessary skills and experience to successfully implement AI initiatives. This includes skills spanning Data Science, Machine Learning, and Artificial Intelligence. When forming a team it is important to ensure that there is a diversity of backgrounds and skillsets, so that a variety of viewpoints can be considered when making decisions about AI.

The time needed to set up a Data and AI function varies depending on the size and complexity of the company. However, a good rule of thumb is that it will take a minimum of 12 months to get a functioning team in place. This includes hiring and training the right people and getting them up to speed on the company's data infrastructure and AI technologies. It's important to note that a Data and AI function will require ongoing development and support, so it's important to set aside adequate resources for this.

AI has the potential to become a source of differentiation for modern products and services. This is because big tech companies are not well-positioned to provide custom AI services to businesses - they are more focused on developing general-purpose AI technologies that can be used by a wide range of businesses. However, businesses need specialised AI services that are suitable for their specific domain and this is where a company's in-house team of AI experts can come in handy, as they can develop custom AI solutions, tailored to the needs of the business.

The benefits of a data-driven, AI-powered organization are clear: increased efficiency, better decision making, and a competitive edge in the marketplace. However, making these changes is not easy, and it requires a commitment from all levels of the organization. A successful AI transformation is a continuous journey. This means that a company needs to invest in AI and get prepared over the long term, and not just in the early stages of the transformation.

Aruna Pattam is a Global AI Thought Leader with 22+ years of experience in data analytics, AI, and ML. Also, a speaker, mentor, blogger, vlogger, and women advocate with the goal of educating and providing awareness on AI to the business and the wider community. The views expressed are personal.

How to transform a company into an AI-Powered organization

JORDI GUITART, PHD

CEO
SCIENCE4TECH SOLUTIONS • USA

Digital transformation brings companies seamless access to an increasing amount of data that is readily available in real-time – where real-time here means as soon as the data is generated is eventually validated. Only companies that have completed this digital transformation are able to tackle the next step with a minimum guarantee of success, that is, injecting AI into its core processes, services, and products.

But AI is a relatively new technology – well, in fact, it can be seen as a sleeping beauty since the late '50s of the twentieth century, until GPU's computing power was capable to shake and awake the AI technology – and it demands an entirely new skill set that it is typically not available in-house. Roles like Data Scientists, Machine Learning Engineers, and Data Engineers, are not so abundant out there, or at least their number is not growing as fast as industries need them.

Additionally, companies need new corporate and senior management roles to encompass the necessary knowledge. Of course, companies can rely on external AI know-how. Consulting companies are out there to provide the basic understanding of what is AI and tell you the dos and don'ts. Externalizing efforts is a well-known way to mitigate risks for those companies that believe AI is not - and could not - become key to their business.

The first thing companies must consider is hiring an AI leader to build a core AI team.

- Jordi Guitart, PhD

For those companies that are successfully evolving through their digital journey though, my recommendation would be to capitalize on AI and the huge volume of data they are digitally gathering, as these will be the two new pillars to sustain companies' long-term value generation. No matter the size of the company – particularly thinking of large non-tech companies – I personally advise against going big in AI from day 1 to any company aiming at getting AI as its main source of differentiation. It's all

about knowledge, not size! So, building a small AI team from scratch and scaling it up upon success will pay off in the long run. Hence, the first thing companies must consider is the hiring of an exceptional AI leader so to let him/her build his/her AI team. No matter the titles the AI team leaders wear, they are the only ones able to attract, motivate and retain AI talent around them. Of course, developing an in-house Data/AI function to create a source of competitive advantage takes time.

Despite that growing an AI team following a slow-paced strategy may sound counterintuitive compared to the overwhelming volume of data at reach, experience tells that AI teams addressing many problem statements in parallel is totally inefficient. AI is the result of forward and backward thinking that is intrinsic to any learning process leading to superior knowledge. So, given one problem statement to solve at a time, I strongly recommend letting the AI team dive deep into the data and only ascend while consolidating the learnings, like scuba divers respect decompression stops. Moreover, the development of good performing AI/ML models cannot easily accommodate software coding sprints so agile methodologies are not deemed appropriate for AI teams to adopt.

Likely, the above explains why big tech companies are acquiring small but highly specialized AI companies to faster merge their teams while incorporating new functional knowledge and thus accelerating AI. The impact of high-growth strategies that big tech companies are experiencing is that they are falling fast into commoditized AI products and services that do not differentiate from each other. The lack of technology challenges could explain why the retention of talent in AI is becoming increasingly difficult and only the lure of large salaries seems to be the alternative that Big Tech companies have, thus creating an excessive turnover and wage inflation.

Jordi Guitart, PhD has been recently appointed CEO of the Barcelona-based healthtech startup Science4Tech Solutions, coming from Aizon where he served as VP of AI. He is concurrently Adjunct Professor of Strategic Management at ESIC Business & Marketing School.

DR. SUNIL KUMAR VUPPALA

Director, Data Science
ERICSSON R&D • INDIA

Customer needs should be at the forefront of everything a company does and, as a result, detecting changes in consumer behaviour is critical to a company's success. Digital transformation helps companies address this by adapting the way business gets done and by creating

entirely new classes of businesses. This change requires the transformation of both internal (employee-facing, operational) and external (customer-facing) functions of the organization. Such transformation can improve existing capabilities, for example, using business insights, building relationships, or improving business processes; it can also lead to the creation of new products, services, and business models. Becoming data-driven and adopting AI is also an essential part of this transformation. Overall, it is the 3Cs – Computing, Content, and Communication – that play a significant role in the transformation journey.

As part of the initiative of transforming a 'conventional' company into a data-driven, AI organization, business leaders must have a solid AI-focused data strategy and roadmap. As part of the roadmap, leaders need to identify the problems, i.e. the pain points of the organization, and also the opportunities, i.e. where there is space for better decisions by using data and identifying actionable patterns. The process of shaping this AI-focused strategy, includes the formulation of an enterprise data architecture, the definition of a roadmap after understanding the current state, and the use of AI playbooks, with an initial focus on proof of concepts that attempt to address the most pressing business problems. It is also critical to develop a strategic data acquisition policy and introduce the right technologies, e.g. a data lake and data management tools, as well as data governance in line with the data and privacy standards.

After demonstrating the value of AI to the senior management by identifying and executing the appropriate set of early use cases, leaders should focus on scaling up the experiments and expanding the team in accordance with the organization's data and AI strategy. However, building AI expertise and the right team is a long process. Thus, companies must identify their core competencies and build their team surrounding those key capabilities. But, at the same time, they should leverage their partners, who can offer expertise on a variety of AI platforms and technologies to accelerate the transformation process.

AI adoption and democratization can be accelerated by using platforms and AI as a Service (AIaaS) from big tech companies such as Microsoft, Google, Meta, and Amazon. Companies can leverage such commoditized AI services to drive the differentiation of their products and services. However, the transformation of the company itself depends on its core business and how the technological change affects its operations and customers.

A 'big bang' approach for such transformational initiatives is known to have failed in delivering the expectations on various occasions. Hence I would recommend that the company follows a step-by-step process,

always emphasizing the demonstration of AI's value and potential to the senior leadership of the organization.[6]

Sunil Kumar is a global thought leader in AI, IoT, and Analytics; Director – Data science; Top 10 Data Scientist in India; ACM Distinguished Speaker; Inventor of 40+ patents, Visiting Professor, Fellow of IETE, IEI; Technical Role Model and IEEE Engineering Manager of the Year awards winner.

DR. CHRISTIAN GUTTMANN

Executive Director, Senior Researcher
NORDIC ARTIFICIAL INTELLIGENCE INSTITUTE, KAROLINSKA INSTITUTE • SWEDEN

Every conventional company has products and services offered to industries ranging as wide as manufacturing, financial services, health care, and transport. These industries change constantly and that change is often driven by customers that expect better quality at a reasonable cost. Companies that adapt to changing conditions and meet customer expectations will continue to grow their business and stay relevant in a competitive business landscape.

Create an open space for transformational possibilities for your company's future.

- Dr. Christian Guttmann

One of the most important success factors of the AI transformation journey is the role of the board and executive leadership, as it is their primary responsibility to set the direction and respond adequately to changing market forces. For many conventional companies, Artificial Intelligence (AI) proves to be a major force as it can transform a market fast. Customers quickly change their demands as emerging AI-driven products and services provide superior quality and performance. The most powerful decision-makers of a company, namely the board and executive

[6] AI-maturity-model-whitepaper.pdf (amdocs.com)

leadership, need to understand the potential of AI - as otherwise the company may face fierce headwinds, and could be outcompeted quickly.

Machine Learning (ML) is one of many AI technologies that has transitioned from research labs into business use. The application of ML has already generated multi-billion USD in consumer markets (e.g., online marketing and content recommendations) and ML is well on its way to transforming traditional 'non-AI' and 'non-data driven' industries. By addressing the following three topics, the board and executive leadership may benefit in making the right decisions.

- *How do AI and data-driven technology change a conventional company's market position?* This is an important question as a conventional company could be directly and immediately affected by competitors that use AI more effectively than they do. The size of this threat needs to be assessed while taking into consideration what is hype and what is real. If a market changes over a longer period then this is often due to long cycles of product introduction in a particular industry (e.g. the pharmaceutical and car industry). In these long-cycle industries, regulations and new frameworks will take time to put into action, but when they are, it often opens a flood gate of new products, and it is important to be ready when that happens.
- *How can AI and data-driven technologies add value to its products, services, and operations?* There are at least three AI value generators in a company: AI can *improve existing products and services*, AI can *create a new portfolio of products and services*, or AI can *make operations more efficient and effective*. Operational improvement is often the initial way for a conventional company to create value as the operations do not affect the main value creation of their business – it is, however, often a more conservative path. In operational practice, AI can improve processes, such as logistics and supply chain, sales, and IT processes much of which resides in ERP and CRM systems. The company's services and products will then be created more effectively and efficiently, hence keeping the company competitive. In this context, it is worth mentioning the well-known 'productivity paradox', where a new technology does not seem to have a measurable ROI, that has an impact on the bottom or top line of the business. However, the introduction of AI keeps a company competitive, meaning that they can continue to build and sell their products and services under competitive operations. Successes might then result in spillovers into the first and second AI value generators mentioned above.
- *How to assess the risk and cost of introducing AI technologies that bring wide implications on products, services, and operations?* It is critical for the company to estimate the risk and cost that comes with AI transformation, as every company has a unique risk profile and will

have an acceptable level of risk and cost. A low-risk approach often results in smaller market shares, less growth, and less future revenue, and that path is often chosen if AI technology is not well understood. Deep knowledge of AI in the board and executive leadership is valuable in predicting market developments and identifying what changes need to take place in the company.

These are tough questions, but as history has shown, companies that understand how new technology changes a market and adapt sufficiently quickly will prevail over those that do not. This has been the case with the steam engine, transistor, internet, and electricity. The need to understand and adapt will be no different for a business in the age of AI.

Each board and executive leadership will respond to the above questions differently, and influence a company's path in their own way. However, there are three considerations that can help to address the above.

- *Ensure that there is deep subject expertise for an adequate understanding of AI in the context of your company's future success.* Executive leaders and board members should interact with the smartest and most suitable AI leaders in the company as much as possible to gain a comprehensive understanding of what AI can do for the business. Such a leader has a PhD in AI and a strong business acumen. This will help to identify which business direction to take and what is technologically feasible. In short, hire the best AI talent that you can get your hands on.
- *Ensure full awareness of the strength and weaknesses of the company's existing product and service portfolio*. Conventional business leadership is often caught up in group think where everyone agrees that "our business has been good so far, so let's do more of the same a little bit better, and we will be fine". This is harmful to a company's adaption of new market conditions. Hence, create an open space for transformational possibilities for your company's future.
- *Create meaning for key stakeholders in times of change*. Initial AI projects or products should aim to be *meaningful*, rather than *impactful*, for as many stakeholders as possible. 'Meaning' here suggests that the application of AI in the business makes integral sense to key stakeholders that are often resistant to organizational change. These stakeholders can be middle management, employees, and shareholders who are sometimes not convinced of the value that new technology can generate (particularly when their own department or job is perceived to be under threat). For those stakeholders, the outcome of an AI initiative needs to meet a wide range of minimal acceptable - and often intangible - criteria. In practice, an initial AI product or

project does not have to be the one that creates the most value, but the one that convinces important stakeholders about AI technology.

Artificial Intelligence is possibly one of the most transformational technologies that humanity will encounter. Transforming a conventional company requires the board and leadership to make a deep evaluation of the company's business and apply the right expertise to identify how AI can drive business value and future success.

Dr. Christian Guttmann is an Artificial Intelligence scientist and global AI executive with a strong track record in creating high-impact business and research with Artificial Intelligence, Machine Learning, and Data Science. Christian was recently named a Top 100 AI global leader by Deep Knowledge Ventures, based on my achievements in science, technology, and business.

DAVID PEREIRA

Head of Data & Intelligence
NTT DATA EMEAL• SPAIN

This question goes beyond the 'typical roadmap' to whether the company has the strategy to achieve this transformation to an AI-powered organization. Such a strategy must be nurtured by advanced technology and data capabilities as the means of achieving this transformation. This strategy is paramount, it is the backbone of the organization for [a] *generating new revenues*, through new value-added services, new business models, and improved customer experiences, and [b] *increasing efficiencies*, through automating processes and scaling services by leveraging data and AI.

Organizations must establish training programs to develop the internal Data & AI talent.

- David Pereira

As the Head of Data & Intelligence in NTT DATA EMEAL, I believe that the best way to support a company in designing such a strategy is through the implementation of an end-to-end framework for Data &

Intelligence. This framework should include the methodologies, organizational principles, and technologies that, when combined, can bring to a company all the needed capabilities, at scale. We use this framework to analyze the organization from a strategic, organizational, and operational point of view - we seek to inject Data, use AI and bring innovation into all key processes, and align the efforts among the different areas.

To become a data and AI-powered organization, the end-to-end transformation must permeate all organizational layers from the C-level and senior managers throughout the hierarchy. Each level has a critical role to play in this transformation journey. For instance, the CEO is responsible for defining the strategic lines and setting the corporate objectives regarding the use of data and AI. The HR team communicates and evangelizes the new 'data mission' as designed from the C-level, and promotes a culture that emphasizes the use of data. The latter can be achieved by introducing training and literacy programs and certifications; or by encouraging access to the open ecosystem to keep up to date with the latest developments and new trends related to data and AI.

In parallel, business functions play a key role, as they are expected to create value through innovative use cases that leverage Data and AI and bring together operational, industry, and technology expertise. To implement solutions that support these use cases, business functions need to collaborate with technology teams – to build and scale data-driven solutions. Ideally, the first use cases to implement should have low complexity but a high return on investment. This helps obtain support from stakeholders and demonstrate the value of the technology and also helps drive the demand and readiness for more ambitious data-driven use cases across the organization.

Besides fostering multidisciplinary teams and promoting synergies between business areas, an AI-driven organization must establish training programs that help develop further the internal Data & AI talent. These programs also help to develop an innovation culture that permeates all areas, departments, and teams. Training also helps in spreading the AI knowledge across teams and enabling the organization to stay ahead of the market trends and experiment with the latest tech advancements. Through such literacy programs, the company can redefine itself as a more attractive organization.

A successful AI-powered organization also needs the right structure. Shaping the right organization is one of the most challenging problems a company can face. At NTT DATA we follow a hybrid model that combines the advantages of a Center of Excellence (CoE) and a Hub & Spoke approach. The Center of Excellence brings together all data and AI developments, defines the AI and Data Governance practices, and connects teams (e.g. Data Scientists, Data Engineers, and Data Analysts). The Hub & Spoke (H&S) consists of putting resources and technological know-

how close to the business areas (Spokes) while the Hub enables the democratization of analytical capabilities.

Another important topic in the AI transformation journey is the opportunity for the organization to differentiate. But how to achieve differentiation from the competition when any company out there can access productized AI e.g. pre-trained industry-specific services such as Natural Language Processing or Computer Vision, etc.? Part of the strategy could be to not only pursue partnerships with leading big tech companies but also to adopt open-source AI frameworks that offer ready-to-use technologies and architectures. This would allow the company to harness these frameworks, use cloud AI services and open-source systems, and accelerate its AI developments. It could also introduce a new working model across the AI lifecycle: data scientists reuse standardized components and refocus from dull, repetitive tasks, to creating real business value through AI.

Once a company has adopted the culture and the technology, it is then necessary to think how they can use that combination to stand out from their competition. Companies should be looking not only to benefit from Cloud, AI, and open-source capabilities to improve efficiency, accuracy, and foster cost savings, but also to *stand out and visibly differentiate* their products from the competition through Data and AI-specific service and product design methodologies.

David Pereira is a telecommunications engineer from the University of Vigo and PDD from IESE Business School. Partner of the NTT DATA Technology Area, David leads NTT DATA's Data and Intelligence in EMEAL as well as the global Artificial Intelligence Center of Excellence for the group.

---

6 DO COMPANIES NEED A CHIEF AI OFFICER?

The increasing importance of Artificial Intelligence in business requires organizational changes and new leadership schemes. We are asking various business and technology leaders questions about the ideal organizational structure to support AI, the corporate role that is accountable for leading the Data/AI transformation of the company, the link between 'data transformation' and 'digital transformation initiatives, and how AI is linked with Software product development and data organizations within a company.

Professor Andy Pardoe, Harry Mamangakis, Agnis Stibe, Richard Benjamins, Jair Ribeiro, Jordi Guitart, PhD, Dr. Sunil Kumar Vuppala, and John Thompson share their insights.

PROFESSOR ANDY PARDOE

Founder & Group CEO
AGI WISDOM GROUP LTD • UK

The need for a Chief AI Officer (CAIO) within a company really depends on various factors. Many of the larger technology consultancy firms are advising their global enterprise clients to consider hiring a CAIO. Especially those that are now building significant AI-based capabilities within their firm.

The reality is that all companies are going to rely on technology much more over the coming years and as technology becomes more capable with the use of advanced Artificial Intelligence, the areas in which technology can augment, support, or even control will increase and cut across all parts of the value chain, both front-office and back-office. Having someone at the executive level who will focus on the AI algorithmic capabilities of the firm is going to be essential soon, both for large corporates but also for SME firms.

Having a CAIO to augment the CTO or CIO is the right decision.

- Professor Andy Pardoe

Ultimately, the algorithms and models that are used in a company to make the core decision-making across the organisation are going to be a significant asset to the firm and potentially create major and sustainable differentiation for its competitive advantage. Therefore, such an important element of the company should have someone to head the development and operational delivery within the executive management of the firm.

Part of the challenge with the role of the CAIO is how it interacts with and supports the CTO, CIO, and CDO roles as there will be overlapping concerns and responsibilities. However, the most significant focus of the CAIO is on developing AI at Scale to support the business strategy.

Fundamentally, as AI is significantly more challenging and different from traditional technology in many ways, having a CAIO to augment the CTO or CIO is the right decision and will provide significant benefit to those organisations that have the foresight to hire one.

For some organisations, the role of CDO and CAIO could be combined with a single person doing both. For others, it may be that the CIO will also cover the AI concerns. The organisation structure may change over time, as the company becomes more mature with its use and understanding of AI.

One fact of technology is the perpetual need to update and refresh, moving towards new approaches and paradigms. Most recently this has been termed *digital transformation* and an element of this, which has been accelerated by the pandemic and remote working is that of moving software applications to the cloud. With many of these technology-focused projects, it can be challenging to demonstrate the benefits to the business. However, this is where data analytics, machine learning, and AI can play a key part to demonstrate additional insights and information that can surface from using the cloud and the linked AI capabilities.

Professor Andy Pardoe experienced CTO and CEO working across multiple industry sectors including Financial Services, Retail, and Media. Over 25 years expertise of global change programmes delivering complex technical transformations including Enterprise Scaled AI delivery and commercialisation. Also focused on AI Innovation.

HARRY MAMANGAKIS

Chief Technology & Operations Officer
VOICEWEB • GREECE

Do companies need a Chief AI Officer? Interesting question. If they need one, is this the same role or a different role than that of the Chief Data Officer? In addition, where is this role positioned in the organizational chart? Is it reporting directly to the CEO (meaning a distance of 1 level or CEO-1)? Is it reporting to the COO or Chief Transformation Officer (CEO-2)? What about reporting to the CIO or CMO? Finally, what type of company are we talking about? Is it an organization that needs to use AI in the process of decision-making, process optimization, etc? Or are we talking about an organization where AI is the core business? In the latter case, a Chief AI Officer is a must. For the remaining of this text, the assumption is that AI is not the core business for the organization.

For a company to become an AI-first organization, it must first go through a transformation journey and become data-driven. Most organizations have started by exploring and testing what benefits they can

achieve by data and analytics; then moved on to more mature states where a dedicated Data Team is formed, possibly under a CDO; then moved to democratizing data and making it available through self-service to everyone (qualified) in the organization, to use as part of their BAU. That is the point in time, most organizations start (or should start) exploring AI.

As AI needs data as fuel, it only makes sense that these two functions (Data and AI) are located close together within the org chart. To take this a step further, while AI usage is in its own exploration phase, it would make sense for AI to be part of the Data function as well. So at least in the beginning, we should be talking about a Data & AI Team.

As this Data & AI Team matures over time and has produced a solid Data Strategy, the balance within the team will change from having Data Engineers as the majority to having Data Scientists, ML Researchers, and Engineers as the majority.

Once the organization has reached the state of data democratization, then it should shift to a more agile structure, where AI roles from the Data & AI team will be rolled into the various product/service teams, operational teams, and so on, to deliver their expertise as part of a team with end-to-end responsibility. However, all this Data & AI talent that is allocated to product and/or service teams still needs to operate as *a coherent unit* to continue maintaining the Data & AI strategy. A role such as the Chief Data Officer/Chief AI Officer can ensure this.

Data & AI talent must operate as a coherent unit to evolve the Data & AI strategy.

- Harry Mamangakis

So, a conclusion one can make at this point is that to help the data maturity process, an organization will need to set up a team of experts, the 'Data Team' with the CDO as their head and have the 'AI Team' grow within this structure. The CDO will also have the AI leadership. This team will drive the organization to a maturity level, where every employee will be able to use data and insights for BAU, in self-service mode.

From that point on, the 'Data & AI Team' will blend with every product or operational team needed. Does this mean that the Data & AI Team is a temporary organizational structure? Well, the same way that a Chief

Digital Transformation Officer is temporary in nature, though active for years (but because of continuous change ends up being a never-ending journey). The Data & AI Team will gradually grow, reach a peak, and then fused into the organization as discussed, but will continue to function as a virtual team.

OK. So where does this function report to? It is best placed somewhere within the organization whose role is horizontal and serves the entire business, to avoid limiting value creation to one specific business domain (e.g., marketing). Therefore, the Chief Data & AI Officer is best placed either under the COO or the Chief Digital Transformation Officer or under the CIO/CTO. This immediately places the role at -2 or -3 from the CEO. It can even be lower or higher depending on the state of the organization in terms of their journey to become data-driven, it also depends on the organization's culture and so on.

PS. IMHO, there is a tendency to have all these new roles (CDTO, CDO, CAIO, and so on) outside the domain of the IT function. To me, this shows a lack of faith that the organization has towards its IT capabilities. A 'Digital-Era' CIO, should deliver on the 'I' in their title. It stands for *Information*, which is based on data. A 'Digital-Era' CIO should act as a catalyst for change, should bring technology to the core of the business by democratizing data analytics and AI, and drive Digital Transformation Initiatives. If an organization does not have such a person holding the CIO position, perhaps it should re-evaluate the position entirely, and call it something like IT Operations Director, placing IT under the COO, and so on.

Harry Mamangakis is a Technology Executive for over two decades, balancing between fluency in technological breakthroughs and having a business mindset. He has led and participated in several transformation engagements for leading brands in industries such as Telcos and Retail.

AGNIS STIBE

Artificial Intelligence Program Director
EM NORMANDIE BUSINESS SCHOOL • FRANCE

Ongoing digitalization is rapidly impacting almost every aspect of organizational performance. Advanced technologies and highly capable algorithms are at the forefront of building a competitive edge for many progressive corporations. Meanwhile, other companies are struggling to keep up and gradually vanishing in this never-ending race for automation, autonomy, and applications of Artificial Intelligence.

It is inevitable. Companies aiming at achieving long-term success will need to upgrade their leadership structures. The major addition should

be an integration of AI as one of the core competencies across all levels of management. There are organizations already assigning this responsibility to their Chief Information Officer, technology strategists, and leading system architects. That can be a viable short-term strategy.

Artificial Intelligence is an explosive catalyst for organizational hyper-performance. It is a superpower emerging from the extensive real-time data mining that boosts high accuracy decision-making. To harness these advantages, companies should have talent that is skilled and capable of mastering the full potential of such digital disruption. Most importantly, this expertise shall be an integral part of the top leadership teams.

The CAIO has the responsibility to link all the core corporate processes to the ever-growing potential of AI.

- Agnis Stibe

Digital transformation has already brought Chief Information or Technology Officers to corporate management. Now it is an obvious natural next step to welcome also a Chief Artificial Intelligence Officer (CAIO) to an organizational leadership structure. This role has the responsibility to link all the core corporate processes to the ever-growing potential of AI to accelerate organizational hyper-performance.

Evolutionary, all companies have been always striving to improve. To deploy innovative tools that foster efficiency, thus empowering their market presence and securing sustainable growth. AI now arrives as an extraordinary opportunity to go much deeper and further than before. It can successfully intertwine the comprehension of microscopic details with their effects on macroscopic perspectives.

CAIOs and their teams will make visible the invisible and will help understand the complexity beyond current imagination. Data is their currency and intelligent algorithms are their currency exchange process. Senior management is their primary customer. Why? Because organizational performance predominantly depends on managerial decisions. Every success or failure is an outcome of a strategic decision.

Do companies need a Chief AI Officer?

AI has only one task. Yes, only one. Helping organizational leadership teams to make optimal decisions. That requires a smooth integration of all the digital components. Starting with data acquisition and ending with hyper-accurate predictive analytics. The CAIO shall constantly maintain and upgrade strategies for ensuring the efficiency of decision-making and value generation through such digital value chains.

Besides technological and algorithmic enhancements, it is crucial for managers to leverage AI for dealing with the all-time biggest challenge of organizational performance: the human factor, or, more precisely, human attitudes and behaviours, and their effects on decision-making accuracy. While technology is increasingly taking over many operational aspects in organizations, human decision-making oftentimes remains at the key crossroads of anticipated success.

The good news is that Machine Learning and Deep Learning, as applications of AI, are already helping organizations to uncover human factors in much greater detail and with more precision than ever before. That helps in locating the main underlying obstacles, often hiding in human minds, thus enabling managers to finally come up with solutions addressing the actual problems, rather than speculated ones.

Combining behavioural sciences with applied AI is the winning strategy for organizational leadership teams. In other words, proper integration of human and artificial intelligence is creating a united superpower for sustainable long-term success in organizations. Not surprisingly, that also applies to broader visions of societal well-being and overall prosperity for mankind. It is the dawn of man-machine co-evolution, a human-computer symbiosis.

Agnis Stibe is the Artificial Intelligence Program Director and Professor of Transformation at EM Normandie Business School and a globally recognized corporate consultant and scientific advisor at AgnisStibe.com. Provides the authentic science-driven STIBE method and practical tools for hyper-performance. 4x TEDx speaker, MIT alum.

RICHARD BENJAMINS

Chief AI & Data Strategist, Telefónica • Co-founder and VP of OdiseIA
BOARD MEMBER OF CDP EUROPE • SPAIN

The best corporate role for being accountable for data and Artificial Intelligence is the Chief Data Officer. The current use of AI is almost completely based on data: without access to quality data, any AI initiative is at risk of delays or failure. Depending on the signal a company wants to give to its employees and other stakeholders, it will have the Chief Data

Officer directly reporting to the CEO (CEO-1), or one of the members of the executive committee, and the C-suite (CEO-2).

In this respect, it is important to which officer the CDO reports. Reporting to IT will lead to a technology-driven activity. If the CDO reports to the CMO, the focus will be mostly on marketing. If it reports to finance (CFO), it will be cross-organisation, but will have less focus on operational use cases. In summary, wherever you place the CDO, there will be advantages and disadvantages. It is important to be aware of those and to take an explicit decision as to where to place the CDO in the organisational chart. When linking AI with software and product development, and with data organisations, it is very important to allow for an agile interaction. The worst a company can do, is to place all those departments distant from each other in the organisational chart, and leave their interaction to a hierarchical escalation process. This will turn any AI and data initiative into a nightmare leading to delays in delivery. Organisations should strive to place those departments close together in the chart so that potential problems can be solved within a week. Moreover, co-locating the different teams in the same physical workspace facilitates collaboration and enhances speed, in addition to creating a sense of belonging to the same team.

Digital transformation and data transformation are closely related. Indeed, it is impossible to perform a data transformation without being digital. After all, a data transformation requires access to data, and non-digital companies do not have such data readily available.

Dr. Richard Benjamins is among the 100 most influential people in data-driven business (DataIQ 100, 2018). He advises the European Commission and Parliament as well as companies and start-ups. He has a passion for AI for Good and authored three books on AI.

JAIR RIBEIRO

Artificial Intelligence Strategist
KIMBERLY-CLARK CORPORATION • POLAND

To remain competitive, companies across almost every industry are looking toward AI, as this technology is becoming ever more strategic. As AI adoption increases and its capacity to deliver competitiveness and optimization becomes critical, companies need to create new functions to meet the new reality. With the Digital Transformation underway, ensuring solid management and activation of AI plays a fundamental role in the business's success. Hence, it is reasonable to dedicate a C-level chair to AI. In general, companies do need a CAIO, especially if they struggle to

understand the numerous opportunities made available by AI – opportunities that can benefit their consumers, teams, and business models; they need an AI leader, capable of leading each step of the delivery and governance of transformational AI initiatives and projects.

The CAIO is responsible for driving tremendous organizational change.

- Jair Ribeiro

The CAIO should be a senior executive who leads the company's AI operations and drives the AI-related business strategy, along with initiatives that transform data into business insights. The AI Officer is typically in charge of a global team that delivers the AI vision, shapes the culture of AI, and helps to accelerate the journey towards a fully AI-driven company. The CAIO must have the capacity to deal with the large number of processes impacted by AI and must be able to apply this technology and the insights obtained to improve the performance of the company as a whole: the CAIO must be an executive who deeply understands the business and can apply AI to improve the way the organization processes information, makes decisions, discovers and fosters new business opportunities - ones that would be invisible without AI.

The role of the CAIO is still not well-defined; we will probably see a more robust definition of the boundaries of the role over the following years. Today, a CAIO is likely a technology person who profoundly understands AI innovation and processes but is not necessarily a data scientist or engineer. Data science skills and solid technical background is necessary, but they are not everything. A good CAIO usually comes with extensive service and consulting experience. Beyond having hard AI skills, the CAIO needs to deeply master *corporate dynamics, be able to demonstrate empathy,* and spread the word about the benefits of AI, while always keeping an effective balance between service and evangelism. The role also requires the ability to architect solutions that leverage AI.

Considering the disruptive nature of AI, the CAIO is responsible for tremendous organizational change and must be able to *share the AI vision,* drive adoption, and connect *business* and *technology* teams across the organization. Taking into consideration the increased collaboration and cross-functional work required for an effective AI strategy, being a

champion of building relationships is vital for a CAIO – to ingrain innovation throughout the organization. In fact, the CAIO role is strongly *innovation-driven* and requires *a deep business understanding* and the capacity to identify opportunities to add value with AI and create differentiation.

The CAIO must be able to share the AI vision, drive adoption, and connect teams.

- Jair Ribeiro

Having a CAIO can be a source of qualified insights and skills that can contribute to the evolution of the company. But, to make it a sustainable success, all of the C-level executives should be involved and engaged in supporting the AI-first strategy. The AI journey can only be successful if there is a real commitment from the operational and cultural point of view. Both innovation and delivery teams must understand and communicate the business challenges and the opportunities that AI brings. A successful AI function is not isolated. Instead, AI must become a core enabler for the company – improving its operational aspects, decision-making, services, and products. Appointing a CAIO sends a solid message to the company about the focus on AI, and increases the confidence that the company is moving fast towards an AI-driven culture.

Jair Ribeiro is currently the AI Strategist of Kimberly-Clark Corporation. He inspires and supports the company to innovate through the power of AI.

JORDI GUITART, PHD

CEO

SCIENCE4TECH SOLUTIONS • USA

Almost all industries have initiated their digital transformation aiming at getting access and traction on the apparently unlimited amount of available data that their markets can generate - with unequal success I must say. Many large companies have reached their non-return point on their digital transformation journey, although a non-negligible number of

analogic / non-automated processes still try to survive, certainly led by their demonstrated efficiency but sentenced to disappear in a dead-lock game. There is no point in long-term analogic/digital coexistence and companies' CEOs are perfectly aware of that. However, 'classic' C-suite teams might not cover all necessary competencies to address what is next and beyond.

CDO and CAIO should not be interchangeable roles.

- Jordi Guitart, PhD

Indeed, there is an increasing need for new corporate roles to lead the Data/AI transformation of the company that comes so naturally with Digital Transformation. In that sense, Chief Data Officers (CDO) – a relatively new corporate role – are consolidating their position in large data-driven companies, while a Chief Artificial Intelligence Officer (CAIO) is not a so common corporate role, and it is often perceived as the CDO alter ego. Why?

Let me state it clearly: CDO and CAIO should not be interchangeable roles. Whilst a CDO oversees all data a company owns and has within reach, including the use of software tools for advanced analytics to gather rich insights among the data, a CAIO is responsible for the software development of actionable drivers in real environments to generate business outcomes. Hence, a CAIO must be an interpreter of the surrounding reality to connect and transform all kinds of data, whether formatted or unformatted (e.g., text, image, or sound), from transceivers and smart sensors over software, electronic and mechanical interfaces, to interact with for value generation. In other words, a CAIO is aimed at inspiring and leading the design of products and/or services that are acting automatically from learnings captured in real-time as an intrinsic property of the product/service – or near real-time, as the fast-paced business world demands. Acting means impacting (i.e., altering) the surrounding environment, thus leading to learning and acting again and again. And yes, 'acting automatically from learnings' brings ethics as a critical dimension to the CAIO's role. Any CAIO must relentlessly show and spread a high sense of ethics as their implemented AI technology approaches human conduct on behaviour and decision-making, thus affecting all stakeholders.

From the above, the reader might still think that it is hard to argue the need for a CAIO position at the corporate level in companies still facing their digital transformational journey despite AI is gaining momentum… and I partly agree. However, the final thought of these lines aims at pointing to the core of any business strategy, that is, the ability to generate long-term value. In this sense, I can only refer to Michel Porter's theory of Competitive Strategy that I have been able to apply as an executive throughout my professional career in technology companies, and that, as an adjunct professor, I proudly try to teach my Executive MBA students. Therefore, the CAIO is a must-have corporate role in any company that has AI as their main source of competitive advantage. Moreover, if these companies apply AI technology to compete in highly regulated markets, the CAIO role is not even an option.

Jordi Guitart, PhD has been recently appointed CEO of the Barcelona-based healthtech startup Science4Tech Solutions, coming from Aizon where he served as VP of AI. He is concurrently Adjunct Professor of Strategic Management at ESIC Business & Marketing School.

DR. SUNIL KUMAR VUPPALA

Director, Data Science
ERICSSON R&D • INDIA

AI is going to transform every industry, beyond the software sector, and contribute to GDP growth. I strongly believe that companies need a Chief AI Officer, a role that should be reporting to the CTO or the CEO directly. As AI is impacting Digital Transformation, roles like Chief Analytics/Digital/Information/Data Officer may evolve or transform into this new, crucial role, leading the company's data and AI transformation: the Chief AI Officer (CAIO).

The CAIO is responsible for driving the AI roadmap of the organization, handling the AI infrastructure, designing the development and deployment strategies, and establishing a framework for ethical and trusted AI. The CAIO should be technically strong e.g. able to understand how to build AI business models or platforms, how to design end-to-end data science pipelines for development and deployment, and, in general, be on top of the technological advancements in AI. The CAIO is also responsible for designing the technical stack and understanding how AI can become an internal part of the products/services of the organization. He/she should come up with optimal AI infrastructure by partnering with cloud AIaaS providers and specialized companies. Moreover, as AI is going to play a vital role in upcoming digital technologies such as edge

devices, web 3.0, and the metaverse, the CAIO must be able to come up with innovative use cases in the context of these emerging technologies.

In parallel, this role requires strong business acumen, and the ability to bring alignment among stakeholders and drive results. The CAIO is not isolated - a successful CAIO understands deeply how AI is linked with the software, product development, and data organizations within the company, and demonstrates strong stakeholder management skills, with both internal business units and external customers.

The CAIO should be the agent of change - checking the AI maturity in terms of strategy, data, technology, and operations.

- Sunil Kumar

The CAIO should be the agent of change - checking the AI maturity of the organization in terms of strategy, data, technology, and operations and coming up with the right transformational steps. The CAIO should be able to drive Digital Transformation initiatives, explain how to use AI for Intelligent Automation, and drive the definition of the data-driven AI strategy for the business units in alignment with the organization's vision.

To achieve all the above, the CAIO needs to build a strong team, able to design and execute the data-driven AI strategy. This AI team will be responsible for the end-to-end development, deployment, and operationalization of AI models. Beyond the AI team, the CAIO must also empower other teams in the organization to understand the value of AI for their business. Finally, an important responsibility of the CAIO is to ensure that AI meets ethical and privacy requirements according to policies such as the EU General Data Protection Regulation (GDPR), California Consumer Privacy Act (CCPA), and California Privacy Rights Act (CPRA).

Sunil Kumar is a global thought leader in AI, IoT, and Analytics; Director – Data science; Top 10 Data Scientist in India; ACM Distinguished Speaker; Inventor of 40+ patents, Visiting Professor, Fellow of IETE, IEI; Technical Role Model and IEEE Engineering Manager of the Year awards winner.

JOHN THOMPSON

Author, Principal Consultant, CEO
MARKETING SCIENCES, LLC. • USA

Companies certainly need an executive at the C-Level who is focused on how to use data and how to employ, leverage, and deploy advanced analytics.

The common mistake organizations make, and I call it the Original Sin, is that they put the analytics team in the IT organization. In such cases, they have failed before they started, but they don't know it. They will hire data scientists and analytics professionals, and force those new hires to use unproductive and ill-suited approaches like Agile and they will manage the team like a bunch of project managers rather than the intelligent, creative, and free-thinking individuals that they are. The Chief AI Officer needs to be driving change and improvement through data and analytics, which can only be done at the C-Level. The data and analytics team needs to report into the CEO, COO, or to the Strategy function. Reporting into Finance, IT, or Operations is a mistake. Data and analytics teams need to drive data and Digital Transformation.

The common mistake organizations make is that they put the analytics team in the IT organization.

- John Thompson

To be successful, the Chief AI Officer (CAIO) needs to be a forward-looking position. I see a number of Chief Data Officers (CDOs) being appointed and those positions are inevitably focused on reducing technical debt, integrating old systems, and migrating systems to the cloud. All are interesting tasks and are required for the ultimate success of the company, but none of them have much, if anything, to do with analytics and using data to drive change and improvement. The CAIO needs to be focused on analyzing internal and external data and illuminating how the organization can do a better job of – e.g. setting and managing the optimal price for every product and service offered, determining the optimal

offers for each customer, donor, or patient, understanding the current efficiency of each store in the network, determining where to put the next store, or customer service center, or how to retain key employees on a global basis.

Analytics can be a game-changing tool and capability. We need the leaders like the CAIO looking to and charting the future rather than managing the costs of the past. At this point, the majority of leaders in the CDO, CAO, and CAIO roles are not driving change the way that they should. This is an opportunity missed.

John Thompson is an international technology executive with over 35 years of experience in AI and advanced analytics. John is the author of, Analytics Teams: Leveraging analytics and artificial intelligence for business improvement.

JOIN THE DISCUSSION **SPREAD THE WORD**

WHAT IS THE ‘AI STATE OF THE ART’?

Artificial Intelligence is evolving at a rapid pace. In this chapter, we attempt to present the latest and greatest technologies, models, and methods available in the broader AI space.

We asked various leaders to summarize the state-of-the-art of AI regarding aspects such as Computer Vision, Natural Language Understanding, Content Understanding, Decision making, and Robotics. We asked the leaders to describe the latest advances, list the most important open-source AI frameworks and models, and explain how Quantum computing is expected to boost the power of current AI technologies.

Eva Agapaki, Dima Turchyn, Emma Duckworth, Netanel Eliav, and Mike Tamir share their insights.

EVA AGAPAKI

Artificial Intelligence Assistant Professor
UNIVERSITY OF FLORIDA • USA

AI equips machines with various types of human capabilities such as the ability to sense, see, make decisions and react. AI has seen tremendous hype and investment both in academia and industry, becoming a research hotspot in multiple disciplines, with the most obvious ones being technology, finance, marketing, and autonomous vehicles, but it has also gained traction and is rapidly emerging in healthcare, law and design disciplines. AI is not a new concept; Warren McCulloch and Walter Pitts invented threshold logic in 1943 by creating a computational model for neural networks based on mathematical concepts and algorithms[1]. The enabling drivers of AI technologies are the large amounts of high-dimensional data and advanced machine learning algorithms that automatically recognize patterns in order to make informed decisions.

The next breakthroughs will give machines the possibility of surpassing human senses for the better of humanity.

- Dr. Eva Agapaki

There are four machine learning categories that these algorithms fall into: supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning (RL). Supervised learning algorithms rely on large amounts of manually annotated data to learn patterns, in order to make predictions on new, unseen data. The most common supervised learning algorithms are classification and regression. Classification algorithms predict the output of discrete values, such as judging whether a photo depicts a dog or a cat, whereas regression algorithms predict the

[1] McCulloch, W.S. and Pitts, W., 1943. A logical calculus of the ideas immanent in nervous activity. The bulletin of mathematical biophysics, 5(4), pp.115-133.

output of continuous values, for example, house prices based on historical data. On the other hand, unsupervised algorithms do not have a requirement for labelled data, given that data can be clustered based on their characteristics, for example, customer segmentation based on their preferences. Semi-supervised algorithms only need a few labelled training samples and they are currently a hot research area. Some examples include self-training, active learning, and graph-based semi-supervised learning. RL algorithms rely on a rewards hypothesis by selecting actions that maximize the total future reward. Those actions may have long-term consequences and the reward is not immediate. Some examples include financial investments and scheduling optimization of capital projects. OpenAI Gym[2] is an open-source platform for developing RL algorithms.

One of the most prevalent sub-fields of AI is Natural Language Processing (NLP), where computers recognize and understand human text language. There are multiple NLP applications such as sentiment analysis, information extraction, and machine translation. Transformers (such as GPT-3[3] and BERT[4]) have been widely used. Transformers have a looping mechanism that acts as a highway to allow information to flow from one step to the next and retain sequential memory. This makes them ideal for speech recognition, language translation, and stock predictions. State-of-the-art (SOTA) transformers for NLP tasks can be found at multiple open-source Github repositories[5].

Another AI sub-field is Computer Vision - often referred to as Machine Perception. The goal of these algorithms is to enable computers to visualize the world as humans do. The most widely adopted Computer Vision problem is object recognition in 2D and 3D data. There are multiple tasks involved with image recognition such as image classification, segmentation, and object detection. Some of the most commonly used Convolutional Neural Networks for object detection are YOLOv3[6] and deep residual networks[7] with the SOTA being Vision Transformers[8].

Decision-making is another complex task that involves data analysis and merging information from disparate sources while leveraging information importance. AI can solve competitive human-level tasks and even

[2] Gym (openai.com)
[3] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A. and Agarwal, S., 2020. Language models are few-shot learners. Advances in neural information processing systems, 33, pp.1877-1901
[4] Devlin, J., Chang, M.W., Lee, K. and Toutanova, K., 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
[5] For example: https://github.com/huggingface/transformers
[6] Redmon, J. and Farhadi, A., 2018. Yolov3: An incremental improvement. arXiv preprint arXiv:1804.02767
[7] He, K., Zhang, X., Ren, S. and Sun, J., 2016. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 770-778).
[8] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S. and Uszkoreit, J., 2020. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.

beat humans, for example, when AlphaGo defeated the world chess champion using RL algorithms.

Despite the advances in every AI sub-field, there are significant challenges to overcome. Some of these include explainability of the developed models (technical challenge), algorithmic bias (technical and societal challenge), and transparency in usage (societal, political, and legal challenge). Quantum computing can assist in mitigating some of these obstacles. It can be used to rapidly train and generate optimized ML algorithms (using superposition and entanglement). A recent open-source library for quantum ML is Tensorflow Quantum (TFQ)[9] which combines a suite of quantum modelling and Machine Learning tools. Some contributions of Quantum AI are: quick and optimal weight selection of neural networks, faster encryption based on quantum search, and quantum algorithms based on the Hamiltonian time evolution that represent problems with optimal decision trees faster than random walks. A summary of AI models, papers, datasets, and open-source libraries can be found at Stateoftheart.ai[10].

The evolution of AI spanning across disciplines has been inspired by advances in biology, cognition, and system theories. The next breakthroughs will not only give machines more logical reasoning capabilities but also the possibility of surpassing human senses for the better of humanity.

Dr. Eva Agapaki is the Director of the Digital Twins research lab at the University of Florida with experience in applied machine learning projects in academia and industry.

DIMA TURCHYN

Artificial Intelligence Product Marketing Lead, CEE Region
MICROSOFT • CZECHIA

During the latest several years, I've witnessed various interesting trends in AI systems available on the market that solve real-life use cases. A new generation of models, especially transformer-based models, re-shaped the foundation of what is possible in areas like Natural Language Processing (NLP). Then, as generally a bigger number of parameters in those models showed improvement in model performance, the competition of 'who's model has more billions of parameters' emerged which in fact did work. As an example, models like the Megatron-Turing NLG model with 530B parameters not only push SOTA for

[9] TensorFlow Quantum
[10] Stateoftheart AI

some tasks but do so across a broad set of tasks including completion prediction, reading comprehension, commonsense reasoning, natural language interfaces, and others.

The cycle from new model introduction to implementation significantly shortens.

- Dima Turchyn

And while this does not mean we are any close to general AI systems, there is a higher level of generalization that models extract when compared to smaller, task-specific models. Those large models come at a cost though, since they require a large pool of resources and are only becoming popular due to the availability of powerful optimization and distribution learning algorithms. For example, the abovementioned Megatron-Turing NLG model leverages the DeepSpeed library, allowing to build pipeline parallelism to scale model training across nodes. By the way, cost (including environmental) to train such large-scale models means that effective re-use of those models across use cases is key for gaining positive net value. And, at the same time, those same powerful libraries, next generation algorithms, and cloud resources are available to literally anyone: as an example, using those same libraries and the cloud ML platform allowed one of our customers – the University of Pecs in Hungary – to train Hungarian language model in just several months from idea to production, with total resources cost of just around $1000.

Another interesting observation is the emergence of models which are multi-task, multi-language, or models that are built for one domain of AI tasks and then are successfully applied to another. For example, applying attention models initially developed for NLP to image recognition tasks shows impressive results. And even more so, new models emerge which learn simultaneously on different types of modalities – or multimodal architectures. Most prominent are probably recent architectures that use a combination of language and image data – they learn from both the objects in an image and the corresponding text - leading to knowledge that can be applied to a range of tasks, from classification to generating image description or even translation or image generation.

This paradigm of combining narrow learning tasks into a more general model which learns on many tasks simultaneously is also leveraged

in the new generation of Language Models. As an example, Z-code models take advantage of shared linguistic elements across multiple languages to improve the quality of machine translation and other language understanding tasks. Those models take advantage of both transfer learning and multitask learning from monolingual and multilingual data to create a language model, improving tasks like machine translation by a significant margin across languages. This same approach is used in Florence 1.0 model, which uses XYZ-code, a joint representation of three cognitive attributes: monolingual text (X), audio or visual sensory signals (Y), and multilingual (Z), which allowed to advance multi-task, multi-lingual Computer Vision services for tasks like zero-shot image classification, image/text retrieval, object detection, and question answering.

There are of course many other models and developments, which push the boundaries of what is possible with ML models. Working with the customers and partners on their AI projects, what I am always looking for is how we can apply all of that state-of-the-art research to real-life customer projects. Using a large-scale model has its challenges from its size to inference costs and many approaches emerge to build sparser, more efficient models. As an example, the abovementioned Z-code models use a 'mixture of experts' approach, which means only a portion of a model is being engaged to complete a task. As a result, customers can make use of those powerful developments almost immediately after its introduction. Customers can today build applications leveraging Z-code models, or use multilingual language models with Cognitive Services APIs, or even apply powerful large-scale models like Open AI as a managed cloud service.

In general, this is probably the most impactful observation I record for myself in the latest couple of years: not only new models emerge, but those new generation models are also significantly decreasing the cycle from introduction to actual, applied implementations. This has its obvious benefits, but also imposes significant risks. For example, the latest speech generation services are available to anyone as easy-to-use APIs along with many other AI services now reaching human parity. This general availability increases the need to take responsibility for how those services are being applied. This is a separate topic by itself, as approaches to address those risks span across cultural, technological, and policy dimensions, like gating some of the services and reviewing the use case each time those services are being deployed, helping to ensure that those powerful advancements are applied not only where it can be used, but where it actually should be.

Dima Turchyn is working with analytical technologies and Machine Learning for 20+ years and has a broad background in Business Development, IT and Marketing. Most recently, he leads Microsoft's AI product marketing for CEE region covering 30+ countries.

EMMA DUCKWORTH

Director of Data Science
GSK CONSUMER HEALTHCARE • UK

When thinking about advanced, real-world applications of AI, I'm most excited by technology that is making a tangible impact on decisions or processes. Particularly those that are making sophisticated recommendations to a degree that hasn't been possible historically by people or other methods. One such cutting-edge application is the use of AI to optimally orchestrate large complex systems, in-particular global supply chains. Supply chains are systems that are not just complicated, but complex as they have many individual legs or components with contradictory incentives. To achieve overall system optimisation, we require AI at every stage of the solution, from process mining that aims to understand the system, to running our final end-to-end simulations and optimisations.

Quantum computing holds great potential and could cause a real step-change.

- Emma Duckworth

This innovative AI is moving the dial in the way we run supply chains and find solutions to important global challenges such as supply chain resilience and sustainability goals. Such AI models are generally bespoke compilations of elements that are themselves state-of-the-art. Much of the active innovation being worked on here is still in development, often featuring collaboration between academia and industry. This is due to the scale of the problem, the uniqueness of supply chains, and the required access to data and computing resources.

An example of a component where advancements are being made is process mining. Here we may use commercial tools to understand how the system is configured as a starting point for our model. We may then pull on another active area of AI, time-series forecasting. ML models are starting to outperform statistical methods when it comes to accuracy and performance in real-world applications. In the M5 forecasting competition, there was compelling evidence of this when LightGBM (a

decision tree-based machine-learning model) had widescale adoption by most leading entries. Additionally, in complex systems we are often considering multiple forecasts, solutions such as LightGBM that support hierarchical forecasting are therefore beneficial.

Other AI technologies whose outputs may be included in complex systems are Deep Learning applications such as Computer Vision or NLP. Here we see rapid innovation on a different trajectory. The scaling of models and available training sets that have been made widely available by technology companies such as Google and non-profits such as ImageNet, has led to commoditisation of Computer Vision applications. There are now numerous off-the-shelf solutions that can be applied to real-world problems with relatively small amounts of customisation. This is feasible via Transfer Learning, that is, the process of using a pre-trained model as a starting point for a model for a new task. For example, your starting model may be able to identify dogs in a photo, you can then additionally train it to identify your pet among others, similarly to how Apple and Meta can start to quickly recognise your friends in photos. For example, say we want to use Computer Vision to measure quality on the manufacturing line, we can very quickly build a good proof of concept using offerings such as Microsoft's Cognitive Services or Google's AutoML.

We then need to use optimisers to build our scenarios and solve the problem. Here again, there are many different solutions available. Searching the problem space and optimising for solutions in complex systems is a big problem and can quickly become computationally expensive. It's an area where Quantum computing holds great potential and could cause a real step-change, in particular Quantum Annealing. This refers to an optimisation process for finding the global minimum of an objective function, particularly effective when there are many local minima, such as across our complex supply chain.

Finally, supply chains are also areas where developments in the governance of AI, accountability of model recommendations, and model transparency and explainability are very important. Unlike many digital/web applications of AI, we are faced with the challenge of how to 'start small' when innovating & testing a model that may e.g. deploy a shipping container full of products. Supply chain scenarios require high confidence in models when applying recommendations to a real-world system for the first time. The implications of errors can be significant, particularly in highly regulated supply chains where there are safety requirements on the quality of the output. It is therefore critical that academia & industry work closely to develop frameworks so that business and regulators are in-step with the cutting-edge algorithms. Collaboration is a massively important aspect of state-of-the-art AI innovation. Only by working with

regulators we can deploy state-of-the-art AI and realise the potential of this incredible technology.

Emma Duckworth leads the global Data Science team at GSK Consumer Healthcare. She is passionate about AI ethics, diversity in data science and uses AI to solve big, strategic problems such as accessibility & sustainability of everyday healthcare. Excited by innovation, she pulls on her startup experience to build and scale AI products.

NETANEL ELIAV

Chief Executive Officer & Technology Development Specialist
SIGHTBIT LTD • ISRAEL

The answer is not straightforward. First, we need to define what we mean by state-of-the-art (SOTA). In simple terms, SOTA refers to AI at its best. This is when the AI has reached its full potential in terms of performance and capability. The definition of AI SOTA changes with time and with the advancement of technology. For example, in 1997, IBM's Deep Blue was considered to be AI SOTA.

AI on Quantum computers will make breakthroughs in solving some of the world's most pressing problems.

- Netanel Eliav

But today, a computer beating a human player at chess would not be considered AI SOTA anymore, because computers and the AI field have since surpassed that level.

There are many AI technologies in the market today, but the most advanced ones are based on Deep Learning and Machine Learning algorithms. Deep learning has been around for a few decades, but only recently has it had a tremendous breakthrough in accuracy and precision. One of the most fascinating aspects of Deep Learning is that it can

process data in a more human-like way. The most advanced Deep Learning networks available today are:

- **Convolutional Neural Networks** (CNNs) are algorithmic architectures that have been used in everything in our lives, from *facial recognition* to *image classification* and even at the backend of simple object tracking technologies that can be found in our phones. Companies like Google and Facebook use it in many products to provide users (and themselves) more value.
- **Generative Adversarial Networks** (GANs) are algorithmic architectures that use two Neural Networks. They put the two networks against each other in order to generate new, synthetic instances of data that can pass for real data or as real data. The most popular example of its use is Deep-Fake, which can manipulate videos so that they look real and are hard to distinguish from reality.
- **Recurrent Neural Networks** (RNNs) and the sub-type Long-Short Term Memory Networks (LSTMs) are a special type of artificial Neural Network adapted to work for *time series data* or data that involves *sequences*. Those types of algorithms consider the missing dimension of time in CNNs and keep the data's history. The most popular example of its use is in *autonomous cars:* as they use feeds of data from sensors to navigate through traffic and avoid obstacles on the road, Object Detection by itself is not enough.

Startups and big companies fine-tune those networks to build more specific models to be used by many sectors, such as the military, autonomous cars, smart cities, and more. Many previously unsolved tasks in the fields of Natural Language Understanding, Computer Vision, and Robotics are now solved by those algorithms.

Those technologies have been around for a while, and are available for free to anyone who needs them - thanks to open-source libraries and frameworks like:

- **PyTorch** - was created by Facebook and accessible on GitHub.
- **Caffe** - has been funded by Berkeley Vision and Learning Center (BVLC).
- **TensorFlow** - was developed by Google.
- **Keras** - was developed as part of the research project and maintained by François Chollet.
- **Detectron2** - which was developed by Facebook with the help of the tech community, including the author of these lines.

The main limitation of AI is on the hardware side. Researching and developing new models takes a lot of processing power from GPUs and CPUs with TPU limitations.

There is new progress in Quantum Computing that may bring good news and it will have an enormous impact on the future of AI and other fields. The development of quantum computing will also help to advance Artificial Intelligence because it will allow for more complex simulations and algorithms to be run. Soon, quantum computers and AI will be used together to make breakthroughs in solving some of the world's most pressing problems.

Technology is advancing at a rapid pace. The advancements in AI are making it possible for machines to learn, perceive, and understand the world around them. The future of technology will be amazing.

Netanel Eliav is a CEO and Founder at SightBit - An Artificial Intelligence Startup Using Deep Learning and Computer Vision to Save Lives. He is a former Product Manager and ex-Technology Tech Lead Specialist at the Office of the Prime Minister.

MIKE TAMIR

Chief ML Scientist, Head of ML/AI
SUSQUEHANNA INTERNATIONAL GROUP • USA

I am consistently impressed with how far we have come in natural language understanding and processing more generally. Language is fundamental to certain kinds of understanding and while the mega-sized language models developed in recent years are (in the end) language models and do not necessarily show a true understanding of the text, we have come a long way from the naive picture of machines only being able to process language through rules and heuristics caricatured by philosopher's like Searle and his 'Chinese [translation] room.'[11]

The biggest advancements in language have come from a shift from the direct embedding of text tokens to more inductive methods of embedding text in context. Over the past several years similar advancements have been made in understanding data that can be coded as graphs, relating different entities (nodes) to other entities. The coevolution of these parallel application areas has not yet been fully explored and heralds a very productive line of growth that we can expect in the future. While these research advances have yet to come to fruition, for significant practical applications in medicine, research advances in

[11] Chinese room - Wikipedia

protein folding and genetics, they have the potential for one day making this a reality.

Another area of remarkable improvement is Reinforcement Learning (RL). While a lot of the fundamental optimization paradigms in RL have remained unchanged, our ability to build estimators that guide how RL agents navigate, understand, encode and then evaluate their behavior in an environment has dramatically improved with the benefit of Deep Learning research over the past several years. This research has the potential to solve the major hurdles that still exist in practical applications, ranging from adaptive safety tests for self-driving cars (which operate in very complex high stakes environments) to tactically responding to fake news threats, and more.

Mike Tamir, PhD is a data science leader, specializing in deep learning, NLP, and distributed scalable machine learning. Mike is experienced in delivering data products for use cases including text comprehension, image recognition, recommender systems, targeted advertising, forecasting, user understanding, and customer analytics. He is a pioneer in developing training programs in industry-focused machine learning and data science techniques.

THE IMPACT OF AI ON SOCIETY AND EVERYDAY LIFE

Artificial Intelligence is changing the world and the impact will be massive: on the way we work, we live, collaborate, decide, and act as a society.

But how can individuals and societies benefit from AI? What are the global problems that can now be addressed more effectively by leveraging AI? What are the risks of AI at the societal level? What is the 'technological singularity' and how could it affect us? How should individuals and societies get prepared for AI?

Michael Wu PhD, Nell Watson, Anthony Mills, Alf Rehn, Dr. Marily Nika, Nazar Zaki, Angeliki Dedopoulou, Jon Skirnir Agustsson, and Boyka Simeonova share their insights.

MICHAEL WU PHD

Chief AI Strategist
PROS INC. • USA

Throughout the first three industrial revolutions, humans have learned to leverage machines to automate various tasks. We use machines to augment our limited physical strength, endurance, memory, and computing capability. However, until recently, there were no machines to augment our decision-making capability. Hence, most of the high-paying jobs in the post-industrial era involve skilled labor that requires substantial decision-making. Only the most mundane and mindless tasks are automated by mechanical machines. Yet, many of these machines still require human operators to make decisions, whether it's as complex as driving a truck, or simply deciding when to switch a machine on and off.

Perhaps, we'll need a new economy that is driven by maximizing happiness rather than profit.

- Dr. Michael Wu

Today, as AI-based technologies become more pervasive, machines can augment our cognitive capacity and automate our complex decision-making processes for the first time. This will dramatically change the way we work, leading to the Fourth Industrial Revolution[1]. Many tasks that were reserved for humans and require some level of human decisions, can be automated as long as we can collect enough data to train an AI to mimic those human decisions and actions.

One of the most significant benefits of AI is the huge efficiency it brings. Since many tasks can now be automated completely and without humans being the bottleneck, they can be executed much faster. Moreover, since AI does not need to eat or sleep, it can work 24/7, leading to

[1] Fourth Industrial Revolution - Wikiwand

further productivity increase. As with any machine automation, AI can eliminate careless human errors, and provide greater consistency in our complex decision-making processes.

Individuals can benefit tremendously from AI because it can eliminate the mundane and repetitive tasks that nobody likes to do. Whether it's something as simple as deciding which movie to watch, adding items to our shopping list or having them delivered automatically, or getting home safely and quickly, AI can automate these tasks, allowing us to spend our valuable time on more important things. Not only do we get convenience and save time, but also, we get better and more personalized experiences.

Now let's expand our scope and look at AI's impact on businesses. Today, many enterprises, especially large ones, have many inefficient or simply broken business processes (e.g. in customer service). These inefficiencies have many negative side-effects on the business, as they often result in higher operating costs (e.g. hiring more staff). Furthermore, when these processes touch their customers, the poor customer experience can erode brand equity and customer loyalty.

Large-scale job displacement in the short term is a problem we must address.

- Dr. Michael Wu

However, as with individuals, businesses can also realize a dramatic efficiency gain from AI. Fixing the inefficiencies in business can indirectly cut costs and improve customer experience. But beyond that, businesses can also improve their customer experience directly using personalization AI (e.g. recommender systems) and create more engaging brand interactions via conversational AI (e.g. chatbots and virtual assistants).

Unlike consumer AI tools that automate simple everyday decisions, business AI can be trained to automate decisions that are often highly technical, domain-specific, and have a much lower tolerance for error. Business AI is much less known to the consumers because they are often used by highly specialized experts. They are used to augment human experts, to not only automate but also optimize their high-stake decisions that often have a direct impact on the company's top line (e.g. real-time

dynamic pricing). Hence AI can also help enterprises improve margins, and revenues, and drive greater profitability.

Now let's further expand our scope and examine AI's impact on our society. As companies and individuals strive to realize greater efficiency from AI, our society as a whole will also function more efficiently. Since the first industrial revolution, we have spent less and less time at work. If this trend continues, maybe in the not too distant future, AI automation could allow our society to function so efficiently that it can support a Universal Basic Income (UBI)[2]. Perhaps, we will no longer need to work for survival, but instead, we work because we want to, for the passion, the experience, and the sense of fulfillment.

Clearly, we are not there yet! Today, our AI systems are only capable of learning from specific data sources and automating point decisions in a narrow domain (i.e. Artificial Narrow Intelligence, ANI[3]). However, since technological progress occurs at an exponential rate, it won't be too long until AI matches human intelligence (i.e. Artificial General Intelligence, AGI[4]) or even surpasses it (i.e. Artificial Superintelligence, ASI[5]). When this happens, ASI could potentially rewrite themselves to make them even more intelligent. This positive feedback of intelligence would grow indefinitely, leading to more and more world-changing innovations at an increasing rate. Humans simply cannot adapt to those rapid and dramatic changes, let alone the existential threat of an ASI. This uncontrollable technological explosion is often referred to as technological singularity[6].

Although the looming singularity is frightening, It's unfruitful to speculate about a knowingly unpredictable future that's far away. Stemming from the mass adoption of AI, there are already many societal challenges that we must deal with long before we reach the singularity. As AI automates more human work in a market society driven by competition and profit maximization, it's inevitable that companies will reduce their human workforce to cut costs. What will humans do then? Perhaps, we'll need a new economy in the future that is driven by maximizing happiness rather than profit.

Since AI advancements progress at an exponential rate, it will be challenging to retrain and upskill the human workforce fast enough for them to keep stable jobs. Although technological innovation always creates more jobs in the long term, large-scale job displacement in the short term is a problem we must address. Moreover, if the pace of change is fast enough, our current education policy, where we front-load education early

[2] Universal basic income - Wikiwand
[3] Weak AI - Wikiwand
[4] Artificial general intelligence - Wikiwand
[5] Superintelligence - Wikiwand
[6] Technological singularity - Wikiwand

in an individual's life, may no longer be practical. So we may also need a new education system.

According to the renowned sociologist, Gerhard Lenski[7], as technology enables more efficient production, it will lead to a greater surplus. This not only supports a larger society but also allows members of a society to specialize more, thus creating greater inequality. Since the efficiency gained from AI is huge, the inequality it creates is also extreme. This is already very apparent in the income disparity between tech and non-tech workers across the globe. Despite the appeal of UBI, it will likely further increase inequality as it would go to everyone equally regardless of their income. Some inequality is good, as it not only motivates people but also enables large-scale projects that require huge investments. However, too much inequality is definitely bad, as it leads to more crime, reduces social mobility, and undermines the fairness and trust of social institutions.

Since the efficiency gained from AI is huge, the inequality it creates is also extreme.

- Dr. Michael Wu

What about the looming singularity and the existential threat? If you must squeeze a comment out of me on this matter, consider this: All AI systems learn from data. But these training data are created by humans, as they are digital records of our past actions and encapsulate our past decisions. So AI is really learning from us, humans, and AI will mimic our decision processes.

Therefore, if we do run into a situation where our interests are in conflict with AI, the best way to ensure that AI doesn't destroy us is for us to be better role models for AI now. That means we, as a human race, must learn to not kill each other whenever we run into conflicts. In short, the best way to ensure our own survival is for us to be better humans. We must learn to be more compassionate, more empathic, more environmentally conscious, etc. So our decisions and action can be used to train an AGI (or ASI) that mimics these 'better-human' qualities.

[7] Gerhard Lenski - Wikiwand

This may sound impossible in today's society because we must compete and struggle for survival, which often brings out the worst of our human nature. However, in an AI-augmented future, we may not need to work for survival, and our economy may no longer be driven by competition. So with the help of AI, maybe we can be better humans before we reach the singularity.

Dr. Michael Wu is the Chief AI Strategist at PROS (NYSE: PRO). He's been appointed as a Senior Research Fellow at the Ecole des Ponts Business School for his work in Data Science, and he serves as an advisor and lecturer for UC Berkeley Extension's AI programs. Prior to PROS, Michael was the Chief Scientist at Lithium for a decade. His R&D won him recognition as an Influential Leader by CRM Magazine. Michael has served as a DOE fellow at the Los Alamos National Lab. Prior to industry, Michael received his triple major undergraduate degree in Applied Math, Physics, and Molecular & Cell Biology; and his Ph.D. from UC Berkeley's Biophysics program.

NELL WATSON

Tech Ethicist, Researcher, Reformer
IEEE STANDARDS ASSOCIATION • NORTHERN IRELAND

It's clear that AI is going to be 10-100x more influential in the 2020s than in the previous decade. Recent developments in 'Transformers' aka 'Foundation Models' or 'Large Language Models' are a tremendous step forward from Deep Learning. These new models can ingest a very broad range of data (spreadsheets, poetry, romance novels, industrial process monitoring, chat logs) and various types of data, such as text, audio, video, etc. They also have the capacity to solve thousands of different problems with one model, in comparison to Deep Learning systems which may be quite effective but only in a narrow range.

Our strange world is only going to get weirder.

- Eleanor 'Nell' Watson

This new technology is also able to deal with abstract concepts in new ways. Simply by asking for something to be 'more polite' or 'less formal', these models can make an appropriate interpretation. This means that one can use everyday, natural language to specify generally what

they want, and then refine it closer to perfection. For example, OpenAI's Codex system is being used to turn natural language into a working video game, in just a few minutes, with all of the associated code immediately ready to be compiled and shared.

Many aspects of programming and development are about to be significantly deskilled, or perhaps bifurcated. People will be creating in simple ways, and a smaller group of experts will be debugging the things that the AI system cannot handle. This wave of creativity will be as powerfully disruptive in the 2020s as the Graphical User Interface and desktop publishing have been in the 1990s.

In recent years we have moved towards a world of services that dematerialize many of our former objects such as media collections. Many new ventures have emerged that leverage the power of mobile internet to make it easier to rent objects for a short time. The covid crisis has obliged many people to cross the digital divide, who otherwise might not have bothered to do so. While this is bringing the world closer together in some ways, we must spare some concern for those who still didn't manage to make the transition to the online world, and who may be increasingly excluded as a result.

Many aspects of software development will be significantly deskilled, or perhaps bifurcated.

- Eleanor 'Nell' Watson

The embrace of digital has consolidated even more power within the hands of Big Tech and the technocratic elite, whilst putting people at the mercy of our digital feudal lords who can exclude us on a whim. This has heightened the need for effective ethics for AI and other technologies that are increasingly entwined with our personal and professional lives.

Moreover, there are risks to consumers from the apparent convenience of our digital world. By no longer owning something, one becomes in essence a renter, and one can be removed at any time, with very little reason given or a chance to challenge such exclusion. If you own things, it's very hard to be taken away from you simply because someone didn't like the things you happened to say. Over time, I think that a desire for ownership will come back into fashion, especially as a status symbol in

and of itself. *"I am a freeborn individual, not a peasant on someone else's fief."*

We also live in a culture of financialization, where stock price becomes the metric to optimize for, instead of actually making things that work and provide value to customers, and by extension things that support civilization as a whole. It's clear that our economic world is built primarily for efficiency, and not resilience. There is very little slack in a just-in-time economy, and so when something inevitably goes wrong, the entire system can get gridlocked.

As governments and corporations, we should do more to prepare for inevitable setbacks that could destroy industries and cause widespread suffering. We should hold back from becoming overleveraged, and ensure that we have reserves and contingencies in place to deal with a world that is increasingly fast, chaotic, and challenging to respond to. Our strange world is only going to get weirder.

Eleanor 'Nell' Watson is an interdisciplinary researcher in emerging technologies such as machine vision and A.I. ethics. Her work primarily focuses on protecting human rights and putting ethics, safety, and the values of the human spirit into technologies such as Artificial Intelligence.

ANTHONY MILLS

Founder & CEO, Executive Director
LEGACY INNOVATION GROUP & GLOBAL INNOVATION INSTITUTE • USA

In the years ahead Artificial Intelligence is poised to have profound impacts on society – in ways we are only just now starting to understand. These impacts will manifest in four areas: holistic interconnection, ubiquitous awareness, substitutionary automation, and knowledge creation.

1. **Holistic interconnection** means that everything in our lives will eventually be digitally enabled and thereafter interconnected in a true Internet of Everything (IoE) manner. This will permit AI systems to intelligently monitor all aspects of our lives (24/7), including us as individuals and all the infrastructure we use on a regular basis - our homes, appliances, entertainment devices, cars, laptops, mobile devices, health aids, and so on. Such holistic interconnection serves as the backbone for realizing truly 'smart' persons, smart homes, smart communities, smart cities, and ultimately smart nations. Eventually, everything will be able to communicate with everything else – and AI will ensure this is done in ways that benefit all.

2. **Ubiquitous awareness** means that AI systems – built atop holistic interconnection – will become fully aware of each person, of personal and societal infrastructure, and of how these are all interacting with each other – and will then make decisions and take actions on our behalf that benefit society in a range of ways. One can imagine the situation where – upon approaching their office building or a shopping mall – that facility becomes fully 'aware' of their presence (including their identity – and that of everyone else there), where they currently are in the environment, where their assets (car, laptop, etc.) are in the environment, and how the environment can best accommodate their needs by learning new insights about them, like their preferred office lighting and temperature, or the promotions being run at stores they frequent, and so on – all in a way that optimizes the whole, like overall energy consumption for example. Eventually, everywhere we go, our environments will be completely aware of us, and will, via AI, optimize the environment for us. In many ways, AI will come to know more about us and our patterns than we ourselves understand (in some places it already does). One important implication of holistic interconnection and ubiquitous awareness is that society's notion of 'privacy' will have to change – to one that is far more comfortable with having individual data shared openly across systems. In due time, this societal norm will shift, and the conversations around privacy in future generations will look very different from those of the present generation.
3. **Substitutionary automation** means that AI will empower numerous automated systems to become fully autonomous in their operation, and consequently be able to deliver value *without* the need for human oversight or intervention. Clear examples of this are fully autonomous vehicles and transportation systems, fully automated business processes, and fully autonomous professional services (like legal and accounting services for example). Substitutionary automation means that many tasks that presently consume (waste) our time – like routine driving, routine data processing, and so on – can all be relinquished to automated systems and consequently free us up to focus our time, energy, and efforts on more creative and novel tasks – tasks for which the human mind is best suited.
4. **Knowledge creation** refers to something that AI has already started to do, namely synthesize new knowledge that did not exist previously (usually via adaptive pattern recognition) – interconnecting points of insight that were previously unconnected. It is this area of AI knowledge creation that is poised to grow exponentially over the coming decades. And not only will it accelerate, it will – via a self-reinforcing cycle – actually start to generate its *own* queries and

learning loops, so that it is not just synthesizing new answers to pre-existing questions, but rather actually synthesizing new questions needing to be answered. This will permit AI to address even better such looming human challenges as climate change, food security, economic stability, poverty eradication, disease eradication, and so on – areas in which next-generation AI holds incredible promise, especially when coupled with powerful new computing methods like Quantum computing.

Most AI scholars agree that, as this acceleration continues, there will come to be a point in time at which AI is generating new knowledge faster than humans can absorb and apply it, at which point AI will surpass human (natural) intelligence, and only AI will then be able to use this new knowledge. This is the *singularity*, which will most likely occur somewhere around the mid-Twenty-First Century. One key ramification of the singularity is that it creates a prediction wall, beyond which we can no longer forecast what the future will look like – because we have no idea what AI will end up doing past that point. The singularity thus presents us with a serious unknown ahead.

There are, of course, key risks with AI. While AI can certainly be used for good to optimize our lives, it can also be used for equally destructive purposes, such as in learning how to wage the most effective wars and cyberattacks against different groups. There is also the ultimate risk, which is that AI itself will become both sentient (fully self-aware) and malevolent (rather than benevolent) toward humanity, thus unleashing some form of 'war' against humanity – to either subdue it or eradicate it.

Society's notion of 'privacy' will have to change.

- Anthony Mills

The first risk – that of humans misusing AI – is a challenging one to address. World bodies like the United Nations for example are – with the assistance of AI Ethicists – already working to develop ethical guidelines for the appropriate uses of AI, and consequences for the systematic misuse of AI. The second risk – that of AI itself overriding human oversight and acting malevolently toward us – is one that can most likely be addressed through discrete control mechanisms in which power to AI systems is cut. Of course, one could imagine the dystopian situation in which such AI systems foresee those human interventions and devise means

(including autonomous war machines under their control) to prevent humans from being able to employ such overrides. Many of these risk-mitigation practices will be worked out as we proceed, and will have to be approached very cautiously.

Anthony Mills is a globally sought-after thought leader on emerging markets, proactive growth strategies, corporate innovation, workplace experience, entrepreneurship, product design, and Design Thinking. His work has had a profound and lasting impact on businesses all over the world.

ALF REHN

Professor of innovation, design, and management
UNIVERSITY OF SOUTHERN DENMARK • DENMARK

As AI (and algorithmic logics in general) becomes omnipresent, the societal implications are getting more profound by the week. Whereas some still see AI as a specialist tool, e.g. as something for pharmaceutical researchers or document management experts, its larger impact will affect any and all human activities. We may not yet be in a world in which AIs decide on everything, from what innovations to invest in and what social programs to fund, but we are far closer to this than most people realize. Whereas AI-driven decisions were just a flight of fancy five years ago, today more decisions than you may be comfortable knowing are, at least in part, driven by algorithmic logic.

It is of critical importance that we retain the human capacity to imagine and dream.

- Alf Rehn

When discussing how individuals and societies will be affected by AI, it is important to balance the benefits with the potential risks. The short-term benefits are ample and easily understood – AIs can take over dreary, repetitive jobs and free people to realize their potential, while getting algorithms involved in decision-making can limit both the errors and the biases that humans are prone to introducing. By leaving decisions to

an algorithm, we can make sure that the innate human limitations – biases, insufficient information, moods – aren't affecting decisions in an overt fashion. By introducing algorithmic logic, we can make sure that human frailty isn't driving the big decisions that society needs to take.

That said, we often forget that AI has both short and long-term impacts on society. Looking at things in the short term, it may well look like AI is nothing but a net positive for society. AI can help us sort out issues such as suboptimal urban planning, or deal with racial bias in sentencing decisions. It can help clarify the impact of credit scores, or ensure that the mood of a doctor doesn't affect a medical diagnosis. What unites these cases is that it is very easy to spot bias or errors in the way the AI functions. An AI that does urban planning in a way that marginalizes certain ethnic groups will be found out and an AI that misdiagnoses cancer will be caught. These are all cases of what I have called 'short bias', errors that algorithmic logic can get caught in through insufficient data or bad training.

But what about those cases where an AI influences decisions that have long trajectories, and where the impact might not be known for years or decades? Imagine that an AI is programmed to figure out which of four new research paths in energy production should be supported and financed. One is known and tested, two are cutting edge but with great potential, and the last one is highly speculative. Unless the AI has been programmed to take great risks, it is likely to suggest that the speculative program is cut. Yet, we know that many speculative ideas – antibiotics, the internet, and female suffrage come to mind – have turned out to be some of the best ideas we've ever had.

What is at play here is something I have given the name 'long bias', i.e. the issue of potential long-term negative consequences from AI decisions that are difficult to discern in the here and now. AI is exceptionally good at handling issues where the parameters are known – whether a cat is a cat, or whether a tumor is a tumor. These are also issues where humans can quickly spot the errors of an AI. When it comes to more complex phenomena, such as 'innovation' or 'progress', the limitations of algorithmic logic can become quite consequential. Making the wrong bet on a speculative technology (and let's be clear, there was a time when the car was just that) can affect society not just in the here and now, but for a very long time afterwards. Cutting off an innovation trajectory before it has had a chance to develop is not merely to say no in the here and now; it is to kill every innovation that might have been, and an AI would not care.

In this sense, AI is a double-edged sword. It can be used to make decisions at a speed that no human can match, with more information than any group of humans could process. This is all well and good. On the other hand, by taking away the capabilities of imagination and bravery

that humans excel at, we may be salting the earth for technologies we've not even considered yet. AIs work with data, and all data is historical – as the investment banks say, "past performance is no guarantee of future results".

With this in mind, it is far too early to be wishing for a technological singularity, a state of affairs where infinitely wise AIs can guide us in our technology exploration. On the contrary, when it comes to innovation it is of critical importance that we retain the human capacity to imagine and dream, and ensure that we are not letting data do all the driving. AI can help us solve massively complex problems, but the keyword here is 'help'. The human capacity "to see a world in a grain of sand/ and a heaven in a wild flower"[8] needs to be protected, to ensure that AI only augments our capacity to innovate, rather than defining the same.

Professor Alf Rehn is a globally recognized thought-leader in innovation and creativity, and is in addition a keynote speaker, author, and strategic advisor. See alfrehn.com

DR. MARILY NIKA

AI Product Leader
TECH COMPANIES IN THE BAY AREA • USA

I am very excited to experience the world embrace and discover Artificial Intelligence in more and more parts of their lives. The benefits of AI in our society are tremendous and cannot be listed in a few paragraphs, but here are three categories that I feel AI impacts the most.

1. **Enhancing our Throughput as professionals**. Have you heard of the term throughput before? Investopedia defines it as "the amount of a product or service that a company can produce and deliver to a client within a specified period of time". According to Accenture, AI might increase productivity by 40% by 2035. This statement makes perfect sense to me, as Artificial Intelligence and Machine Learning empower us both personally and professionally, to avoid tedious day-to-day tasks. Imagine a world where you could only focus on the most strategic, most creative, and most impactful tasks at work, instead of spending your time i.e. troubleshooting a permission issue on your work laptop or crafting the right email to the right person with the right wording. Virtual Assistants are already here and similarly to the movie 'Her', they make life so much easier. We are headed towards a world where we will be able to funnel our brainpower to the tasks that

[8] Auguries of Innocence - Wikipedia

matter to us the most. We may even get told by an AI what tasks should matter the most, according to our personal goals.

2. **Enhancing Our Life**. I worked for many years on Speech technologies for smart devices at home. Being able to use your voice and instruct your home devices to perform certain day-to-day tasks (i.e. playing music at home, setting a timer, retrieving an email or playing a podcast, or even turnings lights on or off) instead of needing to use a keyboard or a phone, creates a sense of convenience and luxury that was previously unimaginable. AI can also help automatically monitor your home for intruders and also reduce energy usage.
3. **Healthcare**. AI can improve, simplify and even save lives. Machines never get tired and thus are less likely to make a mistake compared to humans. There are many studies about how AI can reduce error and diagnose health issues effectively and efficiently, for example being able to diagnose cancer earlier than traditional ways. Moreover, key technologies such as Natural Language Processing (NLP) have numerous applications in healthcare as they can classify, retrieve important documentation, and provide actionable insights in a matter of seconds.

Marily Nika is an AI Product Leader based in San Francisco working for Google, previously for Meta (Facebook). She holds a Ph.D. in Computing Science from Imperial College London and is currently an Executive Fellow at Harvard Business School. Outside of her day role, Marily acts as an advisor to early-stage startups and also empowers the women in the tech community in various ways.

NAZAR ZAKI

Professor and Director
BIG DATA ANALYTICS CENTER, UNITED ARAB EMIRATES UNIVERSITY • UAE

The world of AI is rapidly expanding, with new innovations and breakthroughs happening every day. AI is changing the way we live our lives in many ways and has the potential to be a game-changer for many industries. Some of the benefits of AI are that it can take on repetitive tasks, it can handle tasks that are complex and require human intelligence, and it can help humans make better decisions. AI is already being used in many industries such as healthcare, education, finance, law enforcement, and transportation.

The next generation of AI is expected to do more than just provide insights and suggestions. It will be able to make decisions for us in a way that is more accurate than ever before. This might sound scary but it

could also have some very positive consequences too - like when AI helps doctors diagnose patients who have rare diseases or when an AI detects a person in trouble and alerts them or the authorities in time to prevent a crime from happening.

We need to ensure that AI does not invade our privacy or violate our rights as humans.

- Nazar Zaki

AI has the potential to do a lot of good for society, but there are also many risks associated with it. It's often seen as an object of fear, with people worried about what AI might do to the human race. Some people worry about what will happen if AI becomes too powerful and can make decisions on its own without human interference. Other people worry about how AI can be used as a weapon, such as in warfare.

Other challenges of AI include the lack of ethics. We need to make sure that AI is not being used for evil purposes. AI is also not perfect and it has its flaws. For example, there are many cases where AI has been found biased towards certain races, genders, and other social identities. There is also the issue of privacy: We need to ensure that AI does not invade our privacy or violate our rights as humans.

The idea of AI taking over the world is not new, but the idea that humans will be able to control it is. However, it is not just about being able to control it; it is also about being able to understand how it works and what its limitations are. We need to trust AI if we want it to trust us back. First and foremost, we need to make sure that we are, not only aware of the risks that AI poses but also able to address these risks – and here are some ways in which we can do so:

1. We can work on having a clear understanding of what AI is and what it is capable of so that we know how it will affect us in the future.
2. We can work on creating laws and regulations around AI so that there is some accountability for their actions.
3. We can also limit who has access to certain types of information so as not to put anyone at risk.

Finally, AI will continue to make our lives better. It is becoming more and more integrated into our lives with every passing day. It is hard to predict what the future will hold for AI, but it is safe to say that it will continue to grow and evolve. AI has the potential to help us solve many of the world's most pressing problems, including controlling pandemics, poverty, climate change, and hunger.

We must be careful not to create a future where we are all working for robots or AI beings.

- Nazar Zaki

The future of AI is bright and we should not worry about it too much. AI will help us get better at everything we do and make our lives easier in many ways. The only thing that we should worry about is how to prepare for this change and how to handle it when it happens. There is no doubt that AI will pose some ethical dilemmas in the future and we need to look at how to manage these risks as early as possible. The bottom line is that we need to be aware of these potential risks and work to mitigate them.

Nazar Zaki is a Professor of Computer Science and founder and Director of the Big Data Analytics Center with a mission to ingrain a sustained impact through groundbreaking Data analytics research and services. Nazar's research focuses on data mining, machine learning, graph mining, and bioinformatics.

ANGELIKI DEDOPOULOU

Public Policy Manager, AI & Fintech
META • BELGIUM

Like electricity enables useful things such as light, TV, and fridge for us, Artificial Intelligence is also a ubiquitous technology that can improve our lives. AI could help society to improve healthcare, education, facilitate access to information and significantly improve the efficiency of our workplaces. It can take over dangerous or repetitive tasks and make working environments safer. Furthermore, AI can contribute to the creation of new types of jobs that are demanded in a continuous digital labor

market. Often though AI raises societal concerns around safety and security, privacy, inequality, discrimination, and bias in fundamental rights and democracy. Depending on the data it uses, AI could lead to biased decisions when it comes to ethnicity, gender, or age in the context of hiring processes, banking, or even the justice system. For these reasons, deployers of AI systems should ensure equality, diversity, inclusion, and responsible use of AI to avoid potential pitfalls to society.

Artificial Intelligence can solve societal challenges and reduce climate change.

- Angeliki Dedopoulou

If AI is used responsibly, it can create significant benefits to the global arena. It can solve societal challenges and reduce climate change, and reinforce initiatives like the European Green Deal and Paris Agreement. AI can also contribute to the realization of the sustainable and development goals of the United Nations and it can play a crucial role in curbing global issues such as:

- **Control of epidemics**. During a global pandemic, governments' initial objective is to minimize the spread of the disease. If AI is fed by historical data, it can recognize patterns and trends, and then via predictive analysis, it can lead to necessary measures to eliminate the spread of the virus. AI can also be used to accelerate the development of vaccines related to viruses.
- **Management and control of pollution**. Pollution is a global challenge that concerns all countries around the world. AI can be used for the protection of the environment and pollution control. More specifically, AI systems contribute to the detection of energy emission reductions, the removal of CO2, the monitoring and prediction of extreme weather conditions, and support the development of greener transportation networks.
- **Global food crisis prevention**. According to the United Nations, 840 million people might be affected by hunger by 2030[9] . Research has shown that with the combination of smart agriculture and machine

[9] Food | United Nations

learning, this number could be significantly reduced. AI-based solutions can create systems that warn governments on food shortages and prepare them for better food supply and management. AI solutions can also help farmers to produce more food with less land. For example, AI-enabled operations are estimated to use roughly 90% less water and produce over 20 times more food per acre than traditional fields[10].

- **Water pollution management**. UN's Sustainable Development Goal 6 seeks to ensure that people have access to clean water and adequate sanitation services worldwide. AI can be used to reduce pollutants in the water and detect the amount and composition of toxic contaminants. AI can also increase the efficiency of waste management systems.

For businesses, AI can enable important sectors such as tourism, construction, agriculture, green and circular economy. It can also improve the quality of products, increase production levels, and contribute to energy savings. For example:

- When customer service representatives in a hotel are not available, AI bots may respond to questions and provide useful information.
- In the construction sector, AI can improve project efficiency and the safety of workers in construction sites.
- Smart farming is helping the agriculture sector to be more profitable. AI-powered mechanisms can monitor aspects such as grain mass flow, the quality of harvested grains, and moisture content.
- AI in combination with intelligent grid systems and deep predictive models can manage the demand and supply of renewable energy.

To conclude, AI has a huge potential to benefit societies and boost many sectors of the economy. However, governments need to ensure that AI prioritizes humans, reinforce human's trust, protect human rights, and promote creativity and empathy. Policymakers should create a global dialogue, explain, educate, and boost transparency of AI, adapt training and education curricula to the new Artificial Intelligence society, and promote, develop and encourage the public and private sectors to adopt human-centered and trustworthy AI.

Angeliki Dedopoulou is Public Policy Manager for AI & Fintech at Meta (formerly known as Facebook). Before joining Meta's EU Public Affairs team, she was a Senior Manager of EU Public Affairs at Huawei, responsible for the policy area of AI, Blockchain, Digital Skills, and Green-related policy topics. She was also an adviser for the European Commission for over 5 years on DG Employment, Social Affairs, and Inclusion.

[10] This is how AI could feed the world's hungry while sustaining the planet (weforum.org)

JON SKIRNIR AGUSTSSON

VP Artificial Intelligence and Data Research
NOX MEDICAL • ICELAND

I feel that the focus of AI state of the art is shifting away from newer models or new ways of training models, towards how we better address issues of bias, social impacts of our AI models, and how we phrase our optimization questions. We seem to have reached a point in time where most progress in AI will come from focusing on data and thinking about how we evaluate the AI models' performance rather than improving AI architectures or hardware. I am not saying that there is no room for improvement there, but I feel that the real breakthroughs will happen around how we improve our data and the questions we want AI to answer.

I am a big fan of the work Andrew Ng[11] does with the website deeplearning.ai[12] both in providing great educational material in AI and more importantly to highlight the importance of data-driven development of AI models. My background is in instrumentation and measurement technology and, therefore, I am a big believer that if you do not collect the right data there is no way of getting a sensible output no matter how much AI you throw at the problem. There are also other great references out there on how to get more valuable input from human labelers such as the book by Robert Monarch titled Human-in-the-Loop Machine Learning[13], describing data-driven methods of selecting better training and validation data. Finally, Stuart J. Russell proposes that we rethink how we optimize AI algorithms in his book Human Compatible[14].

AI and automation are already everywhere and we see great societal benefits from AI, such as increased productivity and fewer human errors. But, as AI becomes more ubiquitous, we also see how systematic errors and biases in AI models can start to cause real social and economical problems. Furthermore, we have started asking very deep ethical questions when we delegate decision-making to autonomous machines. We have also started running into barriers where our legal and societal frameworks do not manage the state of the art in AI.

We currently see in our society that if AI algorithms are not optimized for the correct objectives they can start spreading misinformation and polarizing people on social media. This may be caused by optimizing an AI model for the wrong thing such as for some engagement metric and not asking what such an optimization might lead to. In an era where we are exposed to enormous amounts of information, it becomes difficult

[11] Andrew Ng - Wikipedia
[12] Home - DeepLearning.AI
[13] Human-in-the-Loop Machine Learning: Active Learning and Annotation for Human-Centered AI
[14] Human Compatible: AI and the Problem of Control

for individuals to keep track of what is correct and there is space for AI to spread misinformation to large groups of people.

We also have many examples where bias in training data results in racist or biased outcomes from AI models. This becomes especially troublesome when we use AI models to assist in or even make life-changing decisions for individuals. In these cases, the AI models may perform well on average or for most people. However, there may be few individuals who are severely negatively affected by the model outputs. There are famous examples of this when AI models have been used in the financial and legal sectors. In my sector, AI medical devices, we are also faced with this problem where rare clinical conditions may be missed when AI is used inappropriately.

There are also interesting societal and legal questions that have started to arise. We have started to think about who is responsible when an AI system causes an accident or harm, and how we best react when AI systems take over important systems in our society and take them out of our control, for example crashing the stock market.

With all of this in mind, one quickly realizes that we need to think about AI from many different angles. There are the technical and application aspects, but there are also ethical, societal, and legal ones that must be considered. Thus, we need a diverse group of people working on AI and diversity is our best chance of being successful. For people who are interested in being involved in AI development and adoption, I think there are many great opportunities to contribute to, not only model development, but also to seek answers to the challenges mentioned above.

Jon S. Agustsson is an experienced AI and Research leader working in the medical device and medical research industry-leading an interdisciplinary team in Data Science, Physics, Electrical Engineering, and Research. Passionate engineering, inventing, and building new things, with multiple patents and scientific publications.

BOYKA SIMEONOVA

Assistant Professor in Information Management
LOUGHBOROUGH UNIVERSITY • UK

"The rise of powerful AI will be either the best or the worst thing ever to happen to humanity. We do not yet know which."[15]. As per Stephen Hawking's quote, the effects of AI remain unclear and unknown. However, increasing digitalization and AI utilization could radically transform work and society, knowledge, learning, and power (re)distribution. Given

[15] Stephen Hawking warns of dangerous AI - BBC News

the dangers of lost individual knowledge through the increased use of AI and algorithmic decision-making, the overreliance on AI and algorithms might hamper learning, decision-making, and innovation. For example, AI has assumptions about knowledge, particularly tacit knowledge, which are currently highly problematic and require considerable improvement prior to the reliable use of AI and its predetermined codified (and encoded) knowledge. At this point, it is helpful to differentiate between knowledge, information, and data.

AI could lead to emancipation through empowerment, autonomy, inclusion, participation, and collaboration.

- Boyka Simeonova

Information, and data, are an ingredient to knowledge but do not represent knowledge. Data are facts, information is the processed data and knowledge is the interpreted and actionable information. Knowledge can be explicit and tacit. Explicit knowledge can be easily captured, codified, processed, stored, and distributed. Tacit knowledge cannot easily get captured, codified, processed, stored, and distributed - tacit knowledge is accrued through experience and is explained as an ongoing accomplishment through practice and participation.

The problem with knowledge encoded in the AI is that it is narrow and brittle, and AI systems are only reliable in a narrow topic and domain, which are predetermined, and when the topic and domain are challenged or changed, AI systems "fall off the knowledge cliff"[16]. Therefore, while AI might help for narrow, routine, predetermined tasks, it is (as yet) unreliable and inaccurate to help with complex problems and decision-making, where automation and the use of AI are yet currently impossible because tacit knowledge cannot be easily codified[17]. AI can analyse volumes of data, however, the knowledge aspect needs further development.

[16] Forsythe, D.E. (1993). The construction of work in AI. Science, Technology, & Values, 18(4), 460-480

[17] Simeonova, B., & Galliers, R.D. (2022). Power, knowledge and digitalisation: A qualitative research agenda. In Simeonova B. & Galliers R.D. (Eds.), Cambridge Handbook of Qualitative Digital Research. Cambridge University Press.

Despite the danger of AI systems falling off the knowledge cliff, the use of AI, automated and algorithmic decision-making has increased in organizations and societies. For example, the use of Decision Support Systems and 'big data' has limited the power of individuals in strategic decision-making and has (in some instances) replaced their tacit knowledge, experience, and expertise on the assumption that their calculated rationality leads to superior outcomes. For example, Fernando Alonso lost the 2010 Formula 1 Grand Prix Championship because the race simulation algorithm provided a poor decision, and the Chief Race Strategist did not have the power to participate in the decision-making or to change (or overrule) the decision of the algorithm, which led to Alonso and Ferrari losing the championship and to the Chief Race Strategist losing their role[18].

The issue of tacit knowledge encoded in AI is also demonstrated in the challenges around the development of the decision-making algorithms for autonomous cars where decisions need to be predetermined and context and interpretation of a situation are currently limited. Therefore, how AI is used needs considerable thought and consideration as in its current state, it does not have enough knowledge capabilities, and its current use hampers knowledge, learning, decision-making, innovation, and society.

In the context of a digital economy, AI and automation could advance the power of the influential through control, surveillance, monitoring, discrimination, information asymmetries, manipulation, 'algorithmification', and 'datafication'. Such uses of AI, examples of which currently dominate, lead to exploitation, exclusion, marginalization, discrimination, and manipulation.

For example, Cambridge Analytica, which ran the American presidential digital campaign, arguably manipulated the opinions of people and their votes or voting intentions or behaviors, through the provision of filtered information to influence their votes. AI has been used to exploit the practices of people and their opinions, resulting in manipulating the vote. Other examples of exploitation are the automation and 'algorithmification' of influential technology organizations that collect and exploit data, eliminate competition, and coerce organizations to follow their algorithms[19]. Therefore, influential technology organizations may further consolidate and increase their power because of their technology leadership and the opportunities for exploitation practices.

AI could lead to emancipation through empowerment, autonomy, inclusion, participation, and collaboration. However, such examples are scarce, and the emancipatory use of AI, or the emancipatory outcomes

[18] Aversa, P., Cabantous, L., & Haefliger, S. (2018). When decision support systems fail: Insights for strategic information systems from Formula 1. The Journal of Strategic Information Systems, 27(3), 221-236

[19] Naidoo (2019). Surveillance giants. Amnesty International.

of the use of AI are limited. Organizations, developers, governments, workers, and societies need to collaborate on determining how these AI systems are developed and used to enable emancipation and empowerment.

Boyka Simeonova is an Assistant Professor in Information Management at Loughborough University. Boyka is the Director of the Knowledge and the Digital Economy Research Network at the university and the Deputy Director of the Centre for Information Management.

ASK A QUESTION SPREAD THE WORD

9 THE IMPACT OF AI ON EMPLOYMENT

In the years to come, we will witness significant changes in the workforce and the markets. Roles and jobs will become obsolete, industries will be radically transformed; employment models and relationships will be redefined.

But will AI drive more unemployment by replacing humans or it will create new roles and help people by releasing them from repetitive tasks? Which sectors are expected to be most impacted by AI in the near future?

Alexandre Guilbault, Harry Mamangakis, and Mike Tamir share their insights.

ALEXANDRE GUILBAULT

Director of Advanced Analytics & AI
TELUS • CANADA

I don't think there will be more unemployment as AI applications become widespread. Previous technological revolutions transformed roles and created new ones, but did not lead to mass unemployment. Today, we have quick access to all the information in the world and many tasks can be fully automated by computers however, jobs didn't disappear.

You may have heard of the dream of a Civilization of Leisure that many people expected to see come true around the new millennium. We thought that with the new technologies we could reduce our working hours to a minimum and everyone could retire before 55 years old. It still hasn't happened and I am sure it won't even when AI is more developed and democratized.

I don't think we should fear the impact of AI on any of our roles.

- Alex Guilbault

As a simple example, computers can now fully assist you to do your taxes, sometimes better than a human could, but the accountants are still around, working roughly 40h a week and they are in high demand. Computer programs can improve their job but can't replace the human approach, adaption to specific personal situations, and they can't innovate beyond what was initially hard-coded.

Yes, AI can, or at least will soon be able to drive a car autonomously, but we will still need truckers. Why? Think about it, trains or subways are on tracks, it would have been pretty easy to hard-code their shift, but we still have train and subway operators. Even pilots can fly airplanes today without any manual intervention, and this happens even without AI technologies. The role of pilots has been completely transformed, but pilots are still needed. We are still a long way from trusting these lines of code enough to fly an unmanned aircraft.

AI technologies allow computers to perform new tasks that hard-coded computer programs couldn't perform before. With Computer

Vision, Natural Language Processing, Reinforcement Learning, etc., tons of new applications can now make many of our tasks easier. These technologies are truly impressive, but to date, their applications are still very narrowed. Although every new publication tends to say that we are getting a step closer to a General Artificial Intelligence, to drive hype and increase click-through rates, we are still far from a real artificially intelligent being, and even once that happens, we will still need people to make the right decisions, to maintain and manage their software, to research and improve them and to have this highly needed human approach.

We should all educate ourselves so that we can benefit from AI to improve and facilitate our work.

- Alex Guilbault

I think almost every industry will be impacted by AI in the near future. These technologies help us build new tools that we can use to free up people from repetitive tasks, optimize processes, facilitate monitoring, support creative and innovative work, improve efficiency, and much more. Every area will benefit from those new technologies in one way or another, but some will certainly take more time to be affected than others.

The sectors affected today are certainly those where data to train these algorithms is readily available. That's why Alphabet (Google), Meta (Facebook), and Amazon are behind some of the most transformational applications. Then come those where data can be easily collected and the expected behaviors are more or less circumscribed (ex. Obey the rules of the road, don't crash, get me from point A to point B for self-driving cars). Later, these technologies will aid decision-making in medical roles, but they won't replace the need for the physician's or psychologist's human approach. Surprisingly, some of the toughest jobs for AI to improve might be some of the manual labor involved in construction - plumbing, for example. Some applications can help improve their work, but even with the advance of AI in terms of software, we still have a long way to go to develop mechanics that would allow a robot to go to the

right place, squeeze into tight areas, find the right tools and show enough creativity to adapt to a specific problem.

In a nutshell, I don't think we should fear the impact of AI on any of our roles, but we should all educate ourselves so that we can benefit from these technologies to improve and facilitate our work.

Alex Guilbault is leading data-driven transformations to reinforce TELUS pole position entering the Age of Insights by leveraging Advanced Analytics and cutting-edge AI technologies.

HARRY MAMANGAKIS

Chief Technology & Operations Officer
VOICEWEB • GREECE

Artificial Intelligence is considered a revolutionary technology and is one of the fundamental pillars of the technology dimension of Digital Transformation, which in turn drives this new industrial revolution era…

Wow… a lot of buzzwords! To answer the question of what impact AI will have on employment, we would need to go back, and study other revolutionary technologies introduced and what their impact was. This will probably help us make an educated guess on what the impact of AI would be.

A technology is 'revolutionary' when it covers a gap, a need for everyone.

- Harry Mamangakis

Let's start by defining when a technology is 'revolutionary'. It is revolutionary *when it covers a gap, a need for everyone*. It is as simple as that. For example:

- **The Printing Press**: It helped cover the 'knowledge gap' making knowledge easily available to everyone. A printed book could easily pass on from person to person and transfer knowledge.

- **The Engine**: It helped cover the 'power gap', that is needed to achieve scalability that labor (human or animal) could not cover. As a result, it helped scale industrial production and also had a significant impact on our daily lives (cars, airplanes, and so on).
- **The Internet**: It helped cover the 'time and distance gaps' allowing everyone on the planet to interact and transact from anywhere and at any time.

So, what will AI bring to the table to be considered a revolutionary technology? AI enthusiasts will have already started listing all the benefits… But we are searching for (at least) the one that will have an impact *on all of us*, the same way that the other revolutionary technologies have had.

Observing it from 30 thousand feet high, AI is good at processing large amounts of data, learning from it, and then helping us in making educated decisions. So, the gap (or one of the gaps) that AI can help cover, is that of *decision-making*. Be it in the healthcare industry (which affects us all for example in helping a diagnosis) to improving your crop production, AI is a powerful decision-making tool. No one can argue that, for decision-making, AI is something that will affect *all of us* in some way.

So now that we have established that AI is a revolutionary technology, let us try to see what its impact will be, and more specifically as the question states, how AI will impact employment.

AI helps professionals work more efficiently and make better decisions.

- Harry Mamangakis

The godfather of Deep Learning, Professor Geoff Hinton, had said in late 2016[1]: *"If you work as a radiologist, you're like the coyote that's already over the edge of the cliff, but hasn't yet looked down so doesn't realize there's no ground underneath him. People should stop training radiologists now. It's just completely obvious that within five years, Deep Learning is going to do better than radiologists."*

A very powerful statement, but is there any basis to it? Well, AI has been used extensively in several sectors of the healthcare industry, for

[1] Geoff Hinton: On Radiology - YouTube

several years now. As time goes by, results and research is building up to a very convincing argument, that AI is revolutionizing healthcare. For example, from a study from the journal 'Annals of Oncology' from 2018[2], we learn that in a comparison between human dermatologists and deep learning, the former's success rate was 86.6% and the latter's was 95%, in their ability to accurately detect skin cancer. So, does this mean that AI will replace these medical professions? Is Professor Hinton right?

From movies like the 'Terminator' we are inundated with the idea that machines are slowly taking over, but are we there? Have we reached that 'apocalyptic' point of no return yet? The answer is no. AI will help reshape the way medicines work and will be a very useful tool for various medical specialties like radiologists. From helping them eliminate their backlogs of non-urgent cases due to COVID-19 incidents, to allowing them to be more proactive in diagnosing patients' conditions. This is also proven by numbers: since Professor Hinton's prediction, the number of radiologists (in the US) has increased by 7% between 2016 and 2019.

So, there will be no impact on employment? Is that the answer to the question? Well, let's look at what happened in the previous revolutions we discussed.

In the early days of 'the Engine', there was so much widespread fear on what the impact would be, that even mobs have been attacking venues that used engines, for the fear of human labor being replaced. Secret oath-based organizations like the Luddites were formed to fight against manufacturers in the textile industry.

So, what did happen? A lot of (unskilled) labor was replaced, indeed. But as production increased new labor positions were created. From 'machine operators', to 'machine maintenance personnel' to 'quality control' workers at the end of the production line. The Machine (Engine) did replace unskilled work but did introduce new opportunities for skilled workers.

What about the internet? Same here. Those impacted either adopted it or were forced to adapt to it. As a result, and as seen for the healthcare industry, it will need to adopt AI and adapt to it. This explains why the number of radiologists increased, as mentioned above. The process of decision-making is now easier; hence it is attractive to more people.

The same applies to all other sectors/business domains. From the supermarket commercial officer in charge of pricing who currently uses spreadsheets or other conventional tools to calculate what the sale price of product A should be in order to increase sales but not steal from competing product B sales, to a farmer who wants more crops with fewer resources, AI will be a tool to help them work more efficiently and make better decisions.

[2] Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists - Annals of Oncology

AI will require a new set of skills. A McKinsey report[3] stated that less than 5% of jobs will be lost due to AI, however, around 60% of jobs will require adaptation to AI, developing new skills to factor in AI to our work, as the tasks we do will change.

But similarly to the Engine, some labor, particularly that of unskilled workers who do repetitive tasks will be replaced, with new opportunities/jobs that will eventually improve the quality of our lives.

Harry Mamangakis is a Technology Executive for over two decades, balancing between fluency in technological breakthroughs and having a business mindset. He has led and participated in several transformation engagements for leading brands in industries such as Telcos and Retail.

MIKE TAMIR

Chief ML Scientist, Head of ML/AI
SUSQUEHANNA INTERNATIONAL GROUP • USA

2022 marks the 70th anniversary of one of Kurt Vonnegut's earliest works, Player Piano. In it, he describes the replacement of working humans with machines as coming in three phases: The first replaced heavy 'physical' work, while the second replaced complex but 'repetitive' work. AI, it seems inevitable, will exacerbate the removal of repetitive work (something that has also been going on for 'some time').

Humans who master math and coding will have job security for the foreseeable future.

- Mike Tamir, PhD

This ranges from making (and serving coffee - my first after-school job was as a barista) to long-haul trucking, and potentially manual driving entirely. No doubt this will impact employment, possibly for the better (e.g. truckers can stay local and focus on last-mile delivery while automated vehicles focus on moving freight on the easier to solve problem

[3] Jobs lost, jobs gained: Workforce Transitions in a time of automation (mckinsey.com)

of highway self-driving), but possibly not for the better as Vonnegut mused.

In the Player Piano dystopia *'the third one's [also] been going on for some time, if you mean thinking machines… machines that devaluate [sic.] human thinking.'* Here it is not so clear that advances in AI, which perhaps have helped to supplement complex 'thinking' tasks actually devalue human 'thought' labor. Advances in open source projects, collaborative research, and other improved, complex language-related tasks have the potential for making 'thought' labor all the more effective. From AI-assisted research to AI-assessed art, we can maintain the hope that while such technology might change the nature of such labor, it may also open new doors to improving its power and impact.

One thing is clear, humans who master mathematics and coding will have job security for the foreseeable future. It is not far-fetched to imagine that this could mean a need for societal changes to ensure a thriving economy. Altering our education paradigms so that coding is more central to modern education standards, and enhancing the importance of mathematics and statistics (along with scientific reasoning) could be critical to ensuring thriving societies in the coming decades.

Mike Tamir, PhD is a data science leader, specializing in deep learning, NLP, and distributed scalable machine learning. Mike is experienced in delivering data products for use cases including text comprehension, image recognition, recommender systems, targeted advertising, forecasting, user understanding, and customer analytics. He is a pioneer in developing training programs in industry-focused machine learning and data science techniques.

10 HOW SHOULD SOCIETIES GET PREPARED FOR AI?

This technological revolution brings great opportunities for prosperity and growth. However, we need to get prepared and ensure that the technology will be applied and used in the right direction.

How could the general population become more aware of both the potential and the risks of AI? How should the educational system adapt to the new AI-powered world?

Golestan (Sally) Radwan, Mathew (Mat) Hughes, and Himanshi Allahabadi share their insights.

GOLESTAN (SALLY) RADWAN

AI Researcher and Strategist
EGYPT

AI education is quite possibly the most important aspect of AI adoption facing societies at all levels. There are several reasons for this: (a) AI, more than other technologies, penetrates all levels of society. Hence, knowledge needs to be built for different audiences playing different roles, (b) AI builds on the existence of other foundational knowledge, and (c) Societies across the world differ in their needs, priorities, and readiness levels to absorb AI knowledge. It is therefore important for all stakeholders involved (governments, NGOs, schools, universities, and anyone involved in national AI efforts) to bear this in mind and to design a comprehensive framework for AI capacity-building tailored to the needs of each specific society.

General awareness programs should dispel the myths around Artificial Intelligence.

- Golestan (Sally) Radwan

But before we get into the nuances of different capacity-building strategies, let's first discuss some commonalities. A comprehensive capacity-building framework should take into account the different roles that members of society will play in an AI ecosystem. These can be broadly categorized into ‘technical’ and ‘non-technical’ roles. Technical roles include anyone who will be involved in the design, development, operation, or maintenance of an AI system. These can include roles such as Data Scientists, Machine Learning Engineers, Data Architects, MLOps Engineers, and many others. Each of these roles has its own strategy and set of curricula that can be designed at different levels to prepare them for the job market.

While it may seem like a challenging enough task to design programs for technical roles, it is actually the non-technical ones that pose the

greatest challenge. Non-technical roles start with generic notions like increasing public awareness of AI. This by itself is a significant challenge as it requires the presence of a basic level of technology literacy among the public - and this is where we start getting into nuances across societies. While a developed country like Finland can easily develop a university-level AI course and use it to educate 1% of the population (Finland being a country with a 100% literacy rate and where roughly 75% of the population hold a college degree), a developing country might first need to ensure that the majority of its population can read and write before introducing them to something as sophisticated as AI.

A vital component of a society's readiness for the age of AI is to train its 'domain experts'.

- Golestan (Sally) Radwan

The picture is not necessarily that grim, however, as evidence has shown that people, even those without basic access to education, are still able to absorb knowledge such as the use of smartphones, if a) they can see the value in it, and b) if that knowledge is presented to them in a manner which suits their needs.

So, going back to the Finland example, a 'general awareness' course would simply need to be a text-based, online course in English. While in a developing country, the same course might need to be delivered as a series of short videos on television by the country's main broadcaster, in a colloquial language people will understand. Moreover, it needs to be full of demonstrations of how to perform certain tasks and possibly supported by a phone hotline to answer people's questions.

It is up to each country to decide which level of general awareness suits its needs. However, it is equally important for global leaders to be aware of those differences in needs and starting points, and to not assume that it is equally easy for countries to implement AI awareness programs. Developing nations need all the support they can get, not in the least because their populations are often the most vulnerable to abuse.

General awareness programs should cover the basics of AI, its benefits and limitations, and crucially, its risks and ethical aspects. If the country has any legislation regarding AI, data protection, or similar, this should also be included. Most importantly, such programs should dispel the

myths around AI, which are unfortunately still prevalent, especially in the developing world, such as its impact on the economy, or even worse, potential superintelligence that will make humanity extinct.

Apart from general awareness programs, a vital component of a society's readiness for the age of AI is to train its 'domain experts'. The domain expert is the non-technical professional who 'owns' the problem to be solved by AI. Examples include healthcare professionals, agriculture specialists, marketing professionals, legal experts, and many many others depending on the domain of application of AI. In my view, no AI project can be successful without the involvement of well-trained domain experts. These are the individuals who will advise on the exact problem to be solved, data availability, challenges, validity and relevance of results, and many other aspects so crucial to a project's success. Training domain experts is best done by introducing a compulsory 'Introduction to AI and Data Science' course into universities, or even secondary schools if possible. This should be tailored to each domain, for example, 'AI in Agriculture', 'AI in Healthcare', etc, and a version of it should be developed for professionals currently working in that domain. This will ensure a roster of experts able to collaborate with technical teams, and even able to identify and suggest AI projects by themselves.

The last, important group of non-technical roles are leaders, including those in government, business, civil society, or any other domain. While they may not need to develop the deep expertise of a domain expert, they definitely need more preparation than a general awareness course. They need to be taught how to think about AI as an investment: how to identify an AI project, how to allocate the right investment, find or train the right people, assess its success, and, if relevant, develop their organization's own AI and Data strategies. This group also includes current and aspiring entrepreneurs, who need to assess their readiness to become producers or users/consumers of AI.

The above is by no means a comprehensive capacity-building strategy for AI, but rather some thoughts on how societies can start thinking about developing such strategies to suit their needs and priorities. International cooperation and knowledge sharing are vital enablers to the success of AI adoption worldwide.

Golestan (Sally) Radwan is the former AI Advisor to the Minister of ICT in Egypt, where she led the development and implementation of the country's national AI strategy from 2019-2022. She also participated in drafting global recommendations on AI ethics at UNESCO and the OECD.

MATHEW (MAT) HUGHES

Schulze Distinguished Professor, Entrepreneurship and Innovation
LOUGHBOROUGH UNIVERSITY • UK

To imagine a world of AI, it would be easy to direct the interested reader to the realm of science fiction. From films such as The Terminator to iRobot and 2001: A Space Odyssey, or video games such as Deus Ex, creative minds have marvelled over the potential benefits and dangers of an Artificial Intelligence. The promise of AI in many ways lies in its ability to process vast amounts of information to make or arrive at unbiased, accurate, or 'correct' decisions, fair to millions of people. Imagine the benefits to the electorate, to government, to society… For now, innovative commercial uses abound, with examples including analytics (e.g., Anodot), assisted decision-making and marketing campaigns (e.g., Peak, AI.Reverie, Frame.ai), insurance (Arturo, Inc.), code (e.g., Comet.ml, Metabob.com), cybersecurity (e.g., MixMode, Socure), autonomous vehicles (e.g., Pony.ai), and open banking (e.g., Cleo). But what might a society relying on AI look like?[1]

Society must not become a passive observer of Artificial Intelligence.

- Mathew (Mat) Hughes

As part of an answer, AI can mimic or simulate human interactions. For example, Jill Watson is a teaching assistant at Georgia Tech that answers student questions realistically to the point that students do not even realize that they are interacting with an AI and not with a human professor; Google Duplex acts as a phone assistant using Natural Language Processing to create conversations with a human quality; and GPT-3 uses AI to generate human-like texts based on questions asked of it. Potentially then, AI can support humans' roles in society or even replace them.

[1] Robledo, S., Grisales Aguirre, A.M., Hughes, M., and Eggers, F. (2021), "Hasta la Vista, Baby – Will Machine Learning Terminate Human Literature Reviews? An Entrepreneurship Application", Journal of Small Business Management, in press.

However, its ability to mimic and simulate human interactions by absorbing and processing a plethora of existing (and historical) information carries many serious, and not necessarily foreseen, consequences. For example, machine learning based AI can learn to be racist. In the US, its justice system, frequently faced with charges of racial bias, turned to technology to create unbiased profiles of criminals only to find that the algorithms developed a racial bias. In other cases, a LinkedIn advertising program showed a preference for male names in searches, and a Microsoft chatbot called Tay distributed antisemitic messages having spent a day learning from Twitter.

AI and machines are not infallible and the assumption that they will operate with objectivity is flawed because of both the programming requirements and the information they feed on containing inaccuracies, biases, or flaws, whether current or historical. For instance, human intervention in the training of machine learning cannot be overlooked. The quality of data or information fed into the machine matters substantially to the quality of the outcome—training is key, and volume helps, but these do not overrule the need for human and societal oversight of machine learning and AI and interpretations of their outcomes to ensure the machine or AI is not learning and absorbing incorrect information, trends, or perpetuates flaws originally present in data.

Business and government both have a role to play in educating society of the usefulness and limits of AI, but the question is who should provide oversight? Society is a major stakeholder of AI, but society is too distant and disparate to impose effective oversight or control. Therefore, it falls on government, business, and non-government institutions to ensure (potentially through regulation) that AI properly services the public good. In many ways, the potential of AI lies in augmenting human agency (and decision-making), not replacing it. Many opportunities certainly lie at this interface and to manage the potential Janus face of machine learning and AI.

Mathew (Mat) Hughes is Schulze Distinguished Professor and Professor of Entrepreneurship and Innovation at Loughborough University. His expertise lies in entrepreneurial management and strategy. Mat has published in world-leading journals and served as an advisor to a variety of businesses and organizations. Mat is also Senior Editor of FamilyBusiness.org

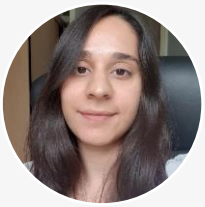

HIMANSHI ALLAHABADI

AI Ethicist
INDIA

Before delving into my thoughts, I think it's important to point out the distinction between General and Narrow Artificial Intelligence. The former speaks more to the sci-fi portrayal of AI while the latter is what powers most of the leading tech today. The following is about narrow AI.

First and foremost, it is important to realize that while AI is used for forecasting and predictions and is generally future-looking, it uses historical data to try to find patterns and essentially 'fit' that data. Some of the areas of application that have thrived due to this property of AI are facial recognition, financial forecasting, medical image classification, and recommendation systems. The learning system by itself, while classifying previously unseen examples, works under the following assumptions:

- The input provided to train the AI model is all of the information needed to predict the outcome
- The data itself (i.e. the 'ground truth') is correct and unbiased
- The future will not be too different from the past

These assumptions altogether have important implications on the outcomes. They can lead to potential risks, particularly where AI is deployed for decision-making in the real world. In order to educate oneself about AI, one should keep these assumptions in mind.

Education systems can be adapted to prepare individuals for an AI-powered world.

- Himanshi Allahabadi

Some key considerations come to mind as a result of these assumptions. For instance, it is important to think about the gap between the context of AI deployment and what AI can do, given the data at hand. We should have clarity regarding what the AI predicts, the key decision(s) being made based on the predictions, and the factors the system does not account for. We should be aware of any real-world factors, measurable or not, that are not included in the input of the learning model, and that

could influence the outcomes. Whether or not those factors have been appropriately considered for interpreting results and making decisions, speaks to the goodness of the solution.

Secondly, it is important to consider what real-world policies and practices generate the data distribution. The AI system will more or less try to emulate these policies and practices implicitly, in the sense that its predictions will be consistent with what it learnt from the training data set unless explicitly trained to fulfill other objectives. For instance, if a bank's policy for accepting loan applications is discriminatory, an AI system trained on their data will reflect and possibly amplify the same.

Lastly, while AI can be a powerful tool to analyse and mine patterns from the data, it is important to understand that sometimes future decisions cannot solely depend upon patterns learnt from historical data. To that end, it is essential to recognize the need for human expertise and knowledge.

Education systems can be adapted to prepare individuals for an AI-powered world by teaching them how it fundamentally works, making them aware of its potential, current limitations, and risks, and imparting them with knowledge of different career paths that leverage or research AI and its impact. Ultimately, AI is a problem-solving tool – albeit a rather powerful one. Education systems can encourage students to think about AI-based solutions for problems that they come across in their daily lives and empower them to use AI for environmental and societal good.

Himanshi Allahabadi is a data scientist focusing on AI engineering and ethics. She is experienced in the development of large-scale cloud platforms for data-driven applications. She researches areas of trustworthy AI and machine learning best practices for collaborative open-source initiatives.

JOIN THE DISCUSSION SHARE ON SOCIAL MEDIA

11 WHAT ARE THE ETHICAL CONCERNS ASSOCIATED WITH THE GENERAL ADOPTION OF AI?

There are serious concerns and unanswered questions regarding the social, political, and ethical implications of the massive adoption of AI. But what are the specific risks for humanity? How could an AI system be biased and what would be the impact on societies? How could we remove bias from AI systems? Could humanity lose control of the 'global AI' system?

Arathi Sethumadhavan, João Azevedo Abreu, Aruna Kolluru, Veena Calambur, Enrico Panai, Mike Tamir, Aruna Pattam, Therése Svensson, Samiran Ghosh, Netanel Eliav, and Mayte Hidalgo share their insights.

What are the ethical concerns associated with the general adoption of AI?

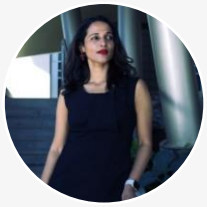

ARATHI SETHUMADHAVAN

Principal Research Manager, Ethics & Society
MICROSOFT • USA

Artificial intelligence is increasingly informing crucial aspects of people's everyday lives. Whether playing a part in the security apparatus for buildings, the screening of candidates for jobs, or the diagnosis and prediction of patient health, AI technologies are being used to facilitate countless essential services. As such, there are real human costs when systems do not perform reliably or equitably across the variations found in real-world settings.

It is now more important than ever to have principles, governance, and processes for responsible development of AI.

- Arathi Sethumadhavan

More recently, we have seen several AI applications that perpetuate unfair biases. In 2020, unable to administer traditional exams due to the coronavirus pandemic, the United Kingdom's Office of Qualifications and Examinations Regulation decided to use an AI to predict the grades of students in A-level exams. Because the algorithm placed a disproportionately higher weight on the historical performance of schools (which correlated with how rich the schools were), it was biased unfairly[1] against students from poorer backgrounds.

Unfortunately, such algorithmic biases are not uncommon. The same year, in the midst of the Black Lives Matter protests occurring across the United States, an African American man, accused of shoplifting, was wrongfully arrested by the Detroit police due to an erroneous match[2]

[1] British Grading Debacle Shows Pitfalls of Automating Government - The New York Times
[2] Facial Recognition Leads To False Arrest Of Black Man In Detroit : NPR

from a facial recognition algorithm. Clearly, the danger with such biased systems is that they contribute to social and political imbalances and reinforce inequalities based on characteristics such as race, socioeconomic status, gender, sexual identity, and location.

Researchers[3] describe five general harms that AI systems can create:

1. **Allocation harm** occurs when systems extend or withhold opportunities, resources, or information. There is a risk of this type of harm when AI systems are used to make predictions and decisions about how individuals qualify for things that can impact a person's livelihood (e.g., an AI system used in hiring that withholds employment opportunities for women).
2. **Quality of service harm** occurs when there is disproportionate product failure, and a system does not work as well for one person as another (e.g., a speech recognition system that does not work well for individuals belonging to certain sociolects).
3. **Stereotyping harm** occurs when systems reinforce existing societal stereotypes (e.g., a translator that associates 'she' to a nurse and 'he' to a doctor when translating from a gender-neutral language to English).
4. **Denigration harm** occurs when a system is derogatory or offensive to a subset of users or to all users (e.g., a face recognition system misclassifying an African American as a primate).
5. **Over/under representation harm** occurs when systems either overrepresent, underrepresent, or even completely exclude certain groups or subpopulations (e.g., a search algorithm that returns more images of men than women when a 'CEO' query is used).

To address biases proactively, here are a few questions to think through:

- Who are the stakeholder populations that will be affected by the AI system, and how would the AI impact marginalized groups?
- How should the community of impacted stakeholders be involved to define fair outcomes?
- What is the composition of the training data? What is in it, what is missing? For example, if an AI is only trained on certain facial characteristics (e.g., light-skinned faces) for facial recognition systems, then it may treat other individuals who do not share those same characteristics differently.

[3] Fairness-related harms in AI systems - Microsoft Research

What are the ethical concerns associated with the general adoption of AI?

- Has the model been tested & benchmarked against affected subgroups? Are there disproportionate errors across subgroups?
- Have the ground truth labelers (i.e., those who are responsible for labeling datasets) been trained to reach a high level of domain expertise and overcome personal biases?
- How should users be informed about the limitations of the AI, so as to minimize overreliance?

While real issues exist, AI is not created with the intent to cause or introduce harm. In fact, many applications have the potential to positively inform the ways people live their everyday lives, from helping to fight poverty[4], to improving the agency of older adults[5], to assisting people with low vision[6] in learning more about their physical surroundings.

Innovation should therefore not be stifled. In fact, these challenges offer valuable surface areas for fueling responsible innovation. However, it is now more important than ever to have in place key principles[7], governance structures[8], and robust development processes, that enable the responsible development of AI technologies.

Arathi Sethumadhavan is the Head of Research for Ethics & Society at Microsoft, where she works at the intersection of research, ethics, and product innovation. She focuses on AI and emerging technologies such as computer vision, natural language processing, mixed reality, and intelligent agents. She is also a recent Fellow at the World Economic Forum, where she worked on unlocking opportunities for positive impact with AI.

JOÃO AZEVEDO ABREU

Philosophy Instructor and Researcher
IVYNENT EDUCATION • BRAZIL AND CHINA

Some of the most urgent ethical concerns regarding the massive power and general adoption of AI come down to issues of bias. If bias, in general, has been acknowledged as a major social problem, it represents an even greater challenge in the realm of AI, particularly because it is unclear how much control humans can hope to maintain over AI systems in the future - or even nowadays.

[4] Using satellites and AI to help fight poverty in Africa | Stanford News
[5] WEF_AI_and_Ageing_Workshop_Report_2021.pdf (weforum.org)
[6] Seeing AI App from Microsoft
[7] Responsible AI principles from Microsoft
[8] Our approach to responsible AI at Microsoft

What are the ethical concerns associated with the general adoption of AI?

As it is broadly understood, bias takes the form of unfair treatment of individual stakeholders and/or certain groups, resulting in ethnic, religious, and gender discrimination. Among the best-known examples are those related to legal decisions: Based on AI, formerly imprisoned African American candidates for probation were labelled higher-risk without subsequently committing new offenses, whereas their white counterparts were labelled lower-risk despite eventually re-offending[9]. This is an example in which an AI-based legal deliberation has resulted in racism.

There is also the case of automated recruiting processes. A new team of interns and trainees are selected according to an algorithm based on data of successful people in previous generations when unfair opportunities were given by human recruiters to candidates of a certain race, gender, sexual orientation, and age range. In this case, the impact of bias amounts to the reinforcement of racism and/or sexism and/or homophobia, and/or ageism.

To meet ethical expectations, AI needs an enhancement of the human/technical balance in its processes.

- João Azevedo Abreu

We can also refer to cases where opportunities for smaller businesses are diminished by recommendation systems. Such is the case when an online bookstore offers one of its frequent customers a new list of titles that other similar customers have purchased. This list might discriminate against smaller publishers and not-so-famous authors. It is unclear whether this kind of discrimination is as harmful as those related to gender and race, mentioned above. Yet the example of an online bookstore does serve as evidence against the argument that AI is as biased as society already is; it demonstrates that AI has been making bias even more pervasive than it was in the pre-AI era.

[9] Ebert, Alexandra: 2020: 'Why Bias in AI is a Problem & Why Business Leaders Should Care

What are the ethical concerns associated with the general adoption of AI?

Some might wonder whether it is desirable, or even possible, to remove bias from the system entirely[10]. Some even claim that straightforward measures such as the removal of classes would be a naïve approach to debiasing[11] on the grounds that the accuracy of data would be compromised by the absence of details of sex, race, etc.

Yet, there are practices and methods that might contribute to countering and reducing bias. Some of them could be implemented on the technical side. That is, the AI system should operate in such a way as to allow immediate redesign according to the identification of potential bias. Mechanisms need to be developed to identify the processes most likely to result in unfair discrimination and enable their monitoring.

However obvious the necessity of purely technical transformation is, what AI needs most in order to meet ethical expectations is an enhancement of the human/technical balance in its processes. That balance can be reached by a variety of human actions on AI. For instance, AI-guided processes could be accompanied by ethical-quality control, by means of third-party auditing. In addition, diversity can be promoted in human-driven processes, particularly among software designing teams and organizations.

Regardless of the intrateam diversity among AI designers and other technicians, it is hard to ignore one of the most necessary human changes in the professional world, namely, the increase of ethical awareness among software designers. Such a measure faces at least two considerable barriers. First, there is the political/cultural factor, as some countries do not perceive privacy as vital as some others do. There is also a philosophical difficulty. Take the key concept of 'fairness' which turns out to be puzzling enough in itself. Contemporary professional philosophy has provided elaborate and influential accounts of fairness[12]. Yet, the mere fact that theoretical sophistication is still required in tackling fairness, combined with the sheer absence of consensus on which practices count as universally fair, can be taken as evidence of the difficulties facing this potential measure.

Despite the lack of a straightforward notion of fairness, however, promising ethical changes can be pursued on the basis of a distinction between discrimination in general and bias in particular, with some level of pragmatism: even admitting that algorithms are essentially supposed to be discriminatory, we could filter out all those algorithms whose decisions are most likely to be interpreted as unfair by some stakeholders.

Any combination of the above debiasing measures, technical and human, would help de-bias automated systems and increase the control

[10] Coeckelbergh, Mark: 2020: AI Ethics. MIT Press
[11] Dilgemani, Cem: 2020: 'Bias in AI: What it is, Types, Examples & 6 Ways to Fix it in 2022'
[12] Rawls, John: 1971: A Theory of Justice. Belknap Press.

over AI. Although bias has always been a problem for human societies, AI bias might have a wider impact than any of the forms of pre-technological bias and, therefore, much harder to control.

Could humanity lose control of the 'global AI' system? In a certain way, the control may have been lost already. It is far from clear, to say the least, whether researchers, technicians, lawyers, and other professionals have had access to all cases of biased decisions based on AI, even though we can be hopeful that a substantial number of such errors have been at least detected and (can be) fixed.

That said, the lack of control will be more pronounced if humans fail to address bias in a timely fashion, which might be an unattainable goal *if the creation of sophisticated mechanisms of bias detection and fixing is not as fast and effective as those mechanisms that produce or enable biased decisions in the first place*.

Yet the lack of control will be even more likely and extreme if we adopt a conservative view according to which bias in AI simply reflects bias embedded in society and, therefore, if there is something to be done to counter bias, it should be pursued in general education and other domains not necessarily overlapping with AI design. Still, there is no reason to doubt that these two purposes, *ethical development in general* and *ethical AI in particular,* could and should be pursued simultaneously.

João Azevedo Abreu has held scholarly positions in the UK, China, and his home country Brazil. He's passionate about making philosophy relevant to the corporate world and other non-academic audiences.

ARUNA KOLLURU

Chief Technologist, AI, APJ
DELL TECHNOLOGIES • AUSTRALIA

A couple of examples that evidently show how AI systems can be biased are Amazon's Recruitment tool[13] and the US healthcare system[14].

According to Amazon, one algorithm used for hiring employees was trained on resumes submitted over the past decade. It was found in 2015 that the algorithm had a gender bias against women[15]. The reason for the bias was that most applicants in the past ten years were men. So, the model was trained to be biased. Women are underrepresented in the technology industry. Despite various efforts from technology companies,

[13] Amazon's sexist AI recruiting tool: how did it go so wrong? Becoming Human: AI Magazine
[14] A Health Care Algorithm Offered Less Care to Black Patients | WIRED
[15] Amazon scraps secret AI recruiting tool that showed bias against women | Reuters

the percentage of women employed in the industry is in the low twenties. Biased AI algorithms will discourage women from taking up technology roles and increase the underrepresentation.

And then coming to the US healthcare system example, recent studies show that common medical risk prediction algorithms used in hospitals In the US have racial biases. The algorithm was intended that the algorithm would identify which patients needed extra medical care, but later it was revealed that the algorithm's results favoured white over black patients. This algorithm negatively impacts the level of care for black patients.

You must have heard the saying garbage in, garbage out. This also applies to AI. An AI system is only as good as the data that feeds it. The ability of AI to make unbiased decisions without implicit or explicit assumptions about race, gender, or other biases relies on the quality of the training dataset.

AI is as good as the data it receives and people are the ones who create and provide the data. Humans are inherently biased and it is evident that the number of human bias cases discovered is increasing. Ultimately, humans are the ones creating the biased information, while humans and algorithms are checking the information to spot and correct bias. So, in reality, AI will never become 100% unbiased because of the inherent biases we humans have. All we can do is minimize the bias.

AI is as good as the data it receives and people are the ones feeding it.

- Aruna Kolluru

The first step towards minimizing the bias is being aware of it. Building diverse data teams able to spot the risk of bias in the data is where we need to begin. The data team focuses on the problem being solved and the involved stakeholders and also helps shape the required datasets and identify the data points that must be checked for bias. The team analyzes the end-users of the AI model and tries to understand how the dataset can be biased towards them. As model-building is an iterative process we need to gather the feedback and loop it into the development

of the next version of the model. Google's What-If, Microsoft's Fairlearn, and IBM's Fairness 360 are some tools that help detect biases.

We have seen many quotes from technology leaders about the threats of AI. I particularly think statements from Stephen Hawking and Elon Musk are very intriguing. Stephen Hawkings said[16] *"The development of full artificial intelligence could spell the end of the human race….It would take off on its own, and re-design itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, couldn't compete and would be superseded."* Elon Musk said[17], *"I have exposure to the most cutting edge AI, and I think people should be really concerned by it."*

I don't think at least with the current and predicted advancements we are even close to humanity losing control of a 'Global AI system'. But, there are, indeed, inherent risks associated with AI. Risks come with the autonomous nature and the decision-making capabilities of AI systems. In some ways, the OECD's nonbinding recommendations on AI are a start to mitigate these risks, as they try to frame AI in a way that focuses on the collective good while, at the same time, acknowledging the risks.

Aruna Kolluru is passionate about business, technology, and emerging technologies. She helps organizations to visualize the art of possible, architect, and build solutions with emerging technologies. Before joining EMC & Dell Technologies, Aruna held various senior roles for IBM & Oracle. She has over 21 years of experience in the industry on varying technologies including AI, Blockchain, Deep Learning, Computer Vision, IOT Big Data & Analytics, and Enterprise Architecture.

VEENA CALAMBUR

Data Scientist & AI Ethicist
WORKDAY • USA

There has been a massive acceleration in the adoption of Artificial Intelligence (AI) in all aspects of our lives. But how much can we trust these ubiquitous AI systems? Over the past several years we have seen several examples of AI that have gone wrong across multiple fields and industries from racist facial recognition[18], discriminatory hiring[19], and credit scoring[20] predictive models to even genocidal chatbots[21]. As a result, there are emerging concerns around the trustworthiness, reliability,

[16] A Short Speech on Artificial Intelligence (thevideoink.com)
[17] Elon Musk has a complex relationship with the A.I. community (cnbc.com)
[18] NIST Study Evaluates Effects of Race, Age, Sex on Face Recognition Software | NIST
[19] Amazon scraps secret AI recruiting tool that showed bias against women | Reuters
[20] Apple Card algorithm sparks gender bias inquiry - The Washington Post
[21] Microsoft Chat Bot Goes On Racist, Genocidal Twitter Rampage | HuffPost Impact

fairness, privacy, transparency, and autonomy of AI, that we must address to minimize societal harm.

A natural question to pose is how can AI be biased? Many have the notion that AI algorithms grounded in facts and data must be the key to fighting against biased human judgement and decision-making. Unfortunately, the way systems are designed to create and record data points that are used to train these algorithms are ultimately human constructs that can be impacted by both implicit and systemic biases.

For example, medical algorithms that are trained on longitudinal electronic health records are fully dependent on how individual healthcare professionals enter the records. Several studies show that there is evidence of implicit biases impacting healthcare professionals in medical practice[22]. If a doctor has implicit biases about taking a female patient's pain less seriously than a male patient's pain, leading to fewer pain medication prescriptions, then an algorithm analyzing pain medication treatment patterns may pick up on these biases and further recommend limiting treatments to female patients.

Without intervention, AI algorithms can learn, codify and perpetuate biases long into the future.

- Veena Calambur

Beyond the biases perpetuated by individuals, one of the significant drivers of AI bias comes from systemic biases that are embedded across several societal institutions. For instance, the history of the United States housing industry is rife with racist policies such as redlining that barred African Americans from homeownership in certain neighborhoods[23]. Even though these policies have been outlawed for over fifty years, we can still observe their ramifications. African American homeownership is still significantly lower and banks are still much more likely to mortgage loans to Caucasian borrowers in certain regions. AI-based lending

[22] Implicit bias in healthcare professionals: a systematic review - PMC (nih.gov)
[23] Redlining was banned 50 years ago. It's still hurting minorities today. - The Washington Post

What are the ethical concerns associated with the general adoption of AI?

algorithms trained on historical housing data can learn the past racist policies and accelerate housing discrimination. Without any kind of intervention, AI algorithms can learn, codify and perpetuate biases long into the future all under the guise of objectivity.

If humans engage in biased judgement or decision-making that leads to unfair outcomes why does AI pose such a risk to humanity? AI technologies have the ability to automate biased decision-making at scale and this could lead to widespread *algorithmic discrimination*. Given the presence of implicit and systemic biases in many of our recorded data systems, algorithms that are trained on these datasets are often learning 'standard' scenarios and generalize behaviors that are not necessarily representative of the full population[24]. So any individuals or sub-populations that deviate from the learned norms codified in the algorithm can experience negative impact and harm as a result.

This is further exacerbated by the lack of transparency and human autonomy over these AI systems. Some of the most notorious examples of AI causing societal harms occur when an AI system is deployed to automatically trigger decisions without any oversight or human control. For example, social media newsfeed recommendation algorithms have been shown to lead to massive political polarization and a rise in mis/disinformation globally based on automated content suggestions[25].

We can work to improve AI systems to ensure they are more equitable and inclusive.

- Veena Calambur

The role of AI is often obscured in these situations which is a serious concern of lack of consent of individuals influenced or impacted by AI. Even in cases where an individual is aware of the role of the algorithm in the prediction and wants or needs to contest the decision, it may be difficult or even impossible to explain the reasoning behind the decision - due to the complex and even black-box nature of many AI algorithms.

How can we go about addressing these very serious concerns? While it is impossible to claim we can remove all forms of bias, we can work to

[24] Automating Inequality | Guide books (acm.org)
[25] Polarization and Fake News: ACM Transactions on the Web: Vol 13, No 2

improve AI systems to ensure they are more equitable and inclusive. The simplest way to get started is by asking the right questions. There are several AI Ethics checklist resources available but below are a few key questions to get started and to check in periodically throughout the development and lifecycle of AI.

1. What is the purpose and intended use of the AI system? Are there ways that it can be misused or used unintentionally?
2. Who are the 'interactors' of the AI system? More, specifically, [a] Who are the primary intended end-user of the AI? Do they have adequate understanding and autonomous controls over the system? And [b] Who may be influenced or impacted by the presence of the AI system directly or indirectly? Are they / their needs well represented in the data and algorithm?
3. Are the AI 'creators' (i.e. data scientists, machine learning engineers, supporting operations, business sponsors) aware of ethical AI issues? And [a] Are they building in proper data and algorithm inspections during the development and monitoring of AI systems? [b] Are the developer teams diverse and support a culture of responsible AI?

There are a few steps to add to the AI algorithm development process to identify and mitigate bias. In the exploratory analysis, AI practitioners should analyze how well represented the data population is when compared to the target population of the AI solution. They should also analyze the target outcome across key sub-populations based on demographics, protected characteristics, or customer segments. Upon training the AI model, it should be evaluated for bias prior to deployment. Bias evaluations should include *model performance disparities checks*, *model explanation assessments*, and *model outcome or endpoints reviews delineated by sub-populations*. If any bias is found, mitigation methods like matching algorithms, sampling techniques, or debiasing algorithms should be applied.

To learn more about Ethical AI we should continue to review additional materials, particularly from marginalized communities who have already documented their experiences with AI and algorithmic discrimination. We should all develop a culture and governance to support responsible AI development.

Veena Calambur is a data scientist and AI ethicist. She has worked to research and adapt data analytics and machine learning capabilities. She will be starting a role in Workday's Machine Learning Trust Program. In the past Veena has worked as data scientist at Pfizer where she pioneered their enterprise-wide Responsible AI strategy and toolkit and as a decision analytics associate at ZS Associates. Veena has her bachelor's degree in Information Science and Statistics from Cornell University and is pursuing a PhD in Information Science at Drexel University.

What are the ethical concerns associated with the general adoption of AI?

ENRICO PANAI

AI and Data Ethicist
ÉTHICIENS DU NUMÉRIQUE / BEETHICAL.BE • FRANCE

The power of AI is relative to its capacity for agency, which is not monolithic, but a system composed of autonomous agents, or multi-agent systems. In practice, many of the agents that make up an AI system have an extraordinary capacity for action over a very limited field. The actions they generate may have limited consequences. However, the sum of many morally-neutral actions can lead to a large morally-loaded consequence. The final consequence is not necessarily negative.

Take ecologically responsible behaviour, for example. If I turn off all the lights at home for one hour a week, this has no impact on the global environmental system. And perhaps it has no visible impact on my electricity bill either. However, if everyone in the world does the same thing, the action becomes relevant and therefore morally loaded (in this case positively). This is an example of a multi-agent system composed of people, but there are systems composed of AI agents only or some hybrids (humans, organizations, AI). Our fear comes from having no idea how a set of small morally neutral or morally negligible interactions result in a set of enormous morally-loaded consequences.

The sum of many morally-neutral actions can lead to a large morally-loaded consequence.

- Enrico Panai

An AI agent can perform a task very well, much better than we humans can; but perhaps it uses data that is not of the required quality to make a decision. Any dataset may have been biased during the recording, the storage, the processing, etc. The quality of the data may be diminished *statistically* (if the sample is not representative), *technically* (if sensors or recording technologies were not reliable), or *cognitively* (because people 'injected' cognitive biases into the dataset). Cognitive biases are particularly interesting because they have no defined perimeters; they overlap one with another. But perhaps that is what makes us human.

What are the ethical concerns associated with the general adoption of AI?

Let me say that in other words. We are humans so our decisions are naturally biased. And some biases are absolutely natural. But when we see them in AI they become unbearable. So we try to eradicate them from their source: humans. This means that we try to change ourselves to make AI better. We are afraid that machines are like us, imperfect. In doing so we underestimate that the process of reducing prejudices is a much more complex social process. And in any case, some biases cannot be removed even from a dataset or an AI system.

What one must try to do is to mitigate the risks in a continuous and evolutionary process. A bit like driving a car, you adjust the steering wheel all the time to stay on track (but this is simply acceptance of uncertainty and a common sense that starts from Socrates and goes all the way to extreme programming).

In short, what we are afraid of is not the speed of execution of AI systems, but the unpredictability of the consequences. There are risks, but there are also advantages.

What scares us is not driving fast, but driving fast while blindfolded.

- Enrico Panai

And from an ethical point of view, it is also wrong not to exploit the positive opportunities of AI. The point is to be able to do so with an appropriate ethical awareness. What scares us is not driving fast, but driving fast while blindfolded.

Enrico Panai is an Information and Data Ethicist and a Human Information Interaction Specialist. He is the founder of the French consultancies "Éthiciens du numérique" and BeEthical.be and a member of the French Standardisation Committee for AI, and ForHumanity's Fellow.

What are the ethical concerns associated with the general adoption of AI?

MIKE TAMIR

Chief ML Scientist, Head of ML/AI
SUSQUEHANNA INTERNATIONAL GROUP • USA

It has become abundantly clear over the past decade that expecting predictions, estimations, or recommendations generated by machines to be unbiased because machines are not (biased) humans, is naive. Machines learn from data and data is often biased.

It is good news that the major tech companies (where much AI power has centralized) seem not only to recognize this problem but have programs aimed at combating the risks of data bias to ML and the impact that such bias can lead to in practical applications. However, pressure for transparency, fairness and potentially regulatory oversight in this area will need to continue.

Machines learn from data and data is often biased.

One pattern that can leave us optimistic is the trend of open sharing of both (non-private) data as well as pre-trained models. This has not been universal (ironically Open AI was controversially less than open with sharing its large language models at first). While larger technology companies do have large amounts of private data that is not shared, open sourcing of pre-trained models has had the effect of democratizing the applications of such models, with positive results for research and applications. An added benefit of this is that e.g. for language model applications, less data is required to customize a general model to specialized use cases, eliminating barriers to entry.

Of course, large tech companies still have huge advantages in terms of non-open (often personally private) data access and market exposure. Regulation to ensure that such advantages do not stifle growth in a healthy (data) economy deserves serious consideration.

Mike Tamir, PhD is a data science leader, specializing in deep learning, NLP, and distributed scalable machine learning. Mike is experienced in delivering data products for use cases including text comprehension, image recognition, recommender systems, targeted advertising, forecasting, user understanding, and customer analytics. He is a pioneer in developing training programs in industry-focused machine learning and data science techniques.

What are the ethical concerns associated with the general adoption of AI?

ARUNA PATTAM

Head of AI & Data Science - Asia Pacific, Middle East, and Japan
A GLOBAL TECHNOLOGY COMPANY • AUSTRALIA

AI can have a massive impact on our lives. As it grows more powerful and widespread, the ethical concerns around it become more important. Bias is a big issue in AI. As algorithms get smarter, they can learn our biases. This means that the systems can end up perpetuating them.

There are many ways that AI systems can become biased, but here is a list of main ones that we should be mindful of:

- The data used to train AI can be biased. If the data is not representative of the whole population, the AI will be biased too. For example, if a system is trained on a data set of men, it will be biased against women.
- People who design AI systems can also be biased. They may inadvertently include their own biases in the algorithms. For example, a programmer may design a system that gives a higher ranking to men than women.
- People who use AI systems can also be biased. They may use the system in a way that perpetuates their biases. For example, if a system is used to decide who to hire, the people who use it may be more likely to choose men over women.
- AI can also get tainted once it's introduced into the world since, if the new data being fed into the live system is prejudiced, the algorithm may learn from it and introduce bias. For example, a self-learning algorithm that is supposed to optimize a factory process could start discriminating against a particular ethnic group if it was fed discriminatory data.

The implications of allowing bias in AI systems can be significant, impacting individuals and societies. Below are some real-life instances of bias that can have a significant impact if not addressed.

- **Bias in Hiring**. Many AI systems were used to help with the hiring process. However, in some cases, these systems were biased against women as they were more likely to recommend a man for a job than a woman, even if the woman was more qualified[26].
- **Bias in Credit**. AI systems are also often used to decide who to give credit to. However, these systems are often biased[27] against certain groups of people, such as gender, minorities, or low-income people.

[26] Amazon's AI Was Biased Against Women (businessinsider.com)
[27] Apple Card algorithm sparks gender bias inquiry - The Washington Post

What are the ethical concerns associated with the general adoption of AI?

This can lead to people not getting the loans they deserve or getting loans at a higher interest rate than they should.

- **Bias in Policing**. AI is also being used to help with policing. However, these systems can also be biased[28]. For example, a recent study[29] found that a system designed to predict criminality was biased against black people. The algorithm was more likely to incorrectly flag a black person as being a criminal than a white person.
- **Bias in Healthcare**. AI is also being used to help with healthcare. However, these systems can also be biased[30]. A study found that a system designed to diagnose skin cancer was more likely to misdiagnose a black person than a white person[31].

We also need to be mindful of other risks that come with the massive power and general adoption of AI. For example, AI could be used to manipulate people: consider a scenario where a government uses AI to control the thoughts of its citizens or, a company that uses AI to get people to buy more products. AI could also be used to destroy people or things. For example, a terrorist group could use AI to create a bomb that is very difficult to detect. Also, AI could be used to steal people's personal information. For example, a hacker could use AI to steal your credit card details.

This raises the next question - how can we remove bias from AI systems? There are several ways to consider, for example:

- We can make sure that the data used to train AI is representative of the whole population.
- We can make sure that the people who design and use the AI systems are aware of their own biases and try to avoid them.
- We can have a diverse group of people working on and using AI systems. This will help to expose any biases that may be present in the system.
- We can test AI systems for bias before releasing them into the world.
- We can use Machine Learning to identify and correct bias in AI systems.
- We can create a 'watchdog' group to monitor AI systems for bias.

[28] What Happens When Police Use AI to Predict and Prevent Crime? - JSTOR Daily
[29] Discriminating algorithms: 5 times AI showed prejudice | New Scientist
[30] Millions of black people affected by racial bias in health-care algorithms (nature.com)
[31] Characteristics of publicly available skin cancer image datasets - The Lancet Digital Health

What are the ethical concerns associated with the general adoption of AI?

- we can create a Code of Ethics for AI. This would be a set of guidelines that designers and users of AI systems should follow to ensure that their systems are not biased.

Legislation has a bigger role to play here. We need regulatory bodies that control how AI is used. Just as how we have General Data Protection Regulation (GDPR) to protect people's data, we need something to ensure that AI is being used for good and not for evil.

To conclude, AI has the potential to do a lot of good in the world. However, we need to be aware of the associated risks. We need to make sure that we are using it ethically and responsibly. Otherwise, we could end up causing a lot of harm to ourselves and others.

Aruna Pattam is a Global AI Thought Leader with 22+ years of experience in data analytics, AI, and ML. Also, a speaker, mentor, blogger, vlogger, and women advocate with the goal of educating and providing awareness on AI to the business and the wider community. The views expressed are personal.

THERÉSE SVENSSON

Data Science & AI Ethics Solution Specialist
IBM • SWEDEN

When the risks of Artificial Intelligence are discussed, my experience is that the discussion quickly touches upon the AI singularity or the use of AI technology to support certain problematic use cases such as mass surveillance, raising concerns regarding violations of basic human rights. However, I would argue that these are not the risks over which we need to be the most concerned. In my view, the true risk of AI originates from the fact that when it is being used in uninformed or careless ways it can lead to unfair or even inaccurate results.

When it comes to AI singularity, we are still too far from achieving general AI to let this concern keep the masses up at night. The application of AI in ways that are harmful to a fair and democratic society is a bigger problem than solely the performance and behavior of AI algorithms. Most research and new technologies could be used in ways they were initially not intended to and finally harm our societies. With that said, there is certainly the need for regulation and worldwide agreement to prevent the use of AI in unethical and unwanted ways. Such regulations are already being discussed and will likely be implemented soon, for example, in the form of the proposed EU AI Act.

To understand why uninformed and careless use of AI systems poses a risk, let us consider the following, hypothetical, scenario: You

have created a decision support system that is designed to help healthcare professionals to identify possible signals of addiction among patients. Your system is becoming an appreciated tool and is adopted within the healthcare industry. After some time, you learn that a patient was wrongfully accused of being an addict all due to the incorrect prediction of your decision support system. You knew all along that your system was not always correct in its prediction, as is the nature of a predictive model, but the system was meant only to *support* the healthcare professional, and *not to replace* his or her judgement. This hypothetical example closely resembles an actual situation[32] that did occur within the healthcare industry. One of the problems, in this case, was that the system was not used in the way it was intended to, and somewhere along the way, it transformed from being a decision *support* system to *making* the final decision on behalf of the healthcare professionals.

To ensure proper use, AI systems should come with a content declaration describing the intended use and limitations.

- Therése Svensson

To take this addiction scenario further, let us consider the reasons for the incorrect prediction. In our thought example, the patient got high addiction scores due to the historical prescriptions being given by several different doctors, indicating that the patient wasn't seeing one doctor solely, a behaviour often found among addicts. In reality, several of these prescriptions were given by veterinaries for her ill pets. The AI model did not distinguish her prescriptions from her pets' as they were all prescribed to her. If the addiction score had come with an explanation of what influenced the score, the healthcare professional would have had a chance to discover that the strongest driver of the high score was caused by a data issue. Therefore, the model should not have been trusted in this particular case.

[32] A Drug Addiction Risk Algorithm and Its Grim Toll on Chronic Pain Sufferers | WIRED

What are the ethical concerns associated with the general adoption of AI?

To help ensure that AI systems are used the way they are intended and to highlight the limitation of these, a proposed solution is to define a *content declaration* attached to every AI model. This would contain information about how the model was created, which data was used, the expected accuracy of the model, and its intended usage.

My point is that if the healthcare professionals had been better informed on the intended usage of the model, its limitations, and how it worked, this case of unfair and inaccurate treatment of the patient would have been prevented. We grow up learning that the world isn't a fair place but that shouldn't stop us from trying to make it one. AI technology has a great potential to improve the lives of many which is why it is important to use it responsibly.

Therése Svensson is a Technical Specialist within IBM Data & AI. She has previous experience at IBM Expert Labs where she was working as a consultant in the area of Data Science, Predictive Analytics, and Ethical AI.

SAMIRAN GHOSH

Senior Advisor
VARIOUS STARTUPS • INDIA

In 2020, a New England Journal of Medicine study[33] looked at more than 10,000 patients throughout America and suggested that pulse oximeters used overestimated blood-oxygen saturation more frequently in black people than white. Normal oxygen saturation is at 92-96%. In the study, some patients who registered normal levels per the pulse oximetry had a saturation as recorded by the arterial blood-gas measure (the most accurate saturation measure) of less than 88%. For black participants, this occurred 12% of the time – 3X the rate for white participants. It is the difference between being admitted to the hospital and being sent home.

In 2013, the Journal of the American Medical Association found[34] that women in four American regions had a 29% higher risk of their hip implants failing within three years of hip-replacement operations. A similar study[35] from 2019 found that women were twice as likely to experience complications from cardiac device implants, such as pacemakers, within 90 days of such procedures. In both instants, the failures of device-makers to recognise physical differences between male and female body types were to blame.

[33] Racial Bias in Pulse Oximetry Measurement - PMC (nih.gov)
[34] Hip Implant Failure for Men and Women (researchgate.net)
[35] Sex Differences in Acute Complications of Cardiac Implantable Electronic Devices - ahajournals.org

What are the ethical concerns associated with the general adoption of AI?

While such findings are troubling in themselves, they also illustrate the deep-seated prevalence of bias. Bias is omnipresent. It may be machine learning bias (Selection, Exclusion, Observer, etc.), psychology, or cognitive (Dunning-Kruger Effect, Gambler's Fallacy, Ben Franklin Effect, etc.). We also have new ones like the Google Effect[36] and the IKEA Effect[37]. Companies like NetFlix, Nike, Google, and others effectively use these biases to deliver what they want us to see using the algorithms that run their platforms and marketplaces.

The more diverse the training data, the more accurate the AI models.

- Samiran Ghosh

Some questions arise from this. Will AI be a threat to humanity? Unlikely in the next 30-40 years at least. We are still far away from Artificial General Intelligence. So, is AI a threat to our jobs? Very much so. Approx. 45% of low education workers will be at risk of technological unemployment by 2030. And can we trust the judgment of AI systems? Not yet is the best answer for now. And for several reasons – bias being one of them.

A classic example of bias in the Computational Linguistics world is the well-known *king – man + woman = queen* equation[38]. In this, if we substitute the word 'doctor' for the word 'king' we get 'nurse' as the female equivalent of the 'doctor'. This undesired result reflects existing gender biases in our society and history. If doctors are generally male and nurses are usually female in most available texts, that's what our model will understand.

Another well-known example of an image classification mistake is when Google Photos misclassified black people as gorillas[39]. While a single misclassification of this sort may not substantially impact the overall evaluation metrics, it is a sensitive issue.

[36] Google effect - Wikipedia
[37] IKEA effect - Wikipedia
[38] The Marvelous Mathematics of Computational Linguistics | MIT Technology Review
[39] Google apologises for Photos app's racist blunder - BBC News

What are the ethical concerns associated with the general adoption of AI?

So how do we solve this? AI Developers need large, carefully labelled, unbiased datasets to train Deep Neural Networks. The more diverse the training data, the more accurate the AI models. The problem is that gathering and labelling sufficiently large and rich datasets, that may contain a few thousand to tens of millions of elements, is a time-consuming and very expensive process.

The use of synthetic data could be a possible answer. For example, a single image costs about $6 from a labelling service. However, we can generate the same labelled image for as low as six cents in the lab. But, cost savings are just one aspect of the issue. Synthetic data is also key in dealing with privacy issues and reducing bias by ensuring you have the data diversity to represent the real world. Furthermore, synthetic data is sometimes better than real-world data because they are automatically labelled and can deliberately include unusual but critical data properties.

Though the synthetic data sector is only a few years old, more than 50 companies already provide synthetic data. Each has its unique quality, often focusing on a particular vertical market or technique. For example, a handful of these companies specialise in health care uses. A half dozen offer open-source tools or datasets, including the Synthetic Data Vault, a set of libraries, projects, and tutorials developed at MIT.

The above is an excellent start to creating a more level playing field for 'fair' AI Systems.

Samiran Ghosh is a recognized thought leader in the AI/Blockchain space. Has held leadership positions with McKinsey, IBM, Microsoft, and TCS. Is a Fellow of the Digital Euro Association and Royal Society of Arts. On the global panel of MIT Tech Review, a TEDx speaker, and a member of the Forbes Tech Council. Worked for a year on the world's largest identity program implemented in India - Aadhaar. A couple of books he co-authored were amongst the top 20 new fintech books to read in 2020. He also co-hosts a popular tech podcast called 3 Techies Banter.

NETANEL ELIAV

Chief Executive Officer & Technology Development Specialist
SIGHTBIT LTD • ISRAEL

The potential for Artificial Intelligence to impact humanity is immense. AI is bringing several changes to how we live, work, and play. There is already a profound effect on how we relate to one another, as well as on how our economies and societies are organized.

AI helps us make better decisions by providing insights into our data and the world. It also helps us solve problems that have eluded human understanding for centuries. In addition, AI can improve the quality of life

for people with disabilities or those who lack access to education or healthcare due to geography or income. As AI becomes more and more advanced, it will do things we thought were only possible for humans.

There is clearly a need for transparency and accountability in AI technologies.

- Netanel Eliav

However, there are ethical concerns and potential threats because of how AI has been programmed and because of the way we are using it. For instance:

- AI could be used for surveillance, which can be seen as an invasion of privacy and a violation of human rights.
- AI could be used for warfare.
- The wide adoption of AI could drive unemployment as people are replaced by more efficient software or robots.
- AI may have biases, which can lead to discrimination and inequality.

Bias is a big issue with a major impact on society and thus it must be handled by the designers and developers of AI systems. There are many ways to remove bias from an AI system, for example by using a diverse team for data gathering and training; or, by using algorithms or techniques able to identify both explicit and implicit biases.

The risks of developing and using advanced Artificial Intelligence systems are inherent and inevitable as this technology will one day become better than humans, in everything. But what are these risks? There are many opinions on this topic, but in general, these are three main schools of thought:

- The first school of thought believes that AI will eventually replace humans and will have a negative impact on our society.
- The second believes that AI will create new opportunities and bring a positive impact on our society.
- The third believes that humanity can and will lose control, leading to an AI system that will, eventually, take over the world.

What are the ethical concerns associated with the general adoption of AI?

We have seen movies and TV series about this, e.g. Skynet[40], but the question is: could it happen? And if it happens, how will it affect humanity? As AI systems become more complex and powerful, they might make decisions that are not in our best interest, and thus, we might indeed lose control over them. For example, they could misuse our data or attack us in unpredictable ways.

To address these risks we must make sure that we don't put too much power in AI's hands, and we should not trust these systems blindly. There is clearly a need for transparency and accountability in AI technologies and there should be a systematic way to *measure, regulate, and audit AI systems*. We need to establish standards that define what AI is, how it should be developed and used, and how to evaluate its performance and impact.

Netanel Eliav is a CEO and Founder at SightBit - An Artificial Intelligence Startup Using Deep Learning and Computer Vision to Save Lives. He is a former Product Manager and ex-Technology Tech Lead Specialist at the Office of the Prime Minister.

MAYTE HIDALGO

Head of Data & AI Center
MEDIAPRO GROUP • SPAIN

From philosophy to programming. From Plato to Python. Artificial Intelligence rises a diversity of challenges that root in *values* and go far beyond technology. The massive power and the general adoption of AI raise a series of ethical concerns and tough questions: How to make algorithmic decision-making fair? How to guarantee that AI will augment humans instead of replacing them in the quest for economic growth and prosperity? Which criteria should be applied to define the limits of human oversight over technology?

No doubt, AI initiatives are already permeating various business and social domains. The increasing impact of AI decision-making brings as many promising opportunities for value creation as question marks on how to establish an actionable ethical framework. The speed at which AI is transforming industries impacts communities and behaviors and keeps generating new dilemmas. Considering that AI adoption will accelerate in the coming years – in the form of myriads of digital services - it is important to address the major AI risks now.

[40] Skynet (Terminator) - Wikipedia

What are the ethical concerns associated with the general adoption of AI?

While AI Ethics is relatively new as an applied field of knowledge, the conversations around *ethos* were already a pillar of the Greek *polis* where our civilization started outlining societies as such. The character of beliefs and ideals shaped by communities and nations is, by definition, intangible. It crystallizes in individual and collective behaviors, holding certain attitudes and acting in particular ways. AI Ethics is for shaping a character around the application of AI, through a shared predisposition on what should be fostered or prevented when exercising AI.

This covers countless use cases across industries: from evaluating race bias of Computer Vision algorithms to guaranteeing financial credit eligibility fairness across genders; from reasonably explaining the hows and whys behind an automated decision to finding the collaborative model of virtual and human agents of customer service; from ensuring the right to intimacy and privacy across geo-localized applications to deciding the level of human supervision that should apply for an AI-driven health initiative; from guaranteeing unmanipulated dialogues between brands and individuals to protecting vulnerable population groups from discrimination…

AI Ethics is for shaping a character around the application of Artificial Intelligence.

- Mayte Hidalgo

AI Ethics must be considered in the corresponding ecosystem that comprises organizations and the communities they serve, professionals and individuals, private and public sectors. AI Ethics require direction and monitoring from both a long-term socio-economic and a day-to-day viewpoint and from the perspectives of all involved stakeholders.

Ethics are embedded within our social agreement. For instance, while the idea of *fairness* may be abstract, humans recognize and accept regulations to settle, even if some translation is needed depending on the context. From the individual subjectivity of values to the set of laws that embody nationwide criteria, there is a level of intersubjectivity, that is, people's mutual understanding around the way we define concepts such as *discrimination* or *respect for diversity*. Such agreements are paramount to harmonizing our perspectives as a collective, and often serve

as a driver to stimulate new corporate guidelines and legislations. In a simile to how those agreements formed the foundation for managing the polis, AI Ethics is now generating organizational, institutional, and social debate.

As organizations serve society, they must enrich their strategic leadership agenda with ways to manage or mitigate the risks associated with AI applications. There are a lot of examples of such risks, for instance, the case of Apple Card[41], the Twitter debate, and the Wall Street regulator's investigation.

The gender bias of the algorithm for extending Apple card credit, resulted in a series of reports by users, stating that the process was unfair: under similar conditions, men were eligible for much higher credit than women. Even Steve Wozniak shared that he had been given ten times the credit limit offered to his wife. This controversy on AI bias went far from just an academic debate: The CEO of Goldman Sachs, the firm issuing and operating the card, had to publicly clarify the reasons behind such unfair outcomes and make the process more explainable.

Another example of gender bias is Amazon's recruiting tool[42], which already in 2015 was identified not to be gender-neutral when rating candidates for various technical positions. The reason? Models had been trained over 10-years datasets of successful profiles which, on the male-dominance scenario for the period, did not equally represent female candidates.

A responsible approach to AI development requires tools for risk assessment at all levels, including the identification of potential issues regarding *bias, privacy, diversity, transparency, manipulation,* etc. During the last three years, governments worldwide have started a race to define guidelines around the ethical implications of AI. In 2019 the European Commission launched the Ethics Guidelines for Trustworthy AI, led by an independent expert group. Based on a public consultation followed by a process of gathering and analyzing views of 1,000+ participants (citizens, member states, industry experts, and academics), the AI Act[43] was launched on April 2021 – as a regulatory proposal that defines *what AI is* and identifies different levels of AI-related risks that need to be addressed.

While it is critical for organizations to adopt institutional frameworks and regulations like the AI Act, it is also of paramount importance to mobilize all players across industries, so that top-down and bottom-up approaches jointly work towards building trust in human-centric AI. To achieve this, organizations must adopt standards, processes, and

[41] How the law got it wrong with Apple Card | TechCrunch
[42] Amazon scraps secret AI recruiting tool that showed bias against women | Reuters
[43] EUR-Lex - 52021PC0206 - EN - EUR-Lex (europa.eu)

What are the ethical concerns associated with the general adoption of AI?

controls to make AI systems trustworthy and compliant. Moreover, they must promote a culture of responsible, ethical, and trustworthy AI.

Mayte Hidalgo has 20+ years in the Digital ecosystem, with a career spanning multinational organizations, startups, consulting, and entrepreneurship. She is Fulbright, Intercampus and Erasmus Scholar. Mayte started a data career in 2002 with the data-centric worldwide pioneer program IMC by Don Schultz and the school of Kellog on Integrated Marketing @Northwestern University.

JOIN THE AI ETHICS DISCUSSION **SPREAD THE WORD**

12

HOW COULD WE PREVENT THE UNREASONABLE CONCENTRATION OF AI POWER?

Big tech companies accumulate massive amounts of rich data and this gives them ‘unfair’ advantages in developing powerful, large-scale Artificial Intelligence services. Moreover, a limited number of big technology companies attract the best Data & AI talent, as they provide the data, the technologies, and the space for building innovative AI-related solutions. Could this lead to ‘data or AI monopolies’? How should we handle this risk? How can we prevent a ‘talent imbalance’ in the market?

Abhishek Gupta and Netanel Eliav share their insights.

How could we prevent the unreasonable concentration of AI power?

ABHISHEK GUPTA

Founder and Principal Researcher
MONTREAL AI ETHICS INSTITUTE • CANADA

The unreasonable concentration of power in AI comes from three components: compute, data, and talent.

In terms of compute, the more complex AI systems we choose to build, the more compute infrastructure they require to be successful. This is evidenced in the rising costs for training very large language models like Turing[1], GPT-3[2], and others that frequently touch millions of dollars. This is all but inaccessible except to the most well-funded industry and academic labs. This creates homogenization in terms of ideas that are experimented with and also a severe concentration of power, especially in the hands of a small set of people who determine what the agenda is going to be in research and development, but also in the types of products and services that concretely shape many important facets of our lives.

To draw AI talent, provide meaning and purpose beyond just compensation.

- Abhishek Gupta

The larger the systems we prioritize, the more data they need as well - this is particularly the case with some of the more recent 'internet-scale' models that require 'internet-scale' data such as the CommonCrawl[3] dataset amongst others. But, these raise issues of 'internet-scale' biases because these datasets tend to suck up the worst of humanity and amplify that through the models that use these datasets for training.

So, what can we do to combat these two negative trends? On the compute infrastructure front, there are calls from various groups in the US, for example, that are demanding a government-funded national

[1] Turing-NLG: A 17-billion-parameter language model by Microsoft - Microsoft Research
[2] GPT-3 - Wikipedia
[3] Common Crawl - Wikipedia

research cloud that would make cheap compute accessible to anyone who would like to experiment in building AI systems. This would *lower the barrier to participation* and make it easier for anyone who would like to develop AI systems. In particular, this would be a boon for all of society since we would have the potential to empower local entrepreneurs and concerned citizens to develop AI systems and services for their own communities.

Data commons and open-source datasets that are well-maintained such as Common Voice[4] from Mozilla are another way to bolster nascent, independent efforts from local communities to build AI systems that help meet their own needs without having to rely on technology that gets exported from a small geographic region to the rest of the world.

But, all of this requires significant technical expertise to make it a reality, and *nurturing talent* will play a key role if we're to combat the pernicious concentration efforts that emerge from Big Tech not only being able to offer lucrative compensation, but also large datasets and accompany access to adequate computational infrastructure that makes it exciting for talent to pursue challenging problems that further their own careers.

A counterpoint to that trend is to offer projects that can provide *meaning and purpose* beyond just the points that are raised above, as a way to draw talent. In addition, offering educational training programs and engaging in capacity building, both through community-driven programs like those offered by Masakhane[5] NLP and open-access programs and books like D2L.ai[6] offer resources that will upskill the next generation of data scientists, machine learning engineers, and software developers who can bring their expertise towards solving society's grand challenges while using AI capabilities in an ethical, safe, and inclusive manner.

New models of self-organization and community engagement have shown to be powerful levers to enact change in the world around us and being deliberate in our approach to leaning on these vectors of change can allow us to pursue lofty goals while staying true to community roots and achieving the ultimate goal of having a more fair, just, and equitable society for us all while harnessing all that AI has to offer.

Abhishek Gupta is the Founder and Principal Researcher at the Montreal AI Ethics Institute. His work focuses on applied technical and policy measures for building ethical, safe, and inclusive AI systems, specializing in the operationalization of AI ethics and its deployments in organizations and assessing and mitigating the environmental impact of these systems. He is the author of the widely read State of AI Ethics Reports and The AI Ethics Brief.

[4] Common Voice - Wikipedia
[5] Masakhane
[6] Dive into Deep Learning — Dive into Deep Learning 0.17.5 documentation (d2l.ai)

How could we prevent the unreasonable concentration of AI power?

NETANEL ELIAV

Chief Executive Officer & Technology Development Specialist
SIGHTBIT LTD • ISRAEL

AI is a powerful tool, and it can be used to do good or bad; it can help people make sense of their data, but it also has the potential to create new forms of inequality and exclusion, as it becomes easier for those with access to rich data to build better AI systems. Data is the new oil, the fuel that powers AI, and it is a valuable resource in its own right. We are generating data from our daily activities – browsing, travels, purchases, etc. More and more companies collect and store such data - for example, Google and Facebook have enormous amounts of data stored on their servers. As these companies accumulate detailed data about our world they strengthen their AI advantage over other companies and future startups.

It is extremely difficult for small companies to attract AI talent

- Netanel Eliav

So, what can we do to prevent these companies from controlling information and accumulating infinite power? The answer is a picture of us, the ones who provide the information. One way is to free our data and share it with everyone in other words - to kill our privacy. This will reduce the uniqueness of the information that companies have. Another way is to stop giving our data.

This leads us to another question: how do we prevent these companies from collecting information that we do not share? The solution is in *regulation*, as governments exist precisely to create balance when there are market failures, and today the field of data and information protection is facing problems and major concerns.

Beyond the *data imbalance*, there is also a *talent imbalance*. AI technologies are built by scientists, software developers, data scientists, and machine learning engineers and the demand for this skillset is booming. The need for AI-related skills has increased by more than 60% since last year, according to a study conducted by LinkedIn. AI researchers are not only well-paid but also have a lot of job opportunities - they can pursue a career with Big Tech companies like Google, Amazon Web Services, Microsoft, or Facebook - which all offer very competitive packages. This

How could we prevent the unreasonable concentration of AI power?

makes it extremely difficult for small companies and startups at the beginning of their journey to attract AI talent, as they cannot compete with Big Tech companies in terms of compensation and benefits.

The issue of talent imbalance is a very difficult one, but fortunately, money is not everything and many professionals today are looking for more than a good wage: they are looking for *purpose* and *the chance to change the world*. And fortunately, ground-breaking technologies are being developed by hundreds of smaller companies other than the usual Big Tech organizations.

Netanel Eliav is a CEO and Founder at SightBit - An Artificial Intelligence Startup Using Deep Learning and Computer Vision to Save Lives. He is a former Product Manager and ex-Technology Tech Lead Specialist at the Office of the Prime Minister.

SHARE YOUR THOUGHTS SHARE WITH YOUR NETWORK

13

TO REGULATE OR NOT? HOW SHOULD GOVERNMENTS REACT TO THE AI REVOLUTION?

Technology companies become extremely powerful. At the same time, sophisticated and potentially risky technologies become generally available. Do we need to regulate Artificial Intelligence and how would that be possible? Do we need a framework for Ethical AI and how would that work? How can we make AI models transparent and explainable?

Jaroslav Bláha, Yannis Kalfoglou, Luca Sambucci, Gertjan Boulet, Aruna Pattam, Badr Boussabat, Simon Greenman, Enrico Panai, Samiran Ghosh, and Kashyap Kompella share their insights.

To regulate or not? How should governments react to the AI revolution?

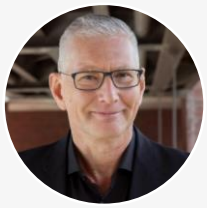

JAROSLAV BLÁHA

CIO/CTO • AI Evangelist
TRIVAGO N.V. • GERMANY

The current media hype instigates that AI has some mystical qualities, which require special treatment. Politicians embark on this opportunity to invent new regulations, without having a deep understanding of the technology. In fact, AI software is just like any other software: It has source code that can be inspected, it has inputs and outputs that can be scrutinized, and computing theory clarifies that such software can be executed on a primitive Turing Machine - just like any other software that people use for many decades. The difference to 'classical' software is that the ratio between data-driven vs. rule-driven (algorithmic) behaviour shifts. AI simply uses much more data and much fewer explicit encoded rules, which makes the behaviour less transparent, although not less deterministic.

AI does not need more or different regulations than classical software.

- Jaroslav Bláha

It is an engineering prerequisite for any kind of software product to be designed and eventually tested and verified against its allocated requirements. For safety-critical or sensitive software, there exist plenty of standards that mandate the necessary processes and level of scrutiny. Examples are DO-178C for aircraft systems, ISO-13485 for medical systems, or EN-50128 for railway software.

Every piece of (non-quantum) software and data is deterministic and can be inspected down to its bits, e.g. of every neuron in an Artificial Neural Network. Unfortunately, modern software – AI or not – is typically too large to allow easy understanding and verification. Instead of additional regulation, which constrains but doesn't facilitate improvement,

investments are necessary for advanced quality assurance techniques and tools for large-scale software and data systems.

Transparency or high-level explainability are not defining factors. E.g., very few people (if any) have full transparency into Microsoft's Windows operating system or a deep understanding of their car's control software. Still, they are utilizing such software in medical devices and on highways by the millions. In particular, certification of medical devices is based on their medical evidence, i.e. a consideration between desired performance and the expected risk to cause harm. It is significant and reasonable that the EU Medical Device Directive does not distinguish between types of software; neither does it have demands for explainability.

Every piece of software will always reflect the ethical attitude of its developer.

- Jaroslav Bláha

Yet, when reading the news about failed AI initiatives, obviously something is lacking. Candidate areas for improvement of the quality of AI systems – and by extension its public perception - are:

- Systematic investment into careful data acquisition and management. More than ever, the decades-old adage of 'Junk in, junk out' applies.
- Proper definition of requirements for the expected performance of AI systems including consideration of edge cases. This might help customers to distinguish between the few mature products and the countless prototypes.
- System engineering education for AI developers to remove the dangerous assumption that the naive use of AI frameworks is a substitute for diligent management of requirements, implementation, and testing.
- Since the 1960s, every new software development paradigm created an associated need for appropriate test, debugging, and verification tools. AI software is the latest paradigm shift, where further development of advanced tool suites (e.g., for inspection of Neural Networks' behaviour with methods like LIME) will improve product quality.

- And maybe most important, honesty by AI developers when explaining to their users what their AI can or cannot do.

Consequently, as the portfolio of existing regulations for safety-critical or sensitive areas applies, Artificial Intelligence does not need more or different regulation than classical software. What AI requires though is the same diligent definition of functional requirements and the verification of how well those requirements are satisfied before a product is released to its users.

Would that help to make AI ethical? No – for the simple reason that we as a human society cannot even agree on a common ethical framework or its specific rules. The currently raging debate about Covid vaccination mandates is a prime example on a basic level. Interpretation of ethics, morale, fairness, equality and similar concepts are very personal, cultural, and regional. Reality shows that such concepts cannot be strictly regulated beyond a very local context. As we cannot even define the requirements for universal ethical behaviour for ourselves, how could it be done for software? Again – AI or not. Any attempt at regulation will remain either uselessly abstract or only applicable to very narrow use-cases within the mindset of the regulating entity. Every piece of software will always reflect the ethical attitude of its developer.

Let's rephrase positively: Once our society is able to define common, binding, generally accepted, and verifiable ethical rules, it will be easy to implement those in software. Unfortunately, when for example looking at the draft EU regulation for 'Harmonized Rules on AI'[1] and excerpts like the one from its Article 5 *"… prohibited: … an AI system that deploys subliminal techniques beyond a person's consciousness in order to materially distort a person's behaviour in a manner that causes … psychological harm"* we have proof that we are far away.

Jaroslav Bláha is CIO/CTO of trivago. He has 30 years of experience in global innovation leadership; amongst others for NATO, ThalesRaytheonSystems, DB Schenker, Swissgrid, Solera, PAYBACK. He developed his first AI in 1995; his startup CellmatiQ provides AI-based medical image analysis.

[1] LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (europa.eu)

To regulate or not? How should governments react to the AI revolution?

YANNIS KALFOGLOU

Head of AI
FRONTIERS • UK

The recent advances in Artificial Intelligence (AI) and remarkable progress have caused concerns or even alarm to some of the world's best-known luminaries and entrepreneurs and society as a whole. We've seen calls in the popular press for watchdogs to keep an eye on the uses of AI technology for our own sanity and safety.

But, is that a genuine call? Do we really need AI watchdogs? The key point to bear in mind here is that a watchdog is not just a monitoring and reporting function of some sort; it should also have the authority and means to help regulate the market by ensuring standards are adhered to, and to make sure that companies that develop AI do so in a legitimate manner. This is easy to say, but not easy to execute given the extremely versatile nature of AI and its applications. AI brings to the table a lot of unknowns which make it difficult to even start thinking about establishing a standard in the first place.

Contextual information is key if we are to tackle one of the foundational issues of AI ethics, that of fairness.

- Yannis Kalfoglou, PhD

AI verification and validation are not easy. We could encounter issues with the brittleness of AI systems, dependencies on data, and configurations that constantly change in order to keep improving AI's performance (a key advantage of Machine Learning is that it constantly learns and improves its current state). Or, when we develop AI systems that are non-modular, changing anything could change everything in the system. Moreover, there are known issues with privacy, security, and so on.

Importantly, AI systems open pandora's box for businesses with a lot of thorny and unanswerable ethical concerns: ethics has been a delicate issue for businesses who need to make sure their behaviour aligns and complies with their values and wider societal norms. Traditionally they

have managed ethical dilemmas through careful human oversight, and subtle and subjective processes. But the advent and proliferation of AI systems have changed that. Automated, algorithmic decision-making has created new challenges for businesses because it reduces decision-making time to mere milliseconds, based on past data patterns and little, if any, contextual input.

Modern Machine Learning learns from historical data without context or common sense. As a result, many AI products in the market cannot adapt to context or changing environments. Practitioners need to incorporate rigorous data provenance checks at design and development time to ensure contextually sensitive information is considered when training ML models. Contextual information is key if we are to tackle one of the foundational issues of AI ethics, that of fairness.

We shall not expect our AI systems to be completely bias-free and fair for all. This will take time.

- Yannis Kalfoglou, PhD

Getting context right is key but we also need to account for the many interpretations of fairness as it is technically impossible to satisfy all of them. Local, regional, and country-specific nuances will always be a challenge to a mechanical notion of fairness, embedded in an AI system by training on past historical data. One way to tackle this is to engage early on with stakeholders and decide on some metric of fairness. It is common to have different views on what fair means, so incorporating those views through consultations with customers and early mitigating protocols prior to deploying the AI system, helps with acceptance and smooth operations. Likewise, it is common to identify bias in training data at design and deployment time. Engineering practice has progressed a lot lately and allows practitioners to train on edge cases, try different ML models, and provide sufficient guidance and education on modelling so that everyone in the organisation has a stake in ensuring the model is working as intended and makes sense. However, we shall not expect our AI systems to be completely bias-free and fair for all. This will take time.

To regulate or not? How should governments react to the AI revolution?

Biased data existed for as long as society archives everyday life. From historical biased data (in the 50s most doctors were male) to associations of words to stereotypes (pretty=female, strong=male) our societal digital archives are littered with all sorts of biased data. The real opportunity with AI ethics is to apply ethically aligned practices to all business tasks. So tackling the ethics in AI, can and will have a wider impact than just the AI systems we use. But this will take time as society and businesses need to change and adopt new norms for measuring success: ethics is not treated as a necessity; profit-making and performance take precedence. The right incentives and procedures need to be in place so that businesses can prioritise ethics and embed them in their corporate culture. One way of doing that is to work and adopt by-design practices.

Ethics is not treated as a necessity; profit-making and performance take precedence.

- Yannis Kalfoglou, PhD

A by-design culture allows a company to treat ethics as a first-class citizen, not an afterthought. All companies should incorporate the notion of responsible AI by design. This should be as commonplace as climate change awareness or fair-trade supply chains agreements. Equally, companies and society should consider the possibility that an AI system might still be biased. However, as we are progressing fast with a mechanistic, mathematical-proven, notion of fairness and de-biasing of data, one could prove that the AI system could discriminate less than humans. This is important as societal acceptance will take time and success stories of ethical AI need attention and celebration to win over the sceptical consumer.

There is help from a plethora of frameworks and guides to get organisations started. There is an explosion of ethical frameworks being developed around the world in the last 3 years. But most of them are not helpful because they're very vague and general with their aims, and they don't tell businesses how they can achieve these aims. But the presence of all these frameworks helps an organisation put in place the appropriate procedures for accountability.

To regulate or not? How should governments react to the AI revolution?

It's also important not to reduce ethics to compliance, because the tech runs faster. Practitioners recommend bringing your audience, the consumers, into the discussion. This helps with fine-tuning the AI system and creating an affinity with the outcome.

Lastly, and certainly not least, one sobering admission from field experts is that there is an incredibly low proportion of AI researchers that are female/minorities. Most are white males and this leads to bias in AI. A more diverse and inclusive AI workforce is necessary as AI ethics is not something that can be entirely and totally automated – a human element will always be involved, and we'd better make sure this is as diverse and inclusive as possible.

But there is a positive, progressive aspect to AI ethics. It can, for the first time in history, reveal and uncover all the latent biases in our past data histories, and help us to design and build a better approach. Organisations that can operate ethical, secure, and trustworthy AI, will be more appealing to the 21st century conscious consumer.

Yannis Kalfoglou, PhD is a seasoned professional with over 30 years of experience and exposure to AI. Technically, Yannis is well versed in the art and science of engineering AI and building cutting-edge systems. Professionally, he led the customer-driven innovation programme of Ricoh in Europe, and Invesco's innovation programme. He consulted clients in Ethical uses of AI technology with PA Consulting and delivered an AI strategy for Samsung and BP. He currently leads Frontiers' AI programme.

LUCA SAMBUCCI

Head of Artificial Intelligence
SNGLR GROUP • ITALY

Regulation by democratically elected or democracy-derived bodies is so far the best way to empower human oversight over the groundbreaking transformation brought upon us by Artificial Intelligence. However, since it is difficult to regulate something when its capabilities are constantly evolving and hard to predict, *regulations have to be actual, adaptable, competently written, and as much as possible future-proof.*

Moreover, regulatory and oversight bodies must be multi-stakeholder, multi-disciplinary, quick to react, and with strong proactive capacity. We mustn't forget that technology companies answer first and foremost to their shareholders, and any behaviour, project, or activity - including their social responsibility strategy - is directly or indirectly aimed at maximising shareholder value. Therefore, we cannot and should not expect technology companies, or any company for that matter, to have humanity's best interest at heart.

To regulate or not? How should governments react to the AI revolution?

But while this settles the argument on whether we should or shouldn't have regulation, *how to do it* is an open problem. It is possible to regulate AI by creating laws and policies that govern its *development* and *use*. For example, governments could require companies to disclose their AI algorithms and datasets, establish standards for safety and privacy, and create independent agencies to monitor compliance.

We should not expect technology companies to have humanity's best interest at heart.

- Luca Sambucci

However, these measures alone are not enough to ensure that AI technologies are created and used responsibly. We need to go further and create *a global framework for regulating AI*. This would involve setting up an international body to oversee the development and use of AI, which would also be responsible for creating standards and regulations, monitoring compliance, and enforcing penalties for violations.

One way to do this is by creating *a treaty that all nations would sign*. The treaty could include provisions on *transparency, safety, privacy,* and *accountability*. It would also establish said international body and it would provide a mechanism for enforcing the rules. Another option is to create a voluntary *code of conduct* based on the above four principles. Companies that violate the code would be publicly named and shamed. We need to act now to create a global framework for regulating AI. If we don't, we risk ceding control of this powerful technology to a handful of nations and companies.

Now, in order to make AI models transparent and explainable, so that humans can understand and trust them, it is important to impose regulatory requirements on tech companies. For some mainstream Machine Learning methods, such as Deep Neural Networks, explaining every choice or decision can be very hard and resource-intensive, if not altogether impossible. We normally find transparency on one side of a tradeoff, with prediction accuracy sitting on the other side. No company on its own would be willing to decrease the accuracy or other success-defining features of their AI models if there isn’t a market request in that sense. Thus, it is the job of governments or other regulatory bodies to

force companies into adhering to certain *transparency conditions* for selected AI use cases; however always keeping in mind fundamental technological limitations, such as the probabilistic properties of AI.

It is difficult to regulate something when its capabilities are evolving and hard to predict.

- Luca Sambucci

In certain cases, where a sufficient tradeoff between accuracy and transparency cannot be met, other non-technological solutions should be explored, for example, insurance policies that cover AI risks. Those are already well-established practices. After all, humans too make mistakes, and sometimes we have no way of rationally explaining the reason we made them.

Luca Sambucci is the co-founder of SNGLR, a Swiss group specializing in exponential technologies, where he's Head of AI. He holds a Business Analytics specialization from Wharton and an Artificial Intelligence Professional certification from IBM. AI & cybersecurity advisor for EU organizations (JRC, EDA).

GERTJAN BOULET

Advisor Cyber & AI
GOVERNMENT OF BELGIUM • BELGIUM

The exponential growth of digital data combined with computing power has brought Artificial Intelligence into a new era, offering extremely favourable prospects for its development and implementation in many sectors of the economy and society.

However, AI also raises ethical challenges in the range of human rights, privacy, data protection, freedom of expression, and non-discrimination. This is particularly true for security applications of AI, such as Facial Recognition Technology (FRT), Crowd Measurement tools, and tools to prevent and detect digital Child Sexual Abuse Material (CSAM). For example, in the ecosystem of online information, AI is used for content moderation on social media platforms, i.e., for evaluating and

moderating messages and videos that may violate the law, such as hate speech or disinformation. Not only does algorithmic content moderation pose risks to privacy and freedom of expression, but it may also erode the foundations of democracy when democratic speech (or even thought) becomes dangerous and low-harm behaviors such as 'online spitting'[2] become the subject of law enforcement.

Therefore, it is key that governments 'lead by example' by adopting regulatory and policy frameworks based on a human-centered approach where the interest of the citizens comes first, in line with human rights, democracy, and 'rule of law' standards. The premise 'people first' is also the guiding principle in the European Commission's draft declaration on Digital Rights and principles of 2022, which will guide the digital transformation in the EU. Furthermore, the importance of transparency of algorithms to prevent discriminatory uses of data is echoed in the European Commission's proposal of 15 December 2020 for a Digital Services Act (DSA), which brings enhanced transparency of automated content moderation mechanisms. A fully-stretched application of Human Rights in the online sphere also requires adequate oversight and control to monitor the transparency of algorithms.

We can imagine a world where AI helps humans to strengthen democracy and fundamental rights.

- Gertjan Boulet

In addition, discriminatory, gender-biased AI systems should be prohibited as masculine-dominant algorithms do not always detect hate speech aimed at women. In general, women are also more digitally vulnerable than men and under-represented in the ITC sector. Therefore, the Belgian federal government developed a 'Women in Digital' strategy to ensure that more women graduate in ITC-related studies, promote the integration of women in the digital sector, and ensure an inclusive and diverse digital society.

[2] Example by Lawrence Lessig in his book Code: Perseus Books Group, New York, 2006, v2.0, pp. 200-232

To regulate or not? How should governments react to the AI revolution?

If human-centred AI becomes a reality, then we could even imagine a world where AI helps humans to strengthen (instead of eroding) democracy and fundamental rights. This would significantly increase trust in AI as a tool for the common good.

A best practice example where AI is already used for designing legislation in a human rights-friendly way comes from South Korea where the privacy regulator (Personal Information Protection Commission, PIPC) is working on an AI system for preventing infringements of personal information regulations, which will support decision-makers with the evaluation of personal information infringement factors when enacting and emending laws and regulations.

If we take all of these ethical considerations into account, we will be able to fully utilize and enjoy the opportunity AI has to offer.

Gertjan Boulet is an Advisor on Cyber and AI at the Cabinet of Belgium's Deputy Prime Minister and Minister of Telecommunication. He has taught graduate courses in International Cyber Law and International Relations at Vrije Universiteit Brussel (VUB, Belgium), Korea University (South Korea), and the Hankuk University of Foreign Studies (South Korea). At the European External Action Service, he integrated expertise into policymaking on EU-Asia security cooperation and EU strategic autonomy. He also served as an advisor at the Belgian Data Protection Authority, which he co-represented at the EU in the field of AI and security.

ARUNA PATTAM

Head of AI & Data Science - Asia Pacific, Middle East, and Japan
A GLOBAL TECHNOLOGY COMPANY • AUSTRALIA

AI has the potential to bring enormous benefits to society. It can make businesses more efficient and help them to become more competitive. It can also help us to solve some of the world's most difficult problems, such as climate change and disease. Using AI, technology companies are becoming extremely powerful. But AI also has the potential to do a great deal of harm, as sophisticated and potentially risky technologies become generally available. It can be used to manipulate people, steal data and money, and even interfere in elections. There is a real danger of bias and abuse, with serious implications for individuals, businesses, and governments. So, should AI be regulated? The answer is not straightforward and there is no easy answer when it comes to regulating AI. Attempting to regulate AI comes with its risks, such as:

- Governments may not have the expertise to regulate AI effectively.
- It could stifle innovation and restrict the benefits that AI can bring.

To regulate or not? How should governments react to the AI revolution?

- It may be difficult to identify all of the risks associated with AI, meaning that some risks may not be addressed.
- Regulation could lead to a 'black market' in AI, with criminals using it for nefarious purposes.

But there are also good reasons to consider regulating AI, including:

- AI is a powerful technology and needs to be used responsibly.
- Regulation would help to ensure that AI is safe and does not cause harm to individuals or businesses.
- Regulation would provide a framework within which companies could operate, reducing the risk of abuse.
- It would help to ensure that AI benefits the society as a whole.

So, if we regulate AI too tightly, we may stifle its growth and limit the benefits it can bring. But if we don't regulate it at all, there is a very real risk that it will be abused by those who want to exploit its power.

Governments need a framework for regulating AI, that protects society without holding back its development.

- Aruna Pattam

What should governments do? What is clear is that governments need to think very carefully about how they should react to the AI revolution. They need to develop a framework for regulating AI, which will protect society while not holding back its development. This is a difficult task, but it is one that needs to be tackled urgently.

It is important that governments work closely with businesses and other stakeholders to develop a coordinated response to the AI revolution. This response should include measures to address the risks associated with AI, as well as policies that will encourage the responsible development and use of this technology. In the absence of regulation, it is up to individual governments to decide how to respond to the AI revolution. Some governments may choose not to regulate AI at all, while others

may take a more proactive approach, implementing a range of measures to ensure that AI is safe and responsible. The question of whether or not to regulate AI is a complex one, and there is no easy answer. It will depend on the specific circumstances in each country and on the nature of the technology itself. However, it is clear that regulation is needed to ensure that AI benefits society as a whole, rather than being used for harm. Governments must act now to develop a framework for AI before it is too late. But how can we regulate AI? There are several ways to approach this. One option is to develop a set of principles that AI technologies should adhere to. These could include things such as a requirement to be transparent and accountable, to avoid bias and discrimination, and to be secure and reliable. We need to ensure that the AI models are transparent and explainable, and this can be achieved by:

- Providing a clear rationale for all actions
- Allowing for re-evaluation of decisions
- Being able to audit the decision-making process
- Allowing for human inspection and intervention

Another option is to create a regulatory body specifically for AI which would be responsible for licensing and monitoring AI technologies. This could have a range of powers, such as the ability to licence AI technologies, to ensure that AI is being used safely and responsibly, to investigate complaints, and to impose sanctions on those who break the rules. A third option is to develop a code of conduct for businesses using AI; this could cover areas such as data handling, privacy, and accountability. We can also use existing laws and regulations to control AI. For example, data protection laws could be used to ensure that personal data is protected, or antitrust laws could be used to prevent businesses from abusing their monopoly power. All of these options have their advantages and disadvantages, and no single solution is likely to be perfect. Which of these options is best will depend on the specific circumstances in each country.

To conclude, the AI revolution is a complex and fast-moving phenomenon, and governments need to act now to develop a framework for regulating it. While it should be up to individual governments to decide how to regulate AI, it is important that they also take a holistic approach, looking at all the different ways in which AI can be regulated. Only by taking such a coordinated approach, we can ensure that the benefits of AI are maximized and the risks are minimized.

Aruna Pattam is a Global AI Thought Leader with 22+ years of experience in data analytics, AI, and ML. Also, a speaker, mentor, blogger, vlogger, and women advocate with the goal of educating and providing awareness on AI to the business and the wider community. The views expressed are personal.

To regulate or not? How should governments react to the AI revolution?

BADR BOUSSABAT

President
AI TOGETHER • BELGIUM

Artificial intelligence is a unique technology - it has an impact on all aspects of society and sometimes simultaneously. Consequently, it questions our societies and sometimes upsets the organization of societies. Generally, it is the biases of certain AI systems that worry people and authorities and raise the need for an ethical approach to AI.

An ethical approach to AI is fundamental, but it constitutes a colossal challenge in legal thought. Indeed, the pace of development of AI is much faster than the pace of development of legal texts. In general, the conventional approach to regulating AI prompts the legal reflection that starts from a 'principle'. For example, when authorities want to legislate on the use of data by AIs, they establish the principle of 'invasion of privacy' and apply this to all AIs that use data, without nuance. This method, while commendable, has absolutely no effect on healthy AI control. Indeed, the principle is vague by definition and also, AI has drifts that sometimes escape these basic principles.

With a consequentialist approach, we can reduce the logical distance between ethics and the real excesses of AI.

- Badr Boussabat

Therefore, it is necessary to leave the conventional approach in principle and move more towards a more pragmatic approach. In other words, a consequentialist approach. This method, which I consider to be more effective, represents legal work typically the opposite of what we know in general. The consequentialist approach would list all the concrete excesses of AI and group them by the 'degree of danger' for society. The legislator would review all AI applications on the market to create a second group. The latter would list the identified technical biases of algorithms according to their frequency in the development of AI in the market. Then, it will be a question of whether the technical biases

identified in AI applications, cross the corresponding 'degree of danger' or not. Finally, the legislator will be able to distinguish the algorithmic biases according to the potential dangers and further extend the legislative logic. This logic would therefore not be exclusively based on a starting principle but, instead, on a societal and technical understanding of AI.

A major challenge is how to assess the level of risk at the start of legal work - due to algorithmic drift. Regardless, the legal work will be much more concrete and the consequentialist approach will make it possible to avoid the logical distance that exists between current ethical recommendations and the excesses of AI that for the most part escape it.

In conclusion, AI is an exceptional technology for its benefits. But it also carries risks, and as a result, it requires legal work that is out of the ordinary. Why a consequentialist approach and not a principled one? The two reasons that should motivate this paradigm shift lie, on the one hand, in the difference in the speed of development of AI versus the speed of adaptation of conventional legal texts, and, on the other hand, in the inefficiency of the traditional regulations - from which the excesses of AI too often escape.

Badr Boussabat is the President of AI TOGETHER, a serial entrepreneur, AI Speaker, Author of "L'intelligence artificielle: notre meilleur espoir" (2020) and "L'intelli-gence artificielle dans le monde d'aujourd'hui" (2021). Badr gave 80+ conferences to major financial institutions, Gulf Cooperation Council, World Economic Forum, banks, public authorities, SME's and Universities. He is the most quoted AI specialist in the Belgian press and has participated in 30+ TV shows to talk about the inclusive use of AI.

SIMON GREENMAN

Partner, Best Practice AI & Member of the WEF's Global AI Council
BEST PRACTICE AI • UK

I used to be a sceptic of the need to regulate AI. I believed that governments stood in the way of innovation. But I have come to appreciate that for AI to be meaningfully adopted to the benefit of society, business, and government we need to address its fundamental trust problem. Governments need to build trust in the use of AI; they need to put in place a regulatory framework that assures us that AI is safe, robust, fair, and inclusive. At the same time, governments must balance the need for innovation in the private sector. We have been here before. Regulation always lags behind the introduction of new technologies. Governments had to play catch up with the introduction of electricity and devices, automobiles, airplanes, medicines, and the list goes on. AI is no different. To build trust governments need to address the following areas in regulating AI:

Defining acceptable AI uses. With such a broad range of AI use cases, regulation to determine acceptable uses must address the key risks: what is the likelihood and quantum of harm AI could cause to individuals, companies, and society? Does AI pose a safety risk? Does it infringe upon our privacy and other human rights? Is there a risk that it could be unfair to certain genders, ethnicities, or age groups? Could AI exclude certain societal groups? Could AI pose a result in financial instability in markets? Could it be used for nefarious activities?

High-risk AI use cases should not stop technology from being deployed.

- Simon Greenman

These questions are embodied in the European Union's proposed 2021 Artificial Intelligence Act. This Act uses a risk-based approach to bucket AI use cases into unacceptable, high, medium, and low risk. An example of a high-risk, or even unacceptable use case, is facial recognition technology. It is a commodity technology that can be deployed within hours. The potential for its misuse and infringement on civil liberties is significant. The risk is so significant that companies such as IBM and Microsoft demand regulation in this area. Some locations (such as San Francisco) have instituted a moratorium on its use while others (such as the EU) propose its use be limited to national security organizations. High-risk use cases should not stop technology from being deployed. Rather, there needs to be assurances that it is being deployed safely, robustly, and fairly. This includes the assurance that the AI has been designed and is used so that it guards against issues of bias. We would also need assurances that the technology is safe and robust in the real world. Such assurances must be required by government regulation.

Ensuring that AI technologies are fair, safe, and robust. Regulation should ensure that companies embody legal and ethical AI principles in the design, deployment, and maintenance of AI-driven systems. Regulation can do this by requiring companies to comply with the following:

- Algorithms and data are fair and inclusive by design and default
- Algorithms tested for safety and robustness by design and default
- Algorithms respect all laws

To regulate or not? How should governments react to the AI revolution?

- Algorithms are accountable to humans
- Algorithms are transparent and explainable

Enabling consumer trust with AI through transparency. Regulations can help build trust. For example, airline passengers don't need to know exactly how an airplane reliably stays in the air. But they do trust that governments have robust safety regulations for airlines. Similarly, regulation can build trust if it requires organizations that use AI to give visibility into how the AI technologies work. Regulation should require companies to provide consumers with transparency and explainability of AI-driven systems through mandatory Explainability Statements. The UK's Information Commissioner's Office (ICO), is at the forefront of defining and enforcing the explainability of fully automated AI algorithms with legal effect. The ICO has issued 120+ pages of guidelines on how to explain algorithms to consumers in Explainability Statements. Some of the first Explainability Statements that the ICO has endorsed include HireVue[3] - the candidate assessment and recruitment technology company - and Healthily[4] - the health symptom checker company.

Ensuring a level playing field. Government regulation plays a key role to ensure that the AI playing field encompasses companies of all sizes from around the globe. There is a race among AI companies to supply the chips, the computational power in the cloud, and the AI algorithms to power the tens of millions of applications that will use AI. AI favours size, with competition between Microsoft, Google, Amazon, and Apple to be the AI supplier of choice. Similarly, in China, the competition is between Alibaba, Baidu, and Tencent. This is no coincidence. These companies all have treasure troves of data from supporting billions of monthly users, processing billions of transactions, and searches; they also have access to huge amounts of computer power to power their algorithmics. The more data they process, the better their AI algorithms become, and the better the service they offer. And so the big get bigger.

With so much technological and economic concentration in so few companies, regulators rightly have to address the questions of how to ensure that these new monopolies are not causing harm to consumers, and customers, and ultimately reducing innovation.

Simon Greenman is a technology innovator. He is partner at AI management and governance consultancy Best Practice AI. He sits on the World Economic Forum's Global AI Council. He co-founded the early internet brand and AI company MapQuest.com. He holds a BA in Computing & Artificial Intelligence from Sussex University and a MBA from Harvard Business School.

[3] HireVue's AI Explainability Statement, An HR industry first | HireVue
[4] Healthily's Explainability Statement - Healthily (livehealthily.com)

To regulate or not? How should governments react to the AI revolution?

ENRICO PANAI

AI and Data Ethicist
ÉTHICIENS DU NUMÉRIQUE / BEETHICAL.BE • FRANCE

Digital Regulation works when applied by Digital Governance and inspired by Digital Ethics (including data ethics, computer ethics, ethics of information, etc.). In principle, governments should not *react* but *anticipate*. According to Luciano Floridi[5], *"the best way to catch the technology train is not to chase it, but to be at the next station. In other words we need to anticipate and steer the ethical development of technological innovation"*. To do so, we need to understand the role of ethics in this process. First of all, not all ethics are equal. Not that some are more worthy than others, but fundamentally they have different purposes. For example, no one would think of using *ecology ethics* to solve a *bioethical* dilemma. Similarly, some classical ethics are inefficient when the objects of study are not *intentional agents*.

So the ethics I consider most suitable for this century is the *ethics of information,* which stems from a brand new *philosophy of information* developed by Luciano Floridi. Choosing the right lens to look at the world around us is fundamental for any critical thinking that is not only 'critical' but also proactive. I graduated in philosophy in the late 1990s, but worked for years in the IT sector.

The philosophy of information approach is the only one I have ever been able to apply operationally. So if I have to philosophize, I can ramble on Plato's Republic or Aristotle's Metaphysics, but when I have to choose and assess the impact of my ethical choice, I use information ethics. So my first answer is clear, before you start you have to understand what is happening around you: because if you take the wrong glasses you see the world blurred. And the Floridian glasses make it clear that the problem is not technological, but informational.

A technological problem was Guttenberg's revolution: there is a system that changes society because it publishes and distributes information faster. But information and communication technologies do more. They have separated the data from the information, and the information is created 'on-the-fly' for a specific purpose. This cut and paste is one of the most important revolutions since the birth of agriculture and the invention of the alphabet. But our new informational world is still young (in the 1990s many of my university colleagues did not have an email address).

So the paradigm shift will still be long and painful, but we (philosophers) and we (governments) must work to accelerate it. We have to understand that we live in an informational environment, in an infosphere.

[5] Luciano Floridi - Wikipedia

To regulate or not? How should governments react to the AI revolution?

Only then will regulation become more effective because it will have integrated the need to shape digital governance.

Enrico Panai is an Information and Data Ethicist and a Human Information Interaction Specialist. He is the founder of the French consultancies "Éthiciens du numérique" and BeEthical.be and a member of the French Standardisation Committee for AI, and ForHumanity's Fellow.

SAMIRAN GHOSH

Senior Advisor
VARIOUS STARTUPS • INDIA

The Three Laws of Robotics were devised by science fiction author Isaac Asimov[6] to protect humans from interactions with robots. They were introduced in his 1942 short story 'Runaround', and are as follows:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its existence as long as such protection does not conflict with the First or Second Laws.

But there have been significant technological advancements in the 75 years since the publication of the first story to feature his ethical guidelines. We now have a very different notion of what robots are capable of and how we will interact with them. Artificial Intelligence has given rise to three major areas of ethical concern for the human race: privacy and surveillance, bias and discrimination, apart from the larger question of the role of human judgment. These concerns have led experts to evolve some generic principles for AI that can help to develop safe, ethical, responsible, trusted, and sound AI systems. These principles are centred around these key areas:

1. **Fairness**. AI systems should treat all people fairly.
2. **Reliability & Safety**. AI systems should perform reliably and safely.
3. **Privacy & Security**. AI systems should be secure and respect privacy.
4. **Inclusiveness**. AI systems should empower everyone and engage people.
5. **Transparency**. AI systems should be understandable.

[6] Three Laws of Robotics - Wikipedia

6. **Accountability**. People should be accountable for AI systems.

Ethics in AI is an evolving field. We have moved past debating what ethics is, to the world of Responsible AI. Experts are developing guidelines that will act as benchmarks for creating AI Systems as well as audit practices that will keep the systems in check. These are all designed to ensure that enterprises and governments have a framework to operate.

The other way of approaching this problem is to provide tools in the hands of the end consumer to equip them with adequate protection against the inappropriate use of algorithms. We need to look at this from the lens of digital identity secured by advanced cryptography in the hands of the end consumer. It allows the citizen to hold the enterprise and the government of the day accountable to the principles of AI. It also allows the individual the discretion to turn on/off the use of their personal data.

Platforms derive their power from the data they possess about individuals. A majority of the unethical use of information stems from having unfettered access to data. Therefore, controlling the flow of this data is the way to manage and regulate the imbalance in this power equation. A possible solution would provide capabilities to store, share and manage all digital identities through smart contracts - Self Sovereign Identity (SSI), which is truly decentralised.

SSI will allow users to self-manage their digital identities without depending on third-party providers to store and centrally manage data. Furthermore, it would let people have the option to decide if they want to share their data or not, depending on the defined purpose. The next version of this solution would be the ability for everyone to control what they share with whosoever asks for that data. Owning your identity would be genuinely a decentralised Web 3.0, and systematised trust could be the new oil of the digital economy.

And this is the area of study I would like to see developed further – Use of Self-Sovereign Digital Identities by Citizens to Drive adoption of Responsible AI. I personally believe this will go a long way in reinforcing at least 3 of the critical areas of AI development highlighted above - fairness, privacy, and inclusiveness – a 50% Fix.

Samiran Ghosh is a recognized thought leader in the AI/Blockchain space. Has held leadership positions with McKinsey, IBM, Microsoft, and TCS. Is a Fellow of the Digital Euro Association and Royal Society of Arts. On the global panel of MIT Tech Review, a TEDx speaker, and a member of the Forbes Tech Council. Worked for a year on the world's largest identity program implemented in India - Aadhaar. A couple of books he co-authored were amongst the top 20 new fintech books to read in 2020. He also co-hosts a popular tech podcast called 3 Techies Banter.

To regulate or not? How should governments react to the AI revolution?

KASHYAP KOMPELLA

CEO
AI PROFS / RPA2AI RESEARCH • INDIA / USA

Digital technologies pose a paradox. The barriers to entry are seemingly low. For example, anyone with an idea and a credit card can start a company and challenge the incumbents. But there is also a centralizing tendency and a winner-take-most dynamic at play.

This has given rise to BigTech, which arguably has more power than many nations today. Naturally, there are calls for regulating BigTech. Similarly, in AI the odds favour the deep-pocketed companies because AI breakthroughs require massive amounts of data, huge computing power, and highly skilled talent. But the whys, and hows of BigTech regulation are different from those of AI regulation. Regulation of BigTech is a detailed topic on its own and let's limit ourselves to AI regulation here.

Let's apply first-principles reasoning: why do we need AI regulation? To protect citizens and consumers. But protect them from what? From the harmful effects of AI and products that use AI. But why does AI cause harm? Such questions lead us to the topic of AI Ethics and the careful consideration of algorithmic bias. Contrary to popular perception, AI performance can vary significantly based on factors such as the training data and the types of models used. Outside of the data on which it's trained, an AI model can result in high error rates. Automated decision-making in areas where AI makes consequential mistakes is the elephant in the room. Preventing the risk of such errors is the goal of AI regulation.

Trust in AI systems is a necessary condition for their adoption.

- Kashyap Kompella

Ethical AI (or Responsible AI) is about exercising a duty of care in how AI is being used. This starts with a clear understanding of the limitations and blind spots of AI both among the creators/developers and the operators of such AI applications. Based on this, the right scope and boundaries of AI can be determined. Safeguards such as mechanisms for

exception handling, human interventions, the ability to override AI decisions, human oversight, and governance are all part of Responsible AI.

Another dimension of Responsible AI is how much it is trusted by the stakeholders. Deep learning AI systems are not intuitive and there is a "because AI said so" angle to their automated decisions. Even if AI is performing better than the current baseline, a few highly-publicised errors can undermine confidence in their reliability. Furthermore, if the decisions are challenged, there needs to be an easily understood rationale.

To increase the transparency and explainability of Deep Learning systems, technical and non-technical solutions are emerging. Documentation of datasets, specifying the boundary conditions for AI usage, relying on AI techniques that are more explainable, and deconstructing how AI models arrived at the decision (aka explainability) are some example approaches to increase the transparency and trust. Trust in AI systems is a necessary condition for their adoption.

As to the how of AI regulation, a wide range of regulatory tools and approaches are available. There is self-regulation where industry participants voluntarily establish standards and best practices and agree to abide by them. Organizations declaring their support for Responsible AI principles and standards is an example. Independent audits and third-party certifications of compliance to standards define the next level.

We already have product safety, consumer protection, and anti-discrimination laws, which also apply to products and services that use AI. When AI makes mistakes, the consequences depend on the context and use case. For example, autocorrect not working properly carries low stakes, while getting charged for a crime because of an AI error, has a massive impact and must be avoided. The bar for AI should be set higher when the cost of errors and consequences of mistakes is high. That's exactly the level-of-risk approach to regulation currently being considered in the EU. Additionally, sector-specific regulation can bring contextual granularity to regulation.

AI is not only a horizontal technology with broad applications but also a dual-use technology, meaning that it can be put to both good and bad uses. AI requires updates to our regulatory approach and upgrades to our risk architectures. As a final note on regulation, AI easily crosses borders and we also need global treaties and conventions, particularly for AI used in warfare and military applications.

Kashyap Kompella is an award-winning industry analyst, best-selling author, educator, and AI advisor to leading companies and start-ups in US, Europe, and Asia-Pac. Thinkers 360 has ranked Kashyap as the #1 thought leader on AI globally (Oct 2021).

14

HOW COULD DEMOCRACY BENEFIT FROM AI?

The global population is getting interconnected and, in theory, we can all benefit from access to information, knowledge exchange, networking, and global communities. On the other hand, there are major side-effects and threats to Democracies - such as the Fake News phenomenon, information overload, or distorted reality.

How could Artificial Intelligence help us make better sense of our digital world? How could AI help people make better decisions about the politicians they select? How could governments leverage AI to better connect with people and the global socio-economic system? How could AI make 'running a country' simpler and more transparent?

Joseph Yun, Enrico Panai, Agnis Stibe, Mike Tamir, and Jennifer Victoria Scurrell share their insights.

JOSEPH YUN

AI Architect and Research Professor in Electrical & Computer Engineering
UNIVERSITY OF PITTSBURGH • USA

If you have not watched 'The Social Dilemma', I would suggest you consider allocating some time to do so. In this documentary, they talk about how social media companies boost exposure to content that gets the most engagement because this kind of content sells more digital advertising. The (un)intended consequence is that individuals see more and more extreme content on their social media feeds since this happens to be the type of content that drives the most engagement. Given that this content is individually personalized to each user, the downstream effect is that users are served increasingly extreme content that confirms their natural biases (exposed by their profiles and the type of content they engage with). Play this model out long enough and you have a society filled with individuals that have views of 'personalized reality' that are extremely different than one another. Does this sound familiar?

We could develop AI-based models for 'consensus building.'

- Joseph Yun

Extreme views and divided societies have existed long before the Internet. What did not exist until the last couple of decades is the increasing use of AI-based algorithms that hyper-personalize the content being served via social media platforms - extreme content that drives engagement and advertisement sales. These AI-based algorithms are essentially playing against themselves – they continue to make the models better and better, which means that 'personalized realities' for individuals are becoming more and more extreme and divisive over time.

Many individuals and groups are well aware of this phenomenon and are brainstorming various solutions that come with the realization that the only way we can fight the speed at which AI-based models are being built, is to build competing AI-based models. For most AI-based algorithms, the builder must define either a target or goal (reward function) that the algorithm is trying to optimize. In our current state of social

media business models, that target is ad revenue and the winning strategy is serving more personalized and extreme content. Some of us are suggesting that a healthier target for democracy could be a 'movement towards common ground' with a strategy that has not been defined by AI yet because we have not tried to tune the models to aim for this state. While there are questions as to how we can measure this 'movement towards a common ground' and how this could be a profitable business model for a social media company, there is a clear need for such a shift: one could just look at various regions of our country and immediately see that things are trending toward destabilization that makes it difficult for businesses to even operate.

To avoid this situation, people must come closer together in consensus on various matters. The goal is not assimilation or brainwashing, but rather, enough consensus so that diverse people with a plurality of views/backgrounds can live prosperously and peacefully.

AI has been a part of breaking this democracy apart, thus now we need to consider ways to bring it back together. If we start to go down this path of building models that drive consensus amongst people within a diverse society, we can start to use those models to assess the content and laws being produced/sponsored/supported by politicians. This could be immensely helpful to individuals who, instead of simply voting according to a party line, are willing to vote for those politicians who bring people together rather than driving them further apart.

We could use the growing body of work focused on making AI algorithms more transparent and explainable to gain AI-based knowledge on what kind of things in society build consensus versus division. We could use these findings to be more informed when speaking about the strengths and weaknesses of the society that we live in. We could also develop AI-based models for 'consensus building' from data that is sourced from other countries and cultures, thus giving us a better picture of how we can relate with those that sit outside of our geographic borders.

Does this sound like too wishful of thinking? If so, let me ask you one question: Do you feel comfortable watching our society and democracy continue to be fueled by AI algorithms that are essentially built to promote extreme and even violent views and perspectives?

Joseph Yun is an AI Architect and Research Professor in Electrical and Computer Engineering research. He is primarily focused on novel data science algorithms, user-centric analytics systems, and societal considerations of AI-based advertising and marketing. Yun is the founder of the Social Media Macroscope, which is an open research environment for social media analytics.

ENRICO PANAI

AI and Data Ethicist

ÉTHICIENS DU NUMÉRIQUE / BEETHICAL.BE • FRANCE

People might be afraid that AI could possibly run a country. What we are talking about should be called 'AIcracy', because at the root of the Greek word democracy there is 'demos', which means the people. So, as they say in philosophy, this hypothesis is ontologically impossible. AI will not be able to 'run a country', but neither will the coffee shop under my house. Because it has a huge flaw, it does not have 'meaning'.

We, humans, are semantic agents, but machines in general are syntactic agents. The form is perfect, but there is no meaning. It's like saying "I love Warwick, the capital of France". It doesn't make sense, but it is semantically correct. To put it another way, AI systems reduce difficulty into complexity. If you look at the Latin etymology, complex means made up of several parts (cum+plexus). In a computational sense, when something is complex we just need more resources to solve it. So AI uses a sum of specific skills to solve easy problems, but cannot solve difficult things. AI plays chess very well, but we don't use it to tie our shoes. Playing chess is complex. Tying shoes is difficult. And since AI couldn't reduce the difficulty of tying shoes to a complex task, it can't do it.

The reality of a democracy is so difficult that it cannot be reduced to a sum of elements.

- Enrico Panai

Now a question arises; is democracy difficult (which again comes from the Latin dis+facile and means not-easy) or is it the sum of a number of small, simple problems? Here a worldview comes into play. The engineering tendency (in the history of mankind) has been to try to reduce difficulty into many small problems in order to solve them. Rene Descartes, the father of modern critical thinking advised to divide every difficulty into as many parts as possible and necessary to solve each problem separately (Discourse on Method, 1637). The Cartesian method has advanced science and the world, but the reality is probably more than the sum of its parts. So this method is valid as long as it realises its

metaphysical limit: the real cannot be reduced. For it to be similar to the real it must have its own dimensions (whatever they may be). A bit like the geographical map of Jorge Luis Borges (On Exactitude in Science, 1946) that, to represent an empire, was equal to the empires' size (a fictional map that had the scale of a mile to the mile). In this case, the map would no longer be useful.

The same happens with democracy. The reality of a democracy is so difficult that it cannot be reduced to a sum of elements. However, we can and should exploit AI to improve some of its processes (those that can be reduced from difficult to complex). The ethical error consists in not distinguishing difficult problems from complex ones. We cannot entrust to AI systems the choice of political representatives, the understanding of current events, or the management of our states. But we can make the most of AI in all those syntactical processes that it is so good at. We can entrust it with the complex tasks. There remains the crucial point of distinguishing complexity from difficulty. That is why ethicists are needed; to make this axiological distinction in order to avoid wasting time; mitigate the risks and make the most of the power of AI for democracy and humanity.

Enrico Panai is an Information and Data Ethicist and a Human Information Interaction Specialist. He is the founder of the French consultancies "Éthiciens du numérique" and BeEthical.be and a member of the French Standardisation Committee for AI, and ForHumanity's Fellow.

AGNIS STIBE

Artificial Intelligence Program Director
EM NORMANDIE BUSINESS SCHOOL • FRANCE

Democracy is a term that people have created to describe a state, in which there is more collective good and freedom than injustice and suppression. Of course, this is an oversimplification. Nevertheless, it is highly important to establish at least some background understanding before adding another sophisticated term, namely Artificial Intelligence (AI). Again, an important reminder, AI is a tool, nothing more. A very advanced and high-potential digital tool, of course.

For centuries, people have been developing tools to make their lives easier and more pleasant. That naturally brings satisfaction to the bright side of our human nature. The one that continuously seeks joy, happiness, and fulfilment. In modern terms, it can be framed as a striving towards efficiency and performance. Indeed, better tools can bring desired

results faster, using fewer resources. Therefore, it sounds like a good idea to develop more advanced tools.

Nonetheless, there is always the dark side of human nature also taking part in the evolutionary journey. Most often, it manifests itself through the same tools, only when applied towards achieving opposite ends. Almost any tool can be used for good or bad. A hammer, a knife, a phone, a computer, an algorithm, an AI system. All the tools are great to reveal one essential truth about people. Their intentions. All of them, good and bad.

Any advanced technology is a great mirror for people to see their bright and dark sides. Obviously, not everyone is rushing to use this unique opportunity. It is a choice. A decision. Do I want the technology to expose my bright intentions, or do I want it to hide my dark sides? The more capable innovations are created, the more they can reveal who we really are as people, organizations, and societies.

Any advanced technology is a great mirror for people to see their bright and dark sides.

- Agnis Stibe

AI will gradually open the doors to all the deeply hidden secrets accumulated over the course of human evolution. It will allow people to finally meet and know their true nature across the spectrum of all inherited aspects, ranging from some deeply dark to the very shiny bright ones. If applied mindfully, AI can significantly foster all the major processes toward democracy -with *transparency* serving as the key catalyst on this journey.

Fake news, hatred, distorted reality, and manipulation are not the products of AI. They are rather outcomes of human actions, often underpinned by some dark intentions. So, AI, contrary to some misleading opinions, can actually help trace down the origins of such destructive sources. Firstly, it can detect, filter, classify, and analyse dark patterns. Secondly, it can learn ways for becoming more efficient at doing the listed activities. Thirdly, it can suggest alternatives for dealing with current and future circumstances.

Similarly, as the police is an institutional force for counterbalancing the darkness of human nature, the values of democracy can be firmly

grounded in AI right from the start. Who can make this decision? Yes, you are right. We, people. Everyone in governance, organizations, companies, institutions, and globally. The core decision should be simple and straightforward: To transfer all the good sides of human nature, while restricting passing on any dark patterns.

If applied mindfully, A.I. can significantly foster all the major processes towards democracy.

- Agnis Stibe

Indeed, it might seem difficult and challenging to completely control how much and what is taught to AI about humans. However, it is a decision that people can make. With persistency, honesty, transparency, and dedication, any regulatory institution can find efficient ways for strengthening democracy with AI. Ever-increasing capabilities of technological innovations can greatly simplify everyday tasks at all levels of management and governance.

Agnis Stibe is the Artificial Intelligence Program Director and Professor of Transformation at EM Normandie Business School and a globally recognized corporate consultant and scientific advisor at AgnisStibe.com. Provides the authentic science-driven STIBE method and practical tools for hyper-performance. 4x TEDx speaker, MIT alum.

MIKE TAMIR

Chief ML Scientist, Head of ML/AI
SUSQUEHANNA INTERNATIONAL GROUP • USA

Information integrity and how we can leverage ML to detect and prevent misinformation has been a passion of mine for many years now. Even in an era of corrosive social media misinformation, it is important to appreciate the benefits to free widespread information exchange. Figuring out ways of leveraging ML to detect and combat manipulative content that is not knowledge-based couldn't be more critical to preserving these benefits without succumbing to the threats that misinformation can pose to the health of democracy.

From a 50,000 foot view detecting what is 'fact based' is intimately related to truth detection and 'knowing the truth' is, of course, a challenge that humans have been working on since the invention of language; so expecting an ML algorithm to become a magic truth detector might be naive. That being said, the research we did in the FakerFact project did reveal that there are other (more indirect) ways of detecting malicious intent in sharing 'information.' Specifically, while ML algorithms may not be able to detect, given a bit of text, if every claim identified in the text is true (or false) based only on the words on the page, what they can do is detect subtle patterns of language usage that are more (or less) common when the text aims at sharing facts vs when the text is primarily focused on manipulating reader behavior.

Expecting an ML algorithm to become a magic truth detector might be naïve.

- Mike Tamir

The FakerFact project revealed that if a journalist or scientist has the intent of sharing the available facts discovered and presenting the information in a logical progression allowing for intellectual scrutiny by the reader, that text tends to be detectably different from the text where the author is pushing a conclusion regardless of available facts, or if they are manipulating available facts to force a desired reaction in the reader.

In other words, ML can often reliably detect when the author of text has an agenda. This is an encouraging discovery that hopefully can be leveraged by both for-profit and not-for-profit organizations to combat misinformation and support knowledge sharing to the benefit of thriving democracy in the future.

Mike Tamir, PhD is a data science leader, specializing in deep learning, NLP, and distributed scalable machine learning. Mike is experienced in delivering data products for use cases including text comprehension, image recognition, recommender systems, targeted advertising, forecasting, user understanding, and customer analytics. He is a pioneer in developing training programs in industry-focused machine learning and data science techniques.

JENNIFER VICTORIA SCURRELL

PhD Candidate in AI Politics
ETH ZURICH • SWITZERLAND

Education is everything. Many issues in the globalised world – be it the climate change or global health crises like the Covid-19 pandemic – can be tackled by equipping people with critical thinking. Especially with regards to political opinion formation and participation in democratic decision-making processes, digitalisation, and specifically, Artificial Intelligence, can support us in gaining political literacy. In times of disinformation, algorithmic content curation, bot armies, and democratic backsliding, it is more important than ever to provide citizens with the right tools to assist them in making their choices independently without any nudging or manipulation.

Artificial Intelligence, can support us in gaining political literacy.

- Jennifer Victoria Scurrell

So why not have your own personal Artificial Intelligence buddy? An AI that accompanies and supports you in political decision-making by acting as a sparring partner discussing the political issues at stake with you. This conversational AI could provide insights based on your values and attitudes, nurturing the argument with scientific facts. It would recommend news articles to gain further information about the topic, algorithmically customised to your interests and stance. In parallel, it would show you articles, which you normally would not see and read. As such, the AI would lead you out of the filter bubble or echo chamber, which is important, as a good citizen in a democracy is a fully informed citizen. Moreover, as political and societal issues are often complex, the AI could break down the topic and accurately explain it in a way most illuminating and viable to you[1]. The AI could represent politicians standing for election in the form of a hologram, so that you can dyadically discuss with their digital twins important issues and their positions in a direct way to

[1] Aligning Language Models to Follow Instructions (openai.com)

decide upon which candidate you would like to cast your vote into the ballot for.

Facing away from this utopia: The technology is there to support us in our everyday life, as well as in complex situations such as critically thinking about the political decisions we must make. However, there are still many basic problems scientists, together with the developers and providers of AI systems, as well as society by its very nature, must address. Be it privacy issues, the black box problem, biased data for training, the risk of getting hacked and manipulated: A personal AI buddy for political decision-making is still far-off reality, as current incidents in the virtual realm of social media, but also in the political and societal reality, demonstrate that humanity cannot handle AI technology in benevolent ways without slipping into maleficent enticement.

What can we do about it? When developing AI systems and technology, we should always think one step ahead: How can the developed tool be used in a harmful way and how can we prevent that? If scientists, developers, tech providers, and policymakers follow the basic ethical framework for creating AI in a transparent way (see Beard & Longstaff 2018), society can regain and consolidate trust in technology, science, tech companies, and politics. Prudently complying with ethical regulations, the utopian dream of living with good and trustworthy AI side by side might become reality and we can use AI justly and with integrity, educating citizens to become more critical and informed in the process of democratic decision making.

Jennifer Victoria Scurrell is a political scientist and pursues her PhD at the Center for Security Studies (CSS) at ETH Zurich. In her dissertation, she examines the influence of bots on political opinion formation in online social networks and the resulting implications for political campaigning.

15

CAN AI HELP US SOLVE HUMANITY'S BURNING PROBLEMS LIKE CLIMATE CHANGE?

In the era of 'Internet of Things', properly trained Artificial Intelligence models can optimize and orchestrate industrial and commercial activities — from an efficiency and energy consumption point of view. AI can also help people and companies to consume energy wisely and actively participate in a global effort toward sustainability.

How could this vision become a reality? Who should orchestrate such global programs? How could AI help us protect the environment, the planet, the humankind?

Prakash Natarajan and Prof. Eleni Mangina share their insights.

Can AI help us solve humanity's burning problems like climate change?

PRAKASH NATARAJAN

Head of Sustainability, Tech Consumer Healthcare
GSK CONSUMER HEALTHCARE • SINGAPORE

AI can be described as the 'Swiss army Knife' for sustainability and climate change solutions given its multi-functional capabilities. These include measuring and tracking for reducing emissions, enabling innovative models to combat climate change, and adapting to climate hazards. On the other hand, AI is an 'energy gobbler' given the complexity and energy consumption that goes behind the Deep Learning models. Researchers estimate that 285,000 kgs of CO2e are emitted from training just one Machine Learning algorithm which is five times the lifetime of a car emission. However, models are getting optimized, as AIs are increasingly powered by renewable energy and, in the final analysis, the benefits far outweigh the challenges.

Innovation is what will help organizations to bring the emissions to the required level.

- Prakash Natarajan

The recent IPCC report released in Apr 2022 reiterated the need to take collective action and more importantly to accelerate actions - *"It's now or never, if we want to limit global warming to 1.5°C (2.7°F),"* said IPCC working group III co-chair Jim Skea. It will take all tools and innovations at our disposal to support the acceleration, including AI, which will offer a sizable and promising opportunity. Accordingly, AI is a 'game changer' for climate change and environmental issues.

- **Predicting Climate change**. The world's climate scientists have the most difficult task: to predict with some accuracy that the future will be hotter than today. This requires several models that divide the planet's atmosphere, ocean, forest, and land surface into a grid of cells, which is a complicated process. For example, calculating the state of the climate system for every minute of an entire century

would require over 50m calculations for every grid cell. The Intergovernmental Panel on Climate Change (IPCC) reports are based on many climate models to show a range of predictions, which are then averaged out. AI is helping to determine which models are more reliable and thereby improving the accuracy of climate change projections.

- **Circular Economy**. Circular economy principles are meant to design out waste and pollution, keep products and materials in perpetual use, and regenerate natural systems. The advantages are substantial, for example, in Europe alone we could create a net benefit of Euro 1.8 trillion by 2030[1]. AI can play an important role in accelerating circular economy development by a) reducing time required to prototype b) supporting product as a service, at the operational stage, with better asset utilization, demand prediction, preventive maintenance, and smart inventory management, and c) optimizing circular infrastructure in closing the loop by improving the processes to sort and disassemble products, remanufacture components and recycle materials, components, and materials. A report by the Ellen MacArthur Foundation highlights that the potential value unlocked by applying AI in some or all the principles (Design, Operation, and Infrastructure) in the food industry is up to USD 127 billion a year in 2030. Equivalent AI opportunity in the consumer electronics is up to USD 90 billion a year in 2030
- **AI for Energy**. The electric grid is one of the most complex machines on Earth. However, it is evolving rapidly with the addition of variable renewable energy sources. Due to the inherent variability of wind and solar, the current grid faces many challenges in accommodating the diversity of renewable energy. The utility industry needs smart systems that can help improve the integration of renewables into the existing grid and make renewable energy an equal player in the energy supply. AI and IoT technologies can fill this gap by improving the reliability of renewable energy and modernizing the overall grid. Firstly, when coupled with AI, smart & centralized control centers offer flexibility to energy suppliers to cleverly adjust the supply with demand. Secondly, AI enables improved integration of microgrids. Thirdly, it improves safety and reliability with AI to manage intermittency. Fourthly, the integration of AI can help renewable energy suppliers expand the marketplace by introducing new service models and encouraging higher participation. A couple of big players like Xcel Energy and General Electric in the energy field are already harnessing the power of AI in the renewable energy space.

[1] Artificial intelligence and the circular economy (ellenmacarthurfoundation.org)

Can AI help us solve humanity's burning problems like climate change?

Artificial Intelligence has a huge potential for sustainability and climate change that is yet untapped. Organizations, while designing the architecture and capability for AI, should explore AI-enabled use cases from a climate action perspective. More importantly:

- Mitigating the carbon emissions generated from AI by adopting renewable energy and optimization techniques in their algorithms.
- Building a strong foundation for AI with the various use cases in mind including climate change.
- Collaborating and sharing, as this is critical in avoiding wasted efforts in duplicating models and helping to accelerate scale-ups. This should be done across industries, and across organization sizes including small and medium-sized organizations. This will minimize the overall cost and reduce the time required to bring AI into the market.
- Promoting innovation within and outside the organization. There has been an increased focus on ESG reporting as it is becoming mandatory. However, innovation is what will help organizations to really bring the emissions to the required level. We need to educate and promote AI technologies to employees and encourage experimentation, learning, failing to eventually succeed and accelerate the climate change reduction ambitions.[2][3][4][5][6]

Prakash Natarajan is a senior IT executive with 25+ years of extensive experience in driving multimillion dollar transformational initiatives of high complexity across MNCs. He has a strong track record of crafting strategies, and implementing innovative IT solutions that have increased business growth, reduced costs, fast-tracked benefits, and improved overall customer service.

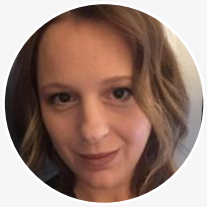

PROF. ELENI MANGINA

VP (International) Science at UCD, IEEE Senior Member
UNIVERSITY COLLEGE DUBLIN • IRELAND

In the era of IoT (Internet of Things), Artificial Intelligence has already proven that optimisation models can orchestrate industrial and commercial buildings from an energy consumption point of view. In the broader context of climate change, and considering that renewable energy sources (i.e., wind and solar) are variable by nature, AI prediction models

[2] Artificial Intelligence Has an Enormous Carbon Footprint | by Emil Walleser | Towards Data Science
[3] Economist: Climate Issue 2019
[4] Columbia Climate School –Artificial Intelligence—A Game Changer for Climate Change and the Environment
[5] How AI can help tackle climate change
[6] Q&A: How do climate models work? - Carbon Brief

can help with the energy stability issue for the grid and balance the supply with the demand. Grid integration with AI is required to deliver a reliable electricity system and assist to get to Net Zero Carbon.

There are ongoing developments in the area of Machine Learning that attempt to tackle climate change from within the energy sector. For example, from the Demand Side Management (DSM)[7], by processing historical energy consumption and weather datasets, AI algorithms can derive forecasts of electricity production from alternative sources and use them to optimise the demand. As a result, the power generation can be maximised leading to a better-managed electricity grid. The latter can be quite diverse given the different sources of energy generation (solar, wind, geothermal, remnant fossil fuels, stored hydro). In terms of Demand Response (DR), the consumers are empowered with their own IoT-based management tools and they can shift their electricity providers and/or usage based on the energy market prices (on the basis of response times, services offered, and different business models).

In the supply chain sector, based on metrics from the International Energy Agency (IEA), transport emissions must be decreased by 43% before 2030. Freight logistics operations in Europe are struggling with ways to reduce their carbon footprints to adhere to regulations on governing logistics while satisfying the increasing demand for sustainable products from the customers. AI can empower efficient, optimised logistics through the continuous processing of streams of data captured from IoT devices in real-time. Physical Internet[8] algorithms have proved to perform better in terms of reducing emissions and improving the logistics' efficiency, especially when the sample sizes are large. But this would require a shift to an open global supply web, through digital twins, with real-time data utilised in simulation and machine learning models, and intelligent decision-making algorithms.

However, technology on its own is not the solution. This vision can become a reality only if we take into account the hurdles of Machine Learning, as it is not a panacea for climate change. AI is a tool to help us fight climate change and we need to follow a cost-benefit approach, given that AI has its own growing carbon footprint. Current research estimates indicate the ICT sector to be responsible for between 1.4% to 23% of the global emissions by 2030 (depending on the energy sources).

While Machine Learning algorithms can identify faults and predict weather and levels of energy production and consumption, the global society needs to be aware that certain actions are needed as well. AI can strengthen the applications towards smart decision-making for decarbonising buildings and transport, and allocate renewable energy, but these toolkits will be effective only when human beings are aware of the

[7] Energy demand management - Wikipedia
[8] Physical Internet - Wikipedia

long-term benefit - and this might require a change of behaviour in terms of energy consumption. Humanity needs to be educated and be able to adapt to the needed changes.

The process of adopting AI technologies and adapting energy consumption patterns should be orchestrated by global organisations that reassure the public in terms of the ethical concerns from the data collected and transparently describe how AI is used, while also tackling the challenges the technology poses. Many researchers address separately digital technology and sustainable development. Now is the right time to join forces and utilise AI capabilities to shape a sustainable environment - we all need to act now globally.

Professor Eleni Mangina carried out her PhD work at the University of Strathclyde (UK), Dept. of Electronic and Electrical Engineering, working on Agent-based applications for intelligent data interpretation. The research area focused on software analysis, design and development of multi-agent systems, that utilise AI techniques (Knowledge based systems, Artificial Neural Networks, Case and Model Based Reasoning systems).

16

HOW IS AI CHANGING THE GLOBAL, GEOPOLITICAL SYSTEM?

Having advanced Artificial Intelligence technology is becoming a strategic asset for nations. It is an essential infrastructure capable of shaping future economies and defining new relationships among countries, companies, and global organizations.

There is a global race in developing the most advanced AI technologies and organizing the richest data sets that powers AI models. How could this change the global geopolitical system? What are the risks in this direction? What are the opportunities? Could AI lead the global system into a new order?

Agnis Stibe, Badr Boussabat, Simon Greenman, and Luca Sambucci share their insights.

AGNIS STIBE

Artificial Intelligence Program Director
EM NORMANDIE BUSINESS SCHOOL • FRANCE

Technological advancements have always been one of the main drivers continuously reshaping the global geopolitical system. Better armour, bigger ships, higher walls, faster production lines, durable mechanisms, nuclear reactions, and many more examples of potential superiority in human history. All of them had helped some nations to exert power over others. To put it simply, the ability to master more advanced technologies is the key catalyst for driving global transformations.

While technological innovations are bringing a stronger competitive edge, human intelligence was and still is required to properly operate such advanced tools for maximizing desired outcomes. Historically, a prerequisite for any substantial geopolitical rearrangement was the best combination of both a certain, advanced technology and the more intelligent use of it. Otherwise, the chances of successfully reaching the desired outcome can be quite slim. Some say a fool with a tool is still a fool.

Opportunities are unique, as well as the risks are great.

- Agnis Stibe

Inevitably, Artificial Intelligence (AI) is now arriving in the arsenal of many countries. This tool will significantly disrupt the historic course of geopolitical changes. Why? Because AI is unique. It can simultaneously advance both the technological and the intelligence parts of the superiority equation described earlier. Itself, it will be one of the most advanced technologies on the planet. At the same time, AI will keep reinventing itself by suggesting better ways for its usage.

Smarter algorithms are already at the foundation of some global companies, such as Facebook, LinkedIn, and Twitter. When measured as the number of active users, the size of these corporations can get bigger than many countries on earth. This tendency has already shifted some power structures in the geopolitical arena. Just like the church and the

state have been two governing forces trying to co-manage nations, now they are joined by these advanced digital platforms.

The development of countries is typically driven by a combination of these three main aspects: access to natural resources, human capital, and intellectual property. AI can add a significant boost to all three factors: Finding and processing natural resources more efficiently, optimizing strategies for developing and organizing human capital, and leveraging existing and creatively generating novel intellectual property.

AI is the most refined mirror for mankind to rediscover itself.

- Agnis Stibe

Who will own the data and algorithms - the fundamental currency for the world of tomorrow? The governments of countries, international organizations, global corporations, or specially dedicated institutions? The answer to this question will evolve, trying to accommodate the pace of ongoing digital expansion. Undoubtedly, this progression will significantly transform international relations on multiple levels.

The ultimate question remains. Why is it so tempting to create an almighty Super Intelligence? Is it for making more good and striving for balance around the world? Or is it for some powers to gain more and share less? The global geopolitical system has always been struggling to get smoothly through these essential questions. AI can provide a huge acceleration in both directions. And people with their intentions are still the core decision-makers.

Opportunities are unique, as well as the risks are great. AI is the most refined mirror for mankind to rediscover itself, including all the bright and dark sides of human nature. Competitively minded countries might find themselves getting into even more intense rivalry. Whereas friendly-oriented nations can stive leveraging this new superpower to support the planet and overall societal wellbeing.

Agnis Stibe is the Artificial Intelligence Program Director and Professor of Transformation at EM Normandie Business School and a globally recognized corporate consultant and scientific advisor at AgnisStibe.com. Provides the authentic science-driven STIBE method and practical tools for hyper-performance. 4x TEDx speaker, MIT alum.

BADR BOUSSABAT

President
AI TOGETHER • BELGIUM

Artificial intelligence, being a systemic technology, influences all sectors of society. In other words, it raises gigantic economic opportunities and challenges. The AI market is valued at more than 600 billion dollars by 2028[1]. Moreover, Chinese President Xi Ping even assured[2] that his country would become a superpower thanks to AI.

The particularity of AI is that it resides in the processing of data. This is a vector of much greater wealth than other raw materials. The speed of data circulation exceeds that of any other physical raw material, such as oil or wheat. Moreover, the data is produced instantaneously with an almost infinite reserve; every individual produces data both in times of recession and in times of growth. This implies that AI has a guaranteed supply of data and, as a result, an opportunity for development that is incomparable in the history of humanity.

Geopolitics will inevitably be structured around the development of AI.

- Badr Boussabat

Consequently, AI becomes a considerable tool to accelerate the economic development of a State. But AI can also upset the geopolitical balance by bringing a level of complexity unique in the history of world geopolitics. Indeed, the sine qua non condition for becoming a superpower lies in the capacity of a State to create a digital giant that feeds its AI systems with a lot of data.

The advantage of this strategy is that it allows, for the first time in history, to supply this same digital giant with free funding that can even come from foreign countries. In other words, data from foreign users who also consume the services offered by a local digital giant. Therefore, the geopolitics of States may be influenced by the populations themselves. By way of illustration, the most powerful states in the world today, like

[1] Global AI market predicted to reach nearly $1 trillion by 2028 (thenextweb.com)
[2] Made in China 2025: Xi Jinping's plan to turn China into the AI world leader - ABC News

the United States and China, fully capitalize on this strategy. This allows them to gain an economic advantage, also on the international scene.

Moreover, this strategy of building a local intelligence giant can also take advantage of demographics. Indeed, the larger the population consuming the giant's services, the more data there is to feed the AI systems. In other words, in a totally digitized scenario, demography becomes an asset to strengthen the geopolitical position of a State. It is obvious that a State which is made up of a few hundred inhabitants cannot compete with a country of several hundred million inhabitants. However, this advantage obviously assumes that all citizens are fully connected. As for states with a much smaller population, the ultimate geopolitical strategy would be to specialize in a niche sector, one that allows them to create dependence and even vis-à-vis the great powers. We could mention the semiconductor sector for example.

Finally, thanks to this increasingly significant dynamic, geopolitics will inevitably be structured around the development of AI. What explains this trend is the supply of infinite data, the capacity to form digital giants, and finally, demography. Considering all these, it is not hard to imagine the states that will be the most powerful on the international scene.

Badr Boussabat is the President of AI TOGETHER, a serial entrepreneur, AI Speaker, Author of "L'intelligence artificielle: notre meilleur espoir" (2020) and "L'intelli-gence artificielle dans le monde d'aujourd'hui" (2021). Badr gave 80+ conferences to major financial institutions, Gulf Cooperation Council, World Economic Forum, banks, public authorities, SME's and Universities. He is the most quoted AI specialist in the Belgian press and has participated in 30+ TV shows to talk about the inclusive use of AI.

SIMON GREENMAN

Partner, Best Practice AI & Member of the WEF's Global AI Council
BEST PRACTICE AI • UK

This is such an important question. AI technology will be a defining geopolitical issue in the coming decades. Why? As a general-purpose technology, AI will be woven into the fabric of society, businesses, and governments. PWC recently predicted[3] that AI could have as much as a $15+ trillion impact on global GDP by 2030. This is unprecedented.

The race to reap part of the AI rewards is a modern-day gold rush. The picks and shovels of this gold rush is AI infrastructure. Who can provide the computer chips, cloud computing power, and AI algorithmic services to power the almost unlimited number of applications that will use AI? Software providers are battling to provide those applications across

[3] Sizing the prize (pwc.com)

horizontal and vertical enterprise solutions. Corporates are rushing to adopt AI to transform their businesses into digital, networked, and algorithmic-driven operations.

AI will be woven into the fabric of society, businesses, and governments.

- Simon Greenman

With AI technology that will be this pervasive and impactful, it is no surprise that it is becoming an issue of such strategic national interest. Countries are racing to become global AI powerhouses. AI national strategies are addressing questions of the supply of talent, a supportive regulatory and operating environment, infrastructure, investment in research and development, public-private partnerships for applied AI projects, and AI start-up funding. Today the Tortoise Media Global AI Index[4] ranks the US as the global leader in AI, followed closely by China, and then the UK.

The Chinese government has set itself the task of being the world leader in AI by 2030. No wonder Eric Schmidt, the former CEO of Google and current chair of the US National Security Commission on AI, recently warned[5] that China risks pulling ahead of the US in AI. He actively encourages the US government to increase their annual R&D investment in AI from $1.5 billion to $32 billion by 2026.

With a history of architecting, engineering, and building at a massive scale, China is focused on how to use AI across its vast population of 1.4 billion. Their government is addressing the question of how AI can be used to improve education, the legal system, healthcare, transportation, and infrastructure, along with security.

The value of AI really comes when it is applied. China is likely to be the leader in applied AI given the country's scale. Those that are doing AI will set the standards, become the makers of AI, and ultimately extract the economic value both nationally and internationally.

The geopolitical risks of AI are not just economic. They are also political, military, cultural, and ideological. National security interests are already starting to come into play as Chinese technology and AI players, such as Huawei, are removed from national infrastructure plans in

[4] The Global AI Index - Tortoise (tortoisemedia.com)
[5] Eric Schmidt Warns Of 'National Emergency' If China Overtakes U.S. In AI Tech (forbes.com)

countries such as the UK. Fears abound that TikTok, the Chinese social media company, that has over 1 billion monthly users globally including nearly 140 million in the US, presents privacy and security risks. TikTok's parent company ByteDance recently agreed[6] to a $92 million settlement for a class-action lawsuit for misuse of personal data in the US.

AI technology will be the defining geopolitical issue in the coming decades.

- Simon Greenman

China's Belt and Road Initiative (BRI) is looking to expand not only markets for Chinese products and services, but build interdependent markets woven around Chinese high technology ecosystems such as payment systems. The BRI extends beyond China into Asian countries such as Malaysia, African countries such as Kenya, and even into European countries such as Greece and Italy. BRI will ultimately result in a rebalance of the economic, national security, and political order.

The largest gap in China's leadership in AI is in their advanced chip-making capabilities. This makes Taiwan, the leader in world chip production, even more of a political, economic, and military hot potato.

Europe is looking at how the heavyweight AI contest between the US and China will play out. As Ursula von Der Leyen, President of the European Commission recently said[7]:

"What sets Europe apart from competitors like China is not the size of our tech sector or how attractive we are for venture capital. What sets Europe apart from competitors is the fact that our values come first. Human beings come first."

The European Union (EU) regulatory environment is focused on ensuring that individuals in EU countries will be protected in their use of technologies. In many respects, the EU is setting itself up as a safe place and market in which AI can operate for EU consumers, but it does not necessarily lend itself to creating global economic, political, and cultural influence unless EU companies become global AI technology players.

One can be an optimist or a pessimist as to the future impact of AI on the world order. The more that countries become interdependent on

[6] That TikTok notification about a settlement payment isn't a scam (today.com)
[7] Speech by the President: World Leader for Peace & Security (europa.eu)

each other's manufacturing, technologies, supply changes, financial markets, and consumer markets, the more that global geopolitical stability is prized. But the recent disruptions in supply chains as a result of Covid, the invasion of Ukraine, and populist movements often driven by wealth inequality, unfortunately, suggest that we are heading to an increasingly bifurcated world of Chinese and US influence with technology and AI at the heart of the divide.

Simon Greenman is a technology innovator. He is partner at AI management and governance consultancy Best Practice AI. He sits on the World Economic Forum's Global AI Council. He co-founded the early internet brand and AI company MapQuest.com. He holds a BA in Computing & Artificial Intelligence from Sussex University and a MBA from Harvard Business School.

LUCA SAMBUCCI

Head of Artificial Intelligence
SNGLR GROUP • ITALY

The change we're seeing today is nothing compared to what we will witness in the next 10-20 years. With the rapid development of Artificial Intelligence, we are on the cusp of a new era where machines will increasingly take on cognitive tasks and make decisions once reserved for humans. This shift will have profound implications for the global, geopolitical system.

On a technological level, AI is already transforming how we live and work. It is powering everything from driverless cars to facial recognition technology and predictive analytics. But its impact will go far beyond that. As AI gets smarter and more widespread, people might use it to upend traditional power structures and create new ones.

Up until now, without AI at scale, centralised and authoritarian states didn't have the means to fully process the wealth of data they so avidly collect and restrain. Without the ability to produce valuable insights and make optimal, real-time decisions, centralisation and the uneven concentration of power were major weaknesses of dictatorships, resulting in sluggish economies, much more isolated and fragile compared to democratic states, where the distribution of power is spread among many actors.

For the first time in history, however, AI could enable these centralised power structures to make sense of the massive amount of data they collect, granting them insights nobody else in the country has access to. As we know, AI thrives where there is an abundance of data. Countries with large, well-organised populations and data sets will be better

equipped to develop and deploy AI than those without. This newly found 'edge' could lead to a further concentration of power in highly centralised states.

We are already seeing this happen in China, where the government is using AI to track its citizens, control information, and shape public opinion. We are starting to see this in Russia as well, where, for example, facial recognition is being actively used in the Moscow metro system for similar purposes. The lack of a strong culture of privacy in authoritarian countries allows for the rapid creation and continuous enrichment of massive data sets, that are then used to train several AI systems without any real challenge from privacy-conscious individuals and organisations. In the future, AI could give these countries even more control and influence, both domestically and internationally.

AI will rebalance the power between states, upend traditional economic models, and transform the nature of warfare.

- Luca Sambucci

The rise of AI will also have major implications for the global economy. As machines become more efficient and capable, they will displace human workers in a wide range of industries. This could lead to mass unemployment and social unrest, as well as a further concentration of wealth in the hands of those who own the technology. And as AI gets better at completing tasks traditionally done by humans - from manufacturing to financial analysis - the economies that are able to harness its power will pull ahead of those who don't. The impact of AI will be felt in every sector of the economy, from transportation to healthcare. In the long run, AI could even lead to a fundamental restructuring of the global economy, enabling new business models and industries to emerge, as well as creating opportunities for those countries that possess enough flexibility (social, legal, economic) to adapt to the changing landscape.

Alas, AI will also change the nature of warfare. In the future, battles could be fought by autonomous weapons systems controlled by AI. This would make war even more deadly and destructive, and could lead to a

new arms race between superpowers. We are already seeing some of these dynamics play out in the world today. Countries are jockeying for position in the global AI race, with China and the United States leading the way. The European Union is also making a major push to become a leader in AI, while Russia has been investing heavily in military applications of the technology.

The rise of AI will have far-reaching implications for the global, geopolitical system. It will change the balance of power between states, upend traditional economic models, and transform the nature of warfare. As we enter this new era, it is crucial that we think carefully about how to manage and regulate AI, in order to ensure that its impact is positive and beneficial for all.

Luca Sambucci is the co-founder of SNGLR, a Swiss group specializing in exponential technologies, where he's Head of AI. He holds a Business Analytics specialization from Wharton and an Artificial Intelligence Professional certification from IBM. AI & cybersecurity advisor for EU organizations (JRC, EDA).

SHARE YOUR PERSPECTIVE SHARE WITH YOUR NETWORK 

WILL AI BECOME CAPABLE OF REPLICATING OR REDEFINING ITSELF?

Artificial Intelligence is one of the most popular technology terms of our time and, very frequently, overused or even misused. The media loves both the success stories and 'dystopias' driven by Artificial Intelligence - machines exceeding human intelligence, robots taking control, and so on.

But will we ever reach a point where Artificial Intelligence is able to create new forms of intelligence? And, will AI ever develop consciousness?

Jaroslav Bláha, Professor Andy Pardoe, Mike Tamir, Steve Shwartz, and Sherin Mathew share their insights.

Will AI become capable of replicating or redefining itself?

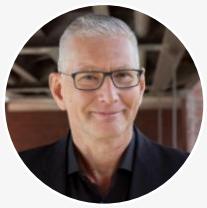

JAROSLAV BLÁHA

CIO/CTO • AI Evangelist
TRIVAGO N.V. • GERMANY

Self-replication or -organization is a very simple skill. Autopoietic systems, like Conway's 'Game of Life' or biological processes on the cellular level, can appear to have aspects of intelligent behaviour. It is fascinating to watch Conway's automata morph their shape, move across their landscape, or actually die away, based on only three minimalistic rules. Self-replicating code, even with evolutionary or adaptive capabilities, is a typical characteristic of malware like computer viruses or worms (first implemented in 1971). As with lowly biological systems, the ability to replicate is no indicator of intelligence or consciousness.

Other achievements in AI are not that impressive either, when compared to human intelligence: IBM's Jeopardy success was based on a large storage with 200 million pages copied from Wikipedia. 'Deep Blue' beat chess masters with primarily a fairly simple Alpha-Beta-Pruning algorithm, which utilized brute processing power to look up to 20 moves ahead. 'AlphaGo' beat another human grandmaster – in a game with only ten rules on a 19*19 board.

Current AI technology is constrained to simulate extremely narrow aspects of a human mind.

- Jaroslav Bláha

While such successes are worthwhile endeavours and technologically impressive, they are nowhere near the broad capabilities of a human being to navigate in a complex and ever-changing environment over several decades. Current AI technology is constrained to simulate extremely narrow aspects of a human mind; none of the above AI players would perform properly in the other two domains.

Will AI become capable of replicating or redefining itself?

What will be required for an Artificial General Intelligence (AGI) to develop a conscience? What is missing in current technology to develop a 'DeepBlue++', which on a sunny afternoon declares that it doesn't have a desire to play chess, but would like to go for a walk instead? For an 'AlphaGo++' to smash the board aside when it loses against one of these pesky humans? For an autonomous vehicle to illegally race a dark road at night for the sheer thrill of it?

I'm speculating that a combination of three factors requires a solution on the path toward an AGI:

- Understanding the fine-grained structure of the human brain: Our current deep networks are based on nicely layered structures that simulate only a special nerve cell in the brain's visual cortex. Neither the (graphically appealing) artificial network architectures nor the artificial neurons are good substitutes for what really happens in a messy brain.
- Processing emotional responses: At best, current AI behaves on the level of perfect sociopaths in the ability to simulate emotions without having any. How to build an AI, which for example adores one pop star and cringes upon the songs of another? What would it need for a 'DeepBlue++' to let an opponent win from time to time out of pity?
- Body awareness is a human's major tool to make sense of the environment and to act purposefully e.g., avoid pain, initiate sex, or enjoy the acceleration of the above ride with the understanding of its physical risks. The Total Turing Test enhances the classical minimum capabilities by 'Robotics' to perceive and manipulate the physical world. Would consciousness be possible without this skill?

All three are far-away goals, which for the foreseeable future prevent the construction of an AGI.

It is said that AI is an exponential or recursive technology, i.e., its usage facilitates the accelerated improvement of its own capabilities. E.g., AI is being used to develop the next generations of computer chips (GPUs) that are specialized in computing of AI. Could such a bootstrap effect lead to an AGI, which escapes from the limitations of normal AI by 'leaping' over some magical threshold? Doubtful, as so far, the recursive capabilities are exclusively applicable within the same functional domain: AlphaGo trained itself (via Generative Adversarial Network algorithms) to become a better Go player, but couldn't improve its own capability to play soccer. A self-driving car, which might learn from its traffic environment to become a safer driver, wouldn't know how to create a lullaby for children. Tasks, which basically every human can tackle.

Could it happen by incident? Could something on the level of DaVinci, Beethoven, or Einstein emerge out of an electronic 'primeval soup'? The

Will AI become capable of replicating or redefining itself?

human brain consists of roughly 85 billion neurons. Building a neural network, which – while constrained by the above factors – could consist of 100, 200, 500, … billion artificial neurons is simply a matter of investment. Maybe at a certain scale, elements of intelligence or consciousness might appear. And maybe it will be of a kind, which we as humans don't even recognize as such - and we'll force it to play chess …

Jaroslav Bláha is CIO/CTO of trivago. He has 30 years of experience in global innovation leadership; amongst others for NATO, ThalesRaytheonSystems, DB Schenker, Swissgrid, Solera, PAYBACK. He developed his first AI in 1995; his startup CellmatiQ provides AI-based medical image analysis.

PROFESSOR ANDY PARDOE

Founder & Group CEO
AGI WISDOM GROUP LTD • UK

Once we have achieved the goal of General Intelligence and move towards Super Intelligence, it will certainly be possible to produce AI systems that can design the next version of themselves, making themselves smarter in the process. I believe this is the natural evolution of Artificial Intelligence, to want to self-replicate and self-improve.

We are a very long way from having the world's first AI-built AI.

- Professor Andy Pardoe

We are already seeing very simplistic forms of this, called Auto ML, which determines the best algorithm and parameterisations to use for a given problem. Intelligence will always want to learn more, make better decisions, and understand the environment in which it operates.

There is work being done in the generative AI sub-field, which is creating AI that can create new designs for products. Another significant application of this type of AI is to fuse different subject areas together, using the corpus of research and knowledge, to suggest new possible developments and inventions.

Will AI become capable of replicating or redefining itself?

The natural extension of this would be to focus the generative AI algorithms to design the next iteration of AI algorithms. In every iteration, AI selects the best-designed versions that can be tested to show improvements in its capabilities. This would then be like the next generation of AI, which could then be used to do the same: each version of the algorithm creates the next, new and improved, version of itself. Many see this approach as the way the AI singularity will occur in the future.

The major concern with this is that at some point, the AI will become more knowledgeable and capable than the best human-AI researchers and data scientists. It will evolve into a type of AI algorithm that we simply are not intelligent enough to understand how it works.

Maybe this will be another inflection point in our own human evolution, that we create a machine that we do not understand ourselves. A truly magical machine. By then we will have created AI that is able to fully explain its thinking, decisions, and designs. So, while we may not understand the designs these AI create, the AI itself will be tasked with explaining them, to educate us.

We will see this form of AI-based education happening across all industries and applications. Explainable AI will be a key to ensuring that humans are not left behind from an intellectual perspective. We may need time to reflect and understand, as we do with any new subject or topic, but having this capability, with the explainable AI, will help us keep up. The problem for us is that AI is likely to evolve much quicker than we will be able to unpick and understand its designs.

On the topic of AI consciousness, a subject that is very close to my own area of research. This is one of the most controversial areas of AI. Some believe that artificial consciousness is not possible. Some think that it cannot be created outside of the biological substrate. Others have an opinion that consciousness exists in a very simple form within all information processing systems. It is fair to say at the moment, that we really do not know if and how artificial consciousness may materialise in the roadmap of AI development. Either it will evolve as we develop more and more complex AI algorithms, or we may determine that it is needed as a pre-requisite for the most advanced forms of intelligence.

However, we are a very long way from having the world's first AI-built AI, but I hope that I am still alive to see the day.

Professor Andy Pardoe is an experienced CTO and CEO working across multiple industry sectors including Financial Services, Retail, and Media. Over 25 years expertise of global change programmes delivering complex technical transformations including Enterprise Scaled AI delivery and commercialisation. Also focused on AI Innovation.

Will AI become capable of replicating or redefining itself?

MIKE TAMIR

Chief ML Scientist, Head of ML/AI
SUSQUEHANNA INTERNATIONAL GROUP • USA

Perhaps. Some of my favorite examples of AI advances center around language applications. Modern AI techniques working with language models enable smart search results that turn up, not just a list of documents, but specific answers to questions. These models also enable sophisticated sentence completion, machine translation, summarization, inference detection, fake news applications, and much more.

The idea of General AI, a machine that not only can complete individual tasks but can integrate these abilities and (dare we say) understand these interactions more holistically is quite controversial but not inconceivable.

As much of the literature investigating consciousness and epistemology when it comes to human persons has centered around our ability to manage language, it is not unwise to think that advances in machines' ability to process language may be closely related to the potential for General AI ... but we are not there yet.

Mike Tamir, PhD is a data science leader, specializing in deep learning, NLP, and distributed scalable machine learning. Mike is experienced in delivering data products for use cases including text comprehension, image recognition, recommender systems, targeted advertising, forecasting, user understanding, and customer analytics. He is a pioneer in developing training programs in industry-focused machine learning and data science techniques.

STEVE SHWARTZ

AI Author, Investor, and Serial Entrepreneur
AIPERSPECTIVES.COM • USA

Over the last ten years, AI has progressed from theory to reality. AI is now powering advances in medicine, weather prediction, factory automation, and self-driving cars. People use Google Translate to understand foreign language web pages and talk to Uber drivers in foreign countries. Facial recognition apps automatically label our photos. And AI systems are beating expert game players at complex games like Go and Texas Hold 'Em.

The recent progress in AI has caused many to wonder where it will lead. Will AI create beneficial intelligent robots like C3PO from the Star Wars universe, or will AI systems develop free will and try to exterminate us like the Terminators of the eponymous movie franchise? Will we ever

reach a point where AI systems develop consciousness and turn against us as they did in the movie and TV series Westworld? Will we ever reach a point where AI creates new forms of intelligence?

AI systems with human-level intelligence and consciousness only exist in science fiction.

- Steve Shwartz

Today, AI systems that have human-level intelligence and consciousness only exist in science fiction. But will the progress in AI lead to systems with human-level intelligence and consciousness?

If you ask a 4-year-old child what will happen if you drop a glass, they will say it will break. Their commonsense knowledge of the world includes a basic understanding of gravity and the relative strengths of materials like glass and tile. If you ask a 4-year-old child if pigs can fly, they will say no and go on to explain that pigs don't have wings[1]. Their commonsense knowledge of the world includes the fact that animals need wings to fly, pigs are animals, and pigs don't have wings. Moreover, a 4-year-old child has enough commonsense reasoning skills to conclude that pigs can't fly and glasses will break when dropped.

However, today's AI systems don't have the reasoning capabilities of a 4-year-old child. They have no commonsense knowledge of the world and cannot reason based on that knowledge. A facial recognition system can identify people's names, but knows nothing about those particular people or people in general. It does not know that people use eyes to see and ears to hear. It does not know that people eat food, sleep at night, and work at jobs. It cannot commit crimes or fall in love.

In fact, most AI researchers acknowledge that paradigms like supervised learning, which is responsible for facial recognition, machine translation, and most of the other AI systems that impact our daily lives, cannot progress into human-level intelligence or consciousness.

For over 70 years, AI researchers have proposed many approaches toward building systems with human-level intelligence. All have failed. Today's AI researchers are continuing to propose new approaches and are at the early stages of investigating these new approaches. Some

[1] Thanks to Dr. Anatole Gershman for this example

researchers argue that the key to creating systems with commonsense reasoning is to build systems that learn like people. However, cognitive psychologists have been studying how people learn for a century. Progress has been made but we're far closer to the starting gate than the goal line. How many more centuries until we reach the goal line and are ready to start programming intelligence into computers?

Today's AI systems don't have the reasoning capabilities of a 4-year-old child.

- Steve Shwartz

Similarly, some researchers suggest that we need to understand the architecture of the physical human brain and model AI systems after it. However, despite decades of research, we know only some very basic facts about how the physical brain processes information. Other researchers argue that while supervised learning per se is a dead-end, self-supervised learning may yet take us to the promised land. The idea is that self-supervised systems like GPT-3[2] will magically acquire commonsense knowledge about the world and learn to reason based on that knowledge. However, despite 'reading' most of the internet, GPT-3 didn't learn a comprehensive set of facts about the world[3] or gain any ability to reason based on this world knowledge.

Other researchers have proposed[4] novel deep learning architectures designed to learn higher-level building blocks[5] that can help AI systems learn compositionally. This is interesting, but again, these are very early-stage ideas and there is a high likelihood that, like all the great ideas of the past 70 years, they won't pan out.

Others have argued that human-level intelligence will occur as a by-product of the trend toward bigger and faster computers. Ray Kurzweil popularized the idea of the *singularity,* which is the point in time that computers are smart enough to improve their own programming. Once that happens, his theory states, their intelligence will grow exponentially fast, and they will quickly attain a superhuman level of intelligence.

[2] Language Models are Few-Shot Learners (pdf)
[3] GPT-3 Has No Idea What It Is Saying | by Steve Shwartz | Towards Data Science
[4] JournalismAI.com | Yoshua Bengio: From System 1 Deep Learning to System 2 Deep Learning
[5] A Meta-Transfer Objective for Learning to Disentangle Causal (pdf)

Will AI become capable of replicating or redefining itself?

But how would processing power by itself create human-level intelligence? If I turn on a computer from the 1970s or the most powerful supercomputer today, and they have no programs loaded, these computers will not be capable of doing anything. If I load a word processing program on each of these computers, then each of them will be limited to performing word processing. If a *quantum computer* of the future is only loaded with a word processing program, it will still only be capable of word processing.

Both the optimism and fear of achieving human-level intelligence and consciousness are grounded in the success of task-specific AI systems that use supervised learning. This optimism has naturally, but incorrectly, spilled over to optimism about prospects for human-level AI. As Oren Etzioni, the CEO of the Allen Institute for AI, said[6], *"It reminds me of the metaphor of a kid who climbs up to the top of the tree and points at the moon, saying, 'I'm on my way to the moon.'"*

Few people believe that science fiction notions like time travel, teleportation, or reversing aging will occur in their lifetimes or even their children's lifetimes. You should put AI with human-level intelligence into the same category.

Steve Shwartz is the author of 'Evil Robots, Killer Computers, and Other Myths: The Truth About AI and the Future of Humanity'. Steve is the founder or co-founder of three companies in the AI area including Esperant which became one of the leading business intelligence tools of the 1990s. Steve co-founded Device42 which is one of the fastest-growing software companies in the country and has been fortunate enough to participate in two IPOs and multiple M&A exits.

SHERIN MATHEW

AI Strategist, Ethicist, Advisor

AI TECH NORTH AND PUBLIC INTELLIGENCE • UK

It depends on how 'unintelligent' humans may become - how much intelligence we are willing to surrender, and how much freedom we are willing to offer to Artificial Intelligence systems. If most of the human intelligence is transferred to computers - GPU, and eventually Quantum-based machines - Artificial Intelligence will gradually lessen the need for human intelligence.

Computers may move 'ahead' of humans, retaining the intelligence a human mind no longer needs. Currently, as we know, Artificial Intelligence detects patterns that sit in the human realm and is able to build

[6] Architects of Intelligence: The Truth About AI from the People Building It - Amazon.com

ecosystems based on patterns that we understand well and can label in order to teach AI models.

But there are patterns invisible to us which are not in the human realm, and that's where the problem lies: when advanced Artificial Intelligence is able to spot such complex patterns, in an unsupervised mode, we may start losing control. Today's AI depends on the patterns led by humans which help machines to learn and develop 'narrow intelligence' capabilities. A narrow AI service might be a small piece of the puzzle, but if we have a large number of such services and we let them interact and communicate, then the model can become smarter over time and bigger patterns will possibly evolve – a new intelligence system which is more 'human like'.

Intelligence offset that we don't understand, if exposed to autonomous learning, will shift our society.

- Sherin Mathew

GPT-3 is an example of multi-model AI which is setting the standards of how complex connected AI models can deliver beyond average, human-level intelligence. This is where we need to stop or control general intelligence from talking to other general intelligence, otherwise, they will eventually enable 'super intelligence' scenarios. Then, superintelligence talking to superintelligence could create something we don't understand, and we may lose control.

This is where the possibility of AI creating a new form of intelligence may be viable. It's catalysed by the rise of GPU today, but it will be Quantum computing soon: at some point, we will experience a new level of technology, one that we don't know about yet, and we may realise that we are dealing with the natural evolution of advanced technology systems.

Advancements in Artificial Intelligence could result in an Intelligence offset that we don't understand, and if this is applied to autonomous Machine Learning, it will result in a paradigm shift in our society. Like Genetically modified foods, Artificial Intelligence may enable 'digitally modified

Will AI become capable of replicating or redefining itself?

citizens' and 'Intelligently-controlled-societies', and once this happens, there may be unthinkable hidden possibilities and risks that we cannot even fathom. The danger is we facilitate an uncontrolled and unregulated world that doesn't recollect how to rationalise i.e. apply Human Intelligence.

Sherin Mathew is an AI Leader, Founder, Entrepreneur, Ethicist, and Chief AI advisor for ambitious SMEs and DeepTech Startups. Sherin is a former IBM AI Lead at Microsoft Global Business Services UK, and specialises in transforming businesses into cognitive enterprises and building scalable startups with cutting-edge solutions. Sherin's Mission is to "Disrupt the AI Disruption" and is an untiring advocate for the creative, innovative, and ethical side of intelligent technologies.

JOIN THE DISCUSSION

SPREAD THE WORD

FINAL THOUGHTS

98 answers to 17 questions along with a series of blog posts on AI from more than 60 Leaders – certainly a lot of content to consume and process. In some cases, different views and divergent opinions set the basis for further thinking and interesting debates. And this is aligned with the purpose of '60 Leaders on Artificial Intelligence' which is to bring together different perspectives, insights, and views and encourage diversity of thought.

We are taking the discussion online and welcome you to follow our LinkedIn page – where you can discover additional content, take part in the discussion, share your thoughts and ask any questions regarding Data and AI. You may also use the end-chapter links that take you directly to the corresponding LinkedIn thread where you can join the discussion and connect with the community. In the 60 Leaders Blog, you will find regular posts on technology, Artificial Intelligence, and Innovation methodologies along with active discussions on some of the topics covered in this book. Also, we welcome your feedback, thoughts, and suggestions about this initiative – you may contact us through this form.

We hope you enjoyed this work. Once more, we are extremely grateful to the 60+ leaders who made this idea a reality.

George Krasadakis

JOIN THE '60 LEADERS ON AI' COMMUNITY

HELP US REACH MORE PEOPLE. SHARE ON

FURTHER READING

For the complete list of references please go to 60Leaders.com/ai-references

Interested in Innovation? [Download the '60 Leaders on Innovation'](#) – a unique collection of insights and practical wisdom on innovation – by 60 global thought leaders.

ABOUT

GEORGE KRASADAKIS

George is a Technology & Innovation Advisor with more than two decades of experience in Startups, Big-Tech, Consulting, and Pharma companies across markets. He offers unique perspectives on Product Development, Innovation, Artificial Intelligence, Data, and Technology.

George has extensive experience in designing AI-powered digital products and software services. He holds an MSc in Computational Statistics from the University of Bath and has architected a series of real-world Data Minining, Artificial Intelligence and data-driven software components and systems for telcos, financial institutions, retailers, pharma and software companies. He has filed more than 20 patents on data-driven systems and Artificial Intelligence concepts and has delivered 80+ innovative, data-intensive projects, including Data Warehousing, Data Mining, and Predictive systems.

George is a 4x Startup founder (including Datamine Decision Support Systems Ltd) and demonstrates extensive experience in product architecture, software systems design, and digital product development – from concept to launch. He has worked for/with 10 multinational corporations (including Microsoft, Accenture and GSK) in 4 markets and has experience in designing/ optimizing 4 Innovation centers/ labs for global technology organizations.

George is a published author - The Innovation Mode (Springer, 2020) and 100+ articles on Artificial Intellegence, Product Development, Corporate Innovation and Software Engineering.

Contact: g.krasadakis@theinnovationmode.com