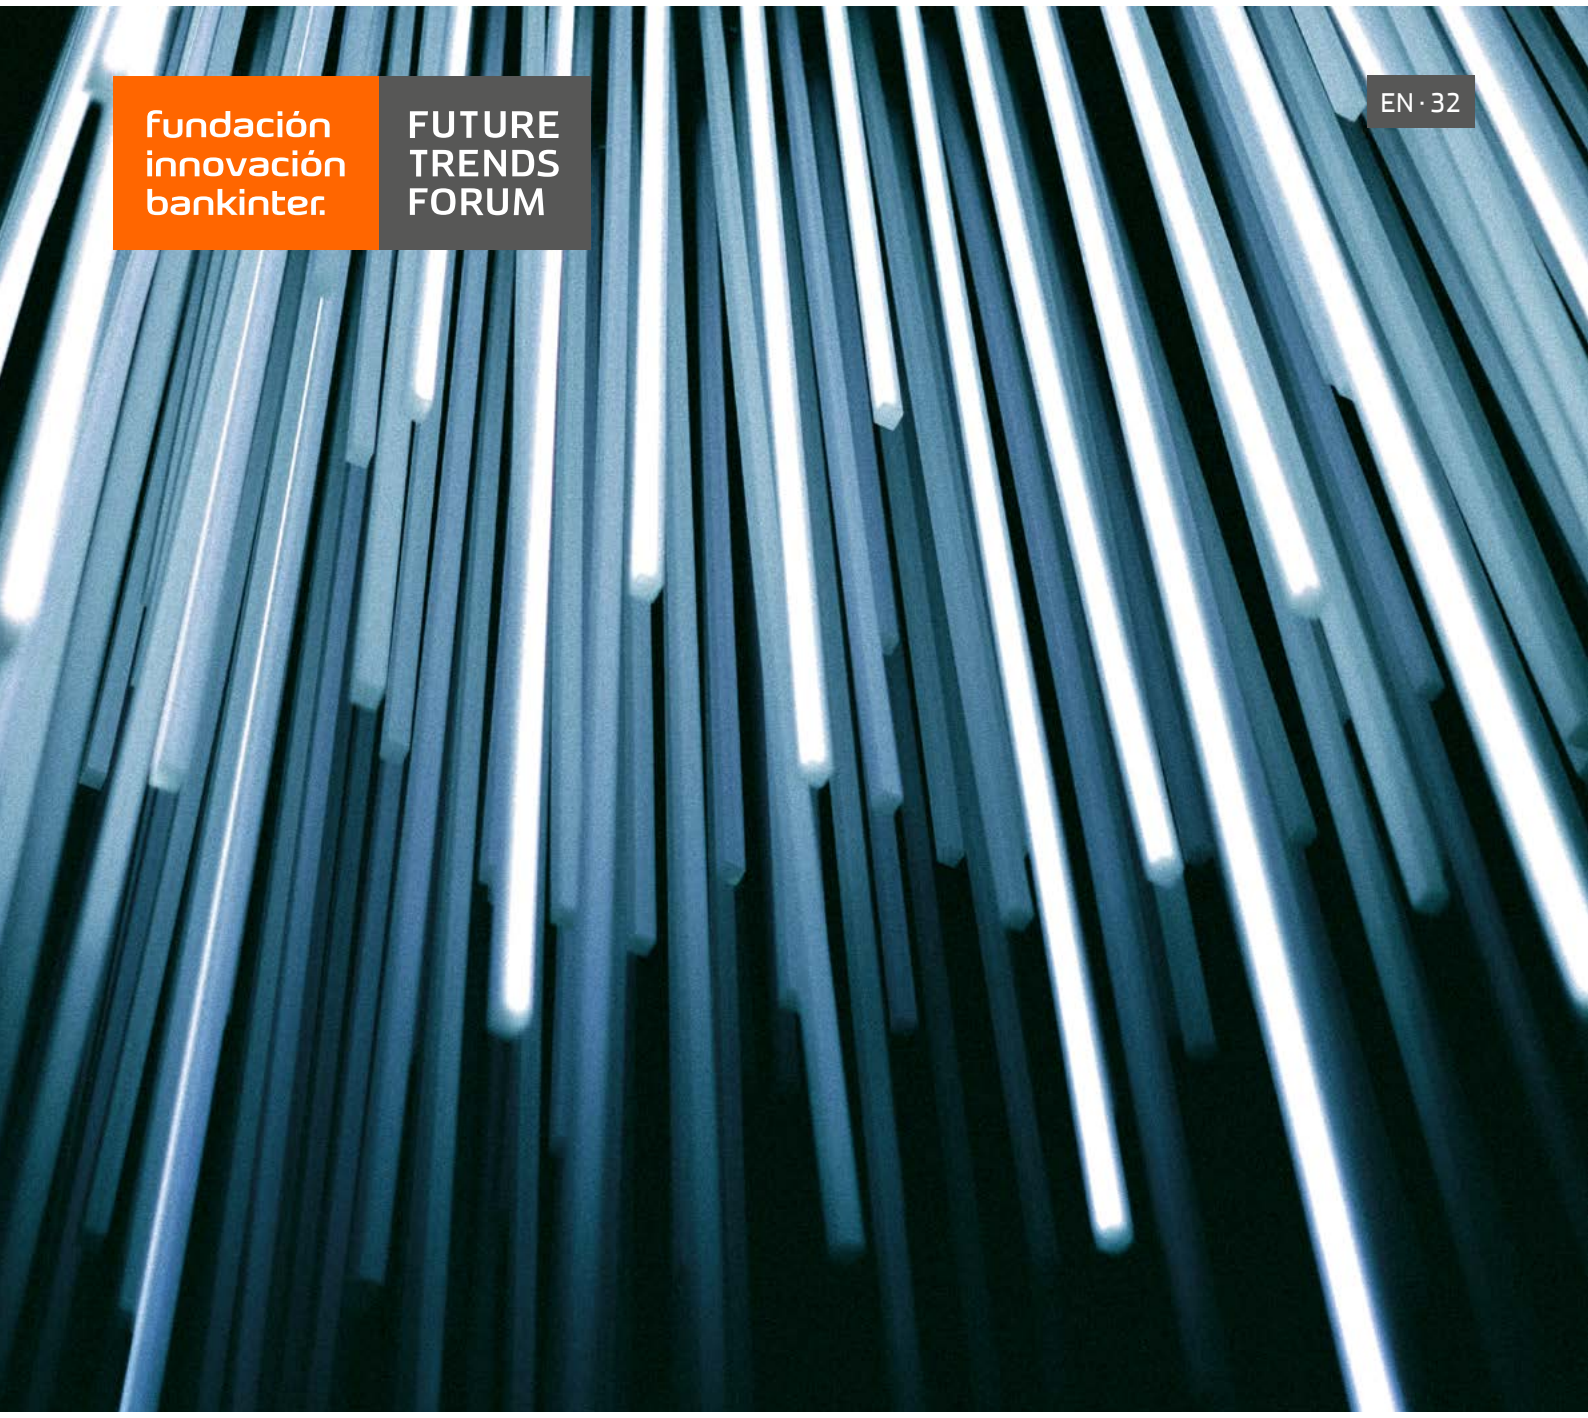

fundación innovación bankinter.

FUTURE TRENDS FORUM

EN · 32

Artificial Intelligence

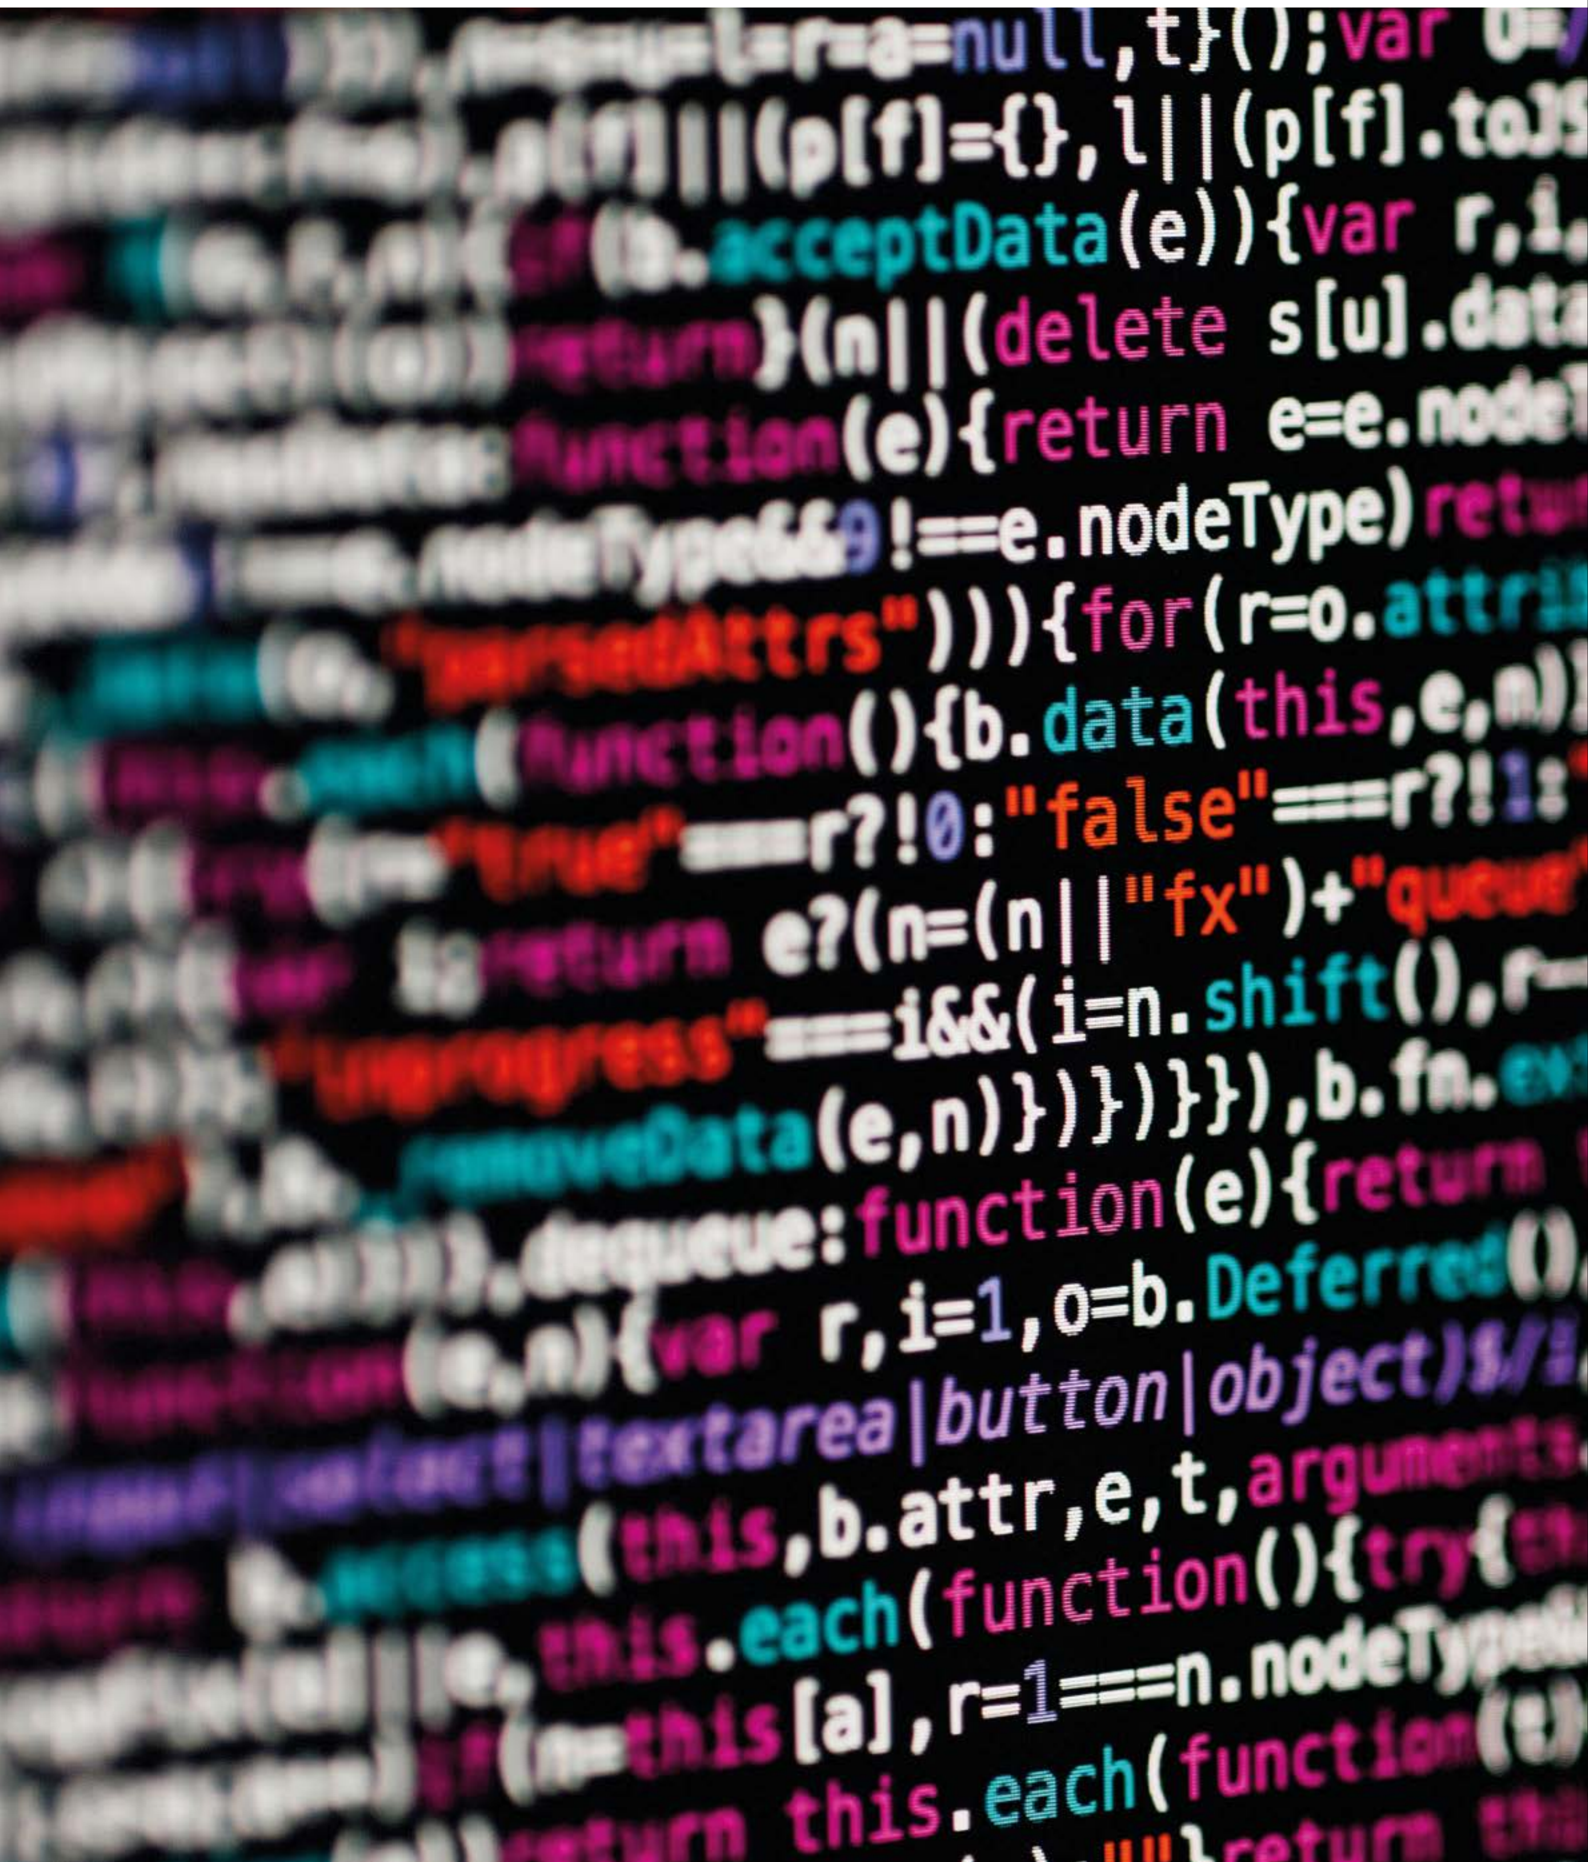
=null,t}();var
(p[f]={},l||(p[f].
acceptData(e)){var r,
}(n||(delete s[u].
(e){return e=e.
!==e.nodeType)
")))){for(r=o.
(){b.data(this,e,
r?!0:"false"===r?!
e?(n=(n||"fx")+"
===i&&(i=n.shift(),
(e,n)})})}}),b.fn.
function(e){
r,i=1,o=b.Deferre
textarea|button|object
,b.attr,e,t,
.each(function(){
[a],r=1===n.node
this.each(funct

Index

01

02

03

04

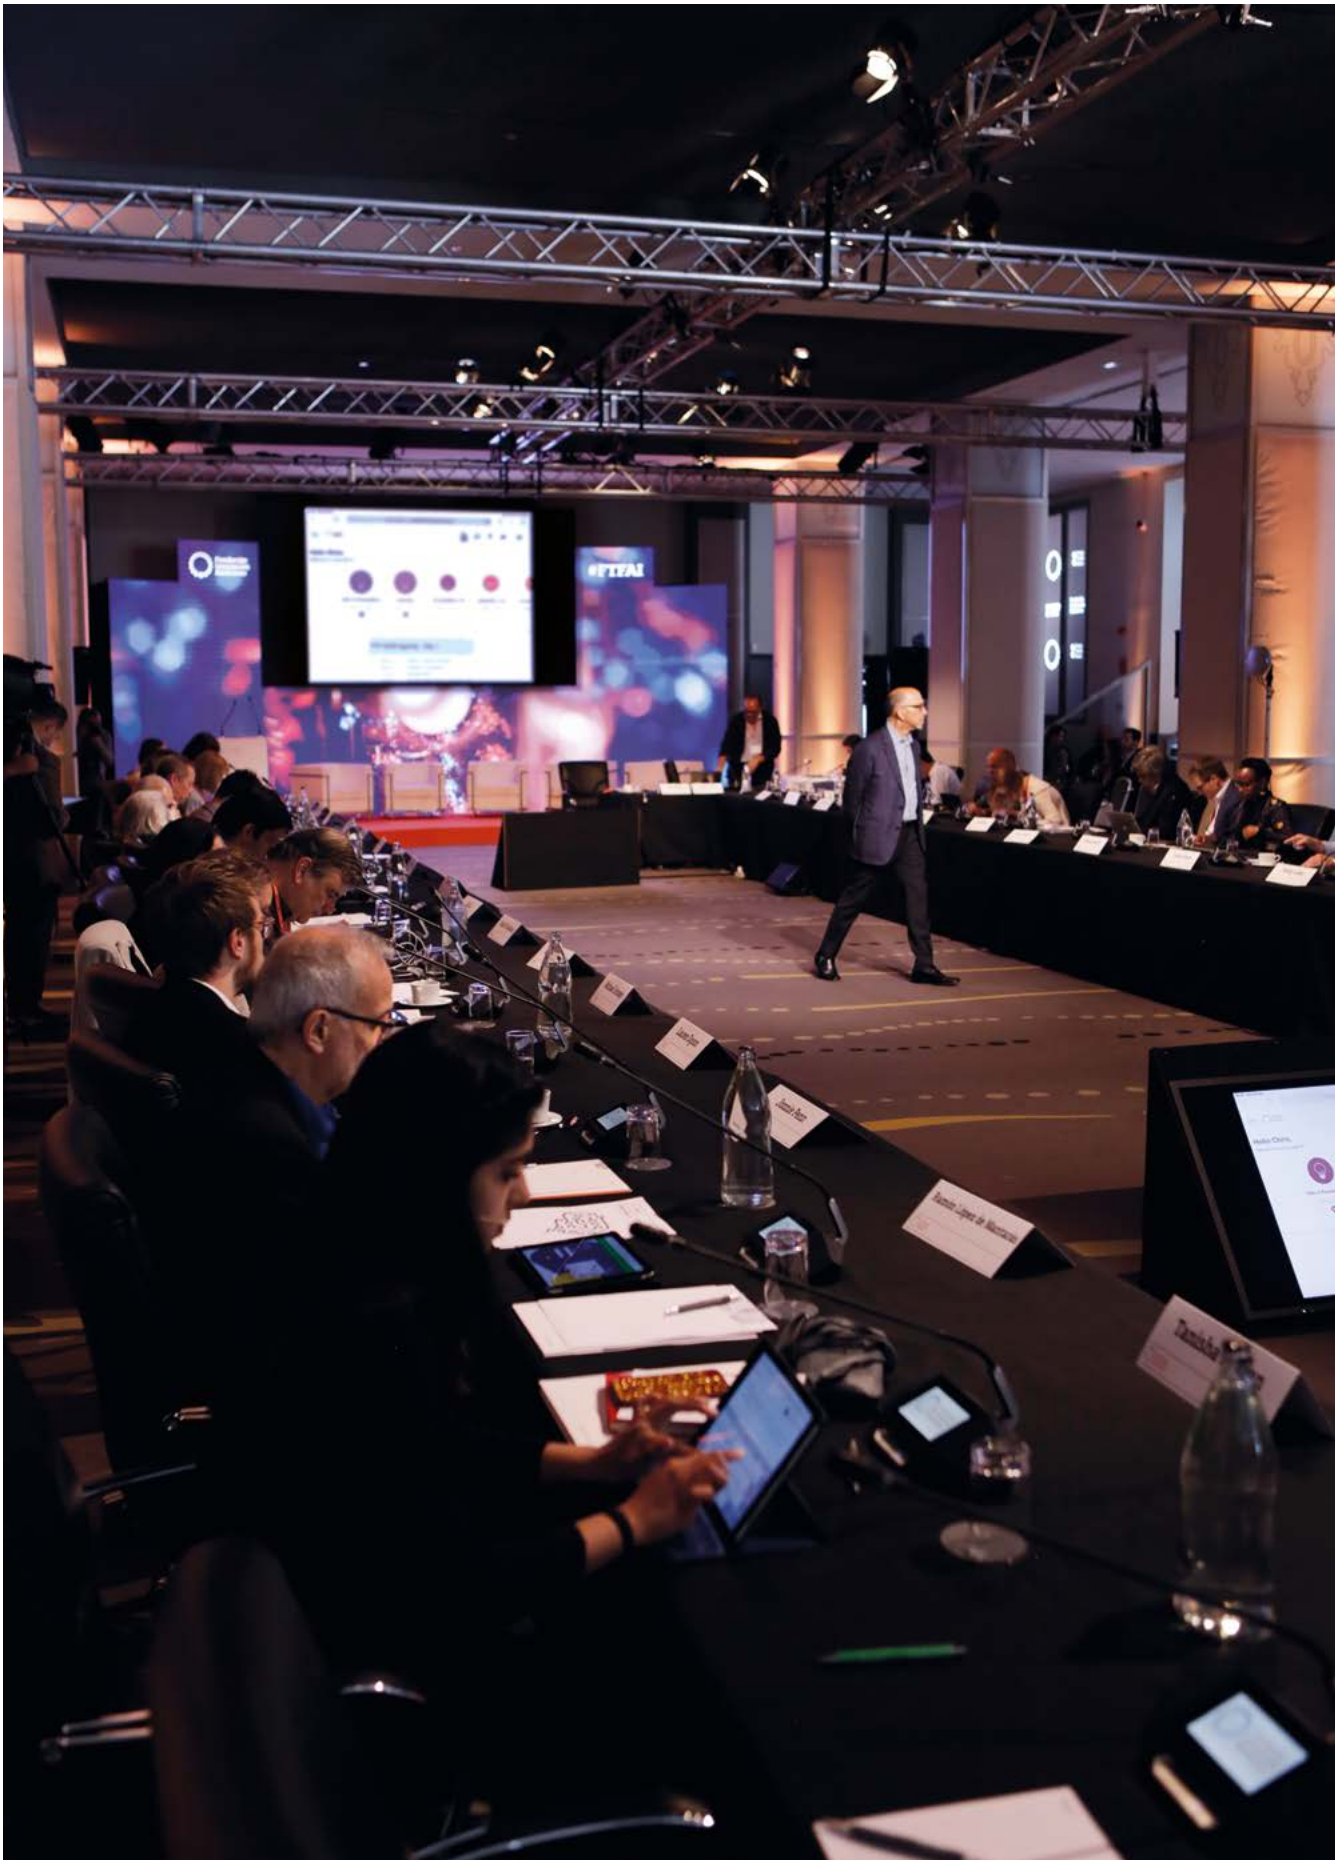
#FTFAI

01

Latest Developments in Artificial Intelligence

It's everywhere, almost ubiquitous. **It chooses our music, television shows, books, and forms of transportation**. It can see, read, listen and answer questions. But what is Artificial Intelligence exactly?

1.1

Introduction: AI, What it is, Why There is a Current Boom

It answers every online search. It obeys our orders at home as well as those of governments, business and all types of organizations. It can both clean the floor and write news articles. It translates texts and can recommend business strategies or a good diet. It analyzes our health, helps us drive or assists us in hiring employees. It learns from us to improve itself. It is Artificial Intelligence or AI. But what is it exactly?

"What is AI? You ask me, gesturing towards my blue screen with your cell phone. What is AI? And you're asking me? You are AI." That is how Gustavo Adolfo Bécquer, the poet, would respond. And he wouldn't wbe wrong. People, and their data, make AI. However, later we will see how AI is more than that. Before Bécquer, another poet, also a Spaniard, had written on the subject. Ramón Llull is regarded a herald to the technology.

This 13th century, medieval polymath was a poet, philosopher, theologian and martyr that was later declared Blessed by the Catholic Church. His ideas on religion led him to devise a sentient, mechanical wheel that would support or refute arguments. Reasoning would be artificially bestowed upon an object that would able to show the truths of the Christian faith so clearly that there would no room for debate. As Llull wrote in *Ars Magna*, the wheel would be able to demonstrate the rightness of Christian dogma.

There were also others before Llull. Some experts trace the origins of AI back to the ancient world. "The first automatons–anthropomorphic robots–that imitated human movement were built millennia ago. According to the Iliad, the Greek god of fire and metalworking, Hephaestus, made two women out of gold. There was 'intelligence in their hearts, strength and speech.' They did part of his work for him. Essentially, he had robots that helped him, which is why he was way ahead of his time." These are the words of our expert and engineer, **Nuria Oliver** from the speech she gave when she accepted her position at the Spanish Royal Academy of Engineering (RAI, in Spanish). Oliver participated in the FTF that led to the creation of this report.

Nuria Oliver
Chief Data Scientist at Data Pop Alliance & Fellow at RAI

By the mid-20th century, one of the so-called fathers of AI as we know it today, Marvin Minsky, had defined it as "the science of making machines do things that would require intelligence if done by men." This technology has since suffered many winters (periods of time defined by a drastic fall in interest and investment in AI) and enjoyed some springs. Today AI is thriving.

AI is basically algorithms. "They are sequences of steps that a machine, instead of a human, executes", explains José Manuel Molina in National Geographic. He is a computer science and Artificial Intelligence professor at the Universidad Carlos III in Madrid. "Those steps," he continues, "are not linear. Instead, they are conditional: the subsequent step is dictated by the fulfillment, or lack thereof, of certain requirements. If X happens then they will perform one action. If Y happens, they will perform a different action." AI's ability to learn on its own is known as deep learning; the algorithms work as a neural network that adjusts automatically. "The technique consists of creating a system on a computer, replicating the human brain: a set of neurons that are interconnected by weights and learn to function like a human being," he explains.

The current boom in AI is related to deep learning and machine learning

The current boom in AI is related to deep learning and machine learning. According to SAS, the multinational software company, it automates analytical model building by analyzing large amounts of data to identify patterns and make decisions with minimal human intervention, essentially learning on its own.

However, the idea that the system truly learns is widely disputed. Sinuhé Arroyo, expert on AI and semantic technologies and founder of Taiger, denies and rejects it. He states that there is no learning within machine learning, but training through data instead, because machine learning is included in the non-symbolic category of AI. It is based on statistics, contrary to the symbolic category of knowledge representation and automatic reasoning, which are based on logic. In this last category of semantics, there is learning. According to him, we can expand a knowledge base, add new concepts automatically, etc. and not many documents are needed for that learning.

According to SAS, while many machine learning algorithms have been around for a long time, the ability to automatically apply complex mathematical calculations to big data –over and over, faster and faster– is a recent development. The AI techniques that are currently available are basically the same as those used thirty years ago. However, infrastructure and the availability of data have changed.

Independent scientist and AI expert Alexandra Kirsch, from Xataka, explains that due to the Internet and the development of faster computers that increase processing power, we can meet new needs and use techniques that were available in the past but were not as relevant as they are today. Therefore the most interesting opportunities for AI application, as well as those with the brightest future, are now emerging. At the same time, the excitement surrounding AI's potential has grown significantly.

Critics question the very name of the technology. They do not consider AI to be intelligent and find the name to be misleading. The experts of this forum reminded us that Artificial Intelligence was originally developed to problem-solve... it was not a quest for intelligence. Therefore, this field should be named "complex information processing."

Pablo Gervás, is the director of the Natural Interaction based on Language research group at the Instituto Tecnología del Conocimiento (Knowledge Technology Institute) of the Complutense University of Madrid. He also agrees that the name "AI" is not very accurate. Gervás states that, even though each of AI's separate capabilities can be considered part of intelligence, theyaren't inherently intelligent. Take autonomous vehicles, for example. It's not that they know how to drive, but rather they know how to turn the steering wheel to remain in a lane, as well as other basic tasks. "We need to explain to people that machines are not like us. We have many skills and know how to combine them, how to switch from one to the other."

In fact, AI in its current form is limited

In fact, AI in its current form is limited. Once it is given parameters, AI is very good at carrying out a specific task. However, it does not possess the human ability of combining knowledge. An AI system may know how to play chess very well, but it will not also be able to detect tumors from an image. If it learns to detect tumors, it will forget how to play chess. This condition is called "catastrophic forgetting."

Ramón López de Mántaras asserts that no program or machine has ever passed the Turing test, which evaluates if a machine can exhibit intelligent behavior indistinguishable from that of a human by analyzing responses to written questions. López de Mántaras also believes that, even if a machine passed the test, it would not be considered intelligent. "In only resorting to that which can be expressed in written language, such as the answers to the Turing test, we create a very narrow definition of intelligence. Intelligence is much more than the ability to maintain a coherent conversation. In any case, the test evaluates aspects of intelligence that can be expressed in words," he explains.

Ramón López Mántaras
Director of the IIIA (Artificial Intelligence Research Institute) of the CSIC

López de Mántaras views AI as a system based on competence without comprehension, an idea that is widely shared and debated throughout the scientific community. It was first explained by the philosopher of mind and cognitive scientist Daniel Dennett. It refers to the fact that, at a functional level, a system can perform to such a degree (competence) that, if it were human, it would appear that the system could comprehend its own actions (in other words, intelligence), when, in fact, it does not. López de Mántaras believes that, a few years down the road, AI may be able to understand. However, he fears that we are approaching another AI winter "due to excessive promises that are being made."

The AI debate always has opposing sides. At the FTF, the views of some experts go beyond those of López de Mántaras. They think that AI will never be able to convince somebody and that this needs to be made clear in order to move forward. Others believe that it is not a matter of AI's ability to comprehend, but rather that it has developed a comprehension model different from humans. Yet other experts find AI's ability to comprehend irrelevant, while a different group views this ability as important because, according to them, AI's purpose is to learn and imitate intelligent behavior.

According to the famous neuroscientist **Antonio Damasio**, AI is nothing more than "a practical tool that humans created in order to expand their abilities and improve their existence." He believes that it is part of how we evolve.

Antonio Damasio
University Professor of Neuroscience at the University of Southern California

Virginia Dignum is a member of the European Commission High Level Expert Group on Artificial Intelligence and a professor within the Computer Science Department of Umeå University (Sweden). She believes that there is no such thing as Artificial Intelligence. "There are computational systems that use AI, which is a field of study and multipurpose technology, not an entity," she states.

The different descriptions from experts illustrate AI as a social construct, a concept born from a desire. This concept is living and changes over time. It reflects the anthropological, social and cultural aspects of humanity at different points in time. That is precisely what we find interesting about AI: the variety of views and visions.

What is this tool of capable of? José Hernández-Orallo is a professor at the Universidad de Valencia and participated in the HUMAINT seminar of the European Commission's Joint Research Centre (JRC-CAS). He outlines the main applications of AI:

- Knowledge representation through ontologies, different types of logic or probable inference (X is a bird, so X can fly)
- Temporary programming and planning as well as probabilistic planning
- Machine learning: linear models, decision trees, neural networks
- Pattern recognition
- Computer vision, facial recognition, biometrics
- Language processing: speech recognition, natural language generation, summarization, retrieval, translation, tagging, sentiment analysis, etc.

The Chinese investor and AI guru Kai-Fu Lee places the development of this set of technologies into four different waves:

- The Internet of AI: the first chapter of implementation, fed by the large amount of data that flows through the web. This data creates detailed profiles of our personalities, habits, demands and desires–the perfect recipe for the creation of personalized content to keep us on a given platform or maximize revenue or profits.
- Business intelligence: AI that can explore hidden correlations that escape our linear, cause-and-effect logic. These technologies can perform better than even the most seasoned experts.
- Intelligent perception: technology that can now see, listen and use thousands of other senses. It collects data that has never been collected before, which is used to create new applications. Intelligent devices and sensors, such as voice interfaces or computer vision applications, are a few examples.
- Autonomous AI: "the most monumental as well as complex wave," according to Lee. This wave includes the three prior ones. The results will be machines that can feel and respond to the world around them. They will be able to move intuitively and handle objects with the same ease as a human being. For now, autonomous vehicles are the most cutting-edge example of this technology.

A multitude of possible applications stem for AI's various capabilities. What is already possible? How are we using and incorporating it, both actively and passively, into our businesses and everyday lives? We will address this topic in the following section.

1.2
Applications. This is How We Live With AI

Almost without us realizing it, AI has been integrating itself into our daily lives. We carry it around in our pockets, on our screens. We have it at home, at school, at work and even as part of our leisure time. It is present in the services, industrial and agricultural sectors. It has had different levels of success, but it is paving the way for customization in areas such as medicine, education, even beauty and, in general, consumption. It also focuses on optimization, whether that is related to traffic, the supply chain, crops as well as mental and physical corporate performance. It can even compose music or create art. **Wilfried Vanhonacker** is the cofounder and former dean of CEIBS (China Europe International Business School)

and founder of SKOLKOVO, a Business Administration School in Moscow, Russia. He thinks that AI and big data together are not only creating the largest emerging market to this date. They go well beyond that. "They will dramatically change most industries. All business processes will be intelligent," he states. Vanhonacker believes that the discovery of new drugs and materials will be the sector with the highest economic potential. Far beyond the realm of business, he notes its applications in security, defense and surveillance.

Wilfried Vanhonacker
Co-founder and Ex-Dean, CEIBS (Shanghai) and MSM Skolkovo (Moscow).

The areas of business management and optimization, marketing and advertising, health and well-being, mobility, finance, legal services and policing, logistics and employment are the AI areas that the FTF experts highlighted:

1.2.1
Business Management and Optimization

"I will try to spend [the prize money] as irrationally as possible.," said Richard Thaler to The New York Times after learning that he won the 2017 Nobel Prize in Economics. The distinguished professor in the Behavioral Science and Economics Department of the Booth Business School at the University of Chicago was in fact recognized for his research on how emotions, not rational thinking, influence economic and financial decisions. Behavioral economics tries to understand how humans act as economic agents in decision-making processes, not always in the most optimal way.

Before Thaler, psychologist Daniel Kahneman was awarded a Nobel prize for his findings on behavioral economics within the framework of a new field: neuroeconomics. As Kahneman and Thaler's research suggests, people, both as individuals and as members of a group, take all types of cognitive shortcuts when making decisions, purchases and investments. In other words, "what people do tends to be much more revealing than what people say," explains **Michael Schrage.** He is an FTF expert, researcher at the MIT Initiative on the Digital Economy at the Massachusetts Institute of Technology and author of several books.

Schrage's work focuses on the behavioral economics of models, prototypes and experiments as a means to manage the risks and opportunities of innovation. In the case of AI, how it can be used to offer people not just better, timelier recommendations and tips, but also more persuasive ones. These recommendations can be used to optimize work and its management, to reduce uncertainty and increase sales.

"The impact that AI and machine learning have on an enriched data setting is that they become the gravity point and the key to the transformation of business capabilities," states the economist. He's referring to value creation, cost reduction, user experience, etc. "Capabilities become the lens through which we should look at how management is going to be affected from here on."

His hypothesis: the essential ingredient is the capacity to learn how to optimize. Optimization is the origin of differentiation with respect to what management and leadership mean in an AI setting. This is one of the key findings that Schrage has discovered together with his MIT colleagues after extensive analysis. This forces us to revisit basic ideas. It doesn't have to do with how we improve, simplify or synthesize our capabilities nor how we make them more efficient, but with what we want the outcome to look like.

"Netflix, Tencent, Google and Amazon lead by example; they are focused on results, whether that is cost reduction, value creation… what do we want the outcome to be?", he repeats. For the researcher, that question is going to create tension among governance, leadership and management. The debate will center around that question. In an ethical world, that becomes the true AI problem for companies all over the world, particularly in China. "Everyone will have to justify the purpose, the reason why they decided to optimize some dimensions instead of others."

Michael Schrage
Research Fellow with MIT Initiative on Digital Economy

Duration: 06:55 min

What does he believe is the true practical, realistic conclusion of this, which some FTF experts do not agree on? The KPIs are the strategy, and the strategy is the KPIs. Without non-financial, explicit KPIs that define what is going to be optimized, there is no strategy. Furthermore, according to his studies, machine learning has driven a huge change in those KPIs; they've gone from being results that helped people make decisions to being used as inputs for machines. That is why he believes that KPIs should not be considered metrics, but rather intelligent software agents that want to learn how to optimize.

The KPIs are the strategy, and the strategy is the KPIs

Back to the discussion around competition versus comprehension. Maybe systems don't know why they're doing something, but they're programmed to learn to optimize. That is why it's important to choose the right KPIs. If machines are going to make decisions, how does that change indicators and what type of data is needed? According to Schrage, the better the performance, the better the data and the better the decisions. And that is, in his opinion, the way companies should manage AI applications.

The better the performance, the better the data and the better the decisions

There's an additional component: how to create virtuous cycles among the KPIs and the digital setting they're in. "The original sin of legacy companies is that they have absolutely no clue how to do this. They can compartmentalize, visualize, create silos... but they're really bad at creating virtuous cycles. That is why they do so bad in stock markets and when competing against digital companies, because Facebook, Uber and all of these organizations truly understand the network effect."

Schrage believes that investing in virtuous cycles is essential. It will force a review of the anthology of optimization, what it means now and what it will mean three to five years from now. It's the new language for strategic planning, he says, in a context where companies are switching to micro services and prioritizing KPIs. He concludes, "The future of optimization is the future of machine learning and AI, and the future of AI and machine learning will be the future of optimization. That is the future of management."

The future of management depends on optimizing machine learning and AI

For **David Weinberger**, technology expert and researcher at the Harvard Berkman Klein Center for Internet & Society, the key is to handle uncertainty in the current setting of change, volatility and ambiguity. The philosopher and writer, a resident in the Artificial Intelligence Google PAIR program, points out that we have adopted techniques and behaviors that only make sense in a connected, digital and open environment.

David Weinberger
Senior Researcher, Harvard Berkman Klein Center for Internet & Society

As for business, Weinberger states that we have discovered how we can be successful by not trying to anticipate what will happen. A good example of this are open platforms created by companies for any developer to leverage on

their products and services, give them different types of usage and integrate them within other products that their customers use in their workflows. "Organizations create these open platforms because they know they cannot anticipate what an entire world of connected users will want to do with their products," he says.

Another example from the business perspective is the launch of products with a minimum set of features, so there is no need to try and predict what users will want to do with their products. These are both examples of a lack of anticipation that, according to the philosopher, is the hallmark of our online tactics and behaviors. However, that hasn't been the case in the offline world so far. "Not anticipating counteracts what has been our most basic strategy since the Paleolithic Age. The success of non-anticipation helps us acknowledge how unpredictable and chaotic our word really is."

In this context, machine learning models help us put that chaos into perspective, since they are not based on reducing complex situations to a bunch of well-known, general factors and principles; instead, they work through the connection of large volumes of data in statistically significant correlations. According to Weinberger, the results can be generated from a network of complex, delicate influences that human thinking cannot even reach. "The success of these models, many of which are incomprehensible to us, allows us to acknowledge and embrace the chaotic, contingent nature of our world," he concludes.

1.2.2
Marketing and Advertising

In the previous section, Weinberger talked about the success of non-anticipation and the need to embrace chaos. However, the Paleolithic inclination to predict is still present in the business world and shapes part of technological development.

"Marketing is obsessed with machine learning, which has to do with pattern recognition and data-driven predictions," states **Adam West,** Marketing Director of Satalia, an AI company that combines optimization and machine learning to solve hard-to-optimize problems.

Adam West
Head of Marketing at Satalia

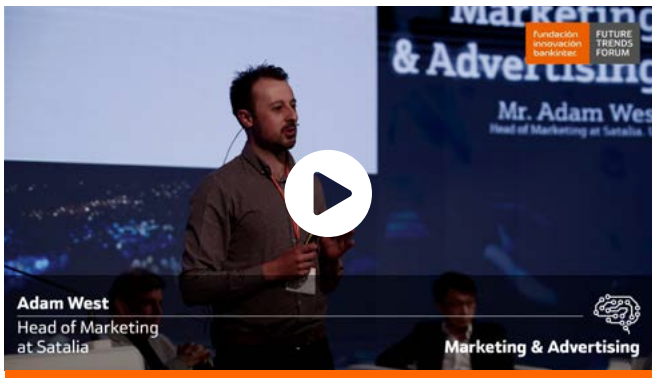

Duration: 07:21 min

Machine learning systems excel at certain forecasts

West points out that there are some predictions that these systems make quite well. For instance, with data on the volume of ice cream sold and daily temperature, machine learning can predict how many ice creams will be sold the following day. However, machine learning is not that good when solving resource allocation or decision-making problems. "It is not going to tell you, based on its prediction, the amount of ice cream you need to manufacture, number of employees you have to hire or how much you need to spend on advertising."

Why does this happen? West says that's the role of optimization, a set of technical skills that is "completely different, not as sexy as machine learning and not as exciting for the industry." The good news, according to research in Satalia, is that the true value lies in the combination of both.

Machine learning has plenty of applications within marketing

West agrees that machine learning has plenty of applications within marketing. It's possible to check a consumer's previous behavior and make a prediction about when they will stop being a customer of a specific company, predict the next product they're going to purchase and offer a product recommendation system, predict demand for a specific store... "That's all great, predictions are awesome, but the truly valuable asset is how to leverage them and what decisions are made with them."

Here are some examples: if predictions say someone will stop being a customer, that information should be used to try and prevent that from happening. Same thing goes for a store's demand. Predictions allow for optimizing the product inventory in order to maximize the customer's experience and match the number of people working assisting customers to that demand. Based on this, Satalia has helped with the routing of vehicles for Tesco supermarket chain and the optimization of PwC's labor force.

"We are obsessed with predictions, but we're not that good at combining such predictions with prescriptive analytics for optimization purposes." He believes that's the area that will see the most dramatic changes in terms of value creation for companies. The expert also mentions Nobel Prize winner Daniel Kahneman: we believe that, as humans, we are good at making decisions and solving these resource allocation problems, but in reality there are plenty of biases that intervene in the process.

That is exactly why we make bad decisions when leveraging on machine learning predictions. According to West, the fact that these systems are exponentially more complex also contributes to it. What is the solution, then? Focusing not just on machine learning but also on its combination with optimization and optimization algorithms.

West states that companies no longer invest heavily in advertising. Instead, they invest in operational capacity. The reason is that the way brands are built is changing. Traditionally, a company with a mediocre product would invest heavily in advertising, packaging and merchandising to wrap it in emotional and social associations. That was how the product stood out.

"That's how brands have been built, generally, over the last 30 or 40 years. Neither Coca-Cola nor Pepsi are competing for the functional benefit of water with sugar, but instead for the associations with the brand. Same thing goes for Nike and Adidas; they're not competing over the clothes, which are the same, but over the associations with the brand." West believes, however, that this is changing because it is no longer that easy to emulate the functional benefits of a product.

Brands stop using advertising and communications and focus directly on product innovation

West talks about the "Amazons and the Googles of the world" and how, even though they still have associations, they haven't achieved them by investing directly in them, but rather in product innovation, research and development. "The functional benefit is superior now, and that is how they differentiate from each other." There is therefore a transition when creating value. Brands stop using advertising and communications and focus directly on product innovation, which, unlike in previous cases, is not easy to emulate given the complexity of obtaining the right talent to build these types of capabilities.

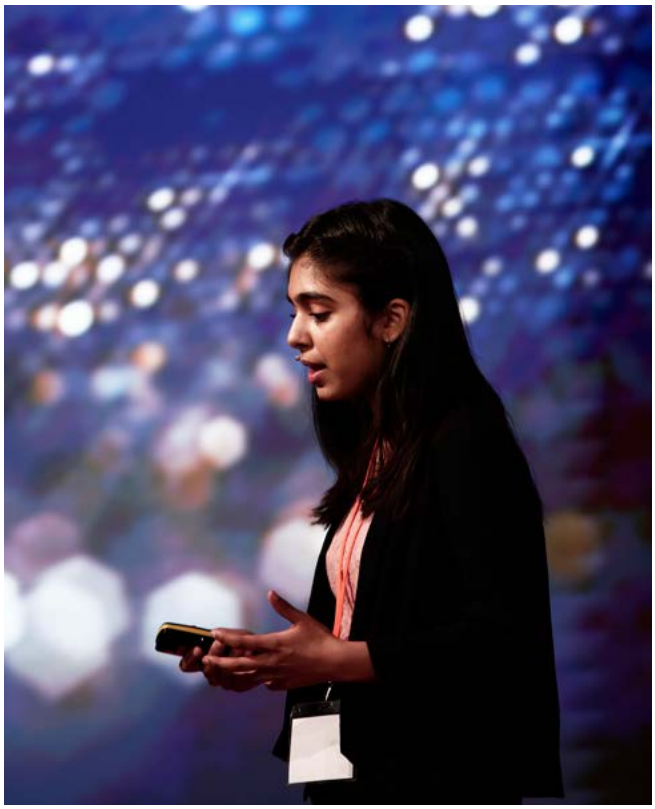

1.2.3

Health and Wellness

In previous sections we have mentioned behavioral economics and how it can influence people's decision-making processes due to known cognitive biases. What if those biases were used to help people change their bad habits and take up habits that would lead them to a healthier, more satisfying lifestyle?

"Changing habits is hard. Behavioral science, among other cognitive sciences, is helpful because it allows us to understand what makes those changes difficult," states **Oliver Smith,** Health Strategy Director at Telefónica Alpha.

Oliver Smith
Strategy Director, Health Moonshot at Telefonica Innovation Alpha

Smith would rather not talk about "bad habits" *per se*. He believes these exist because we've somehow found them useful or we've found "something good about them at some point, whether we were aware or not." The problem is that these habits might lead us to poor health choices, which often make us less happy. The expert explained that behavioral sciences can provide answers to questions such as why it's hard for us to depend on our will when we go on a diet, or why peer pressure (explicit or implicit) often prevents us from not drinking alcohol.

Duration: 12:13 min

In order to explain these tools—here is the challenge—we need to be able to measure what is happening at the right moment so we can make sense out of it. This is where AI comes in. Our cell phones compile data on our behavior that AI can then analyze with increased sophistication and detect patterns that would otherwise be missed. "At Alpha Health we believe that this combination will allow us to develop recommendations and services that will adapt to each person individually, from prevention to treatment, to help people change their habits."

His initial focus is on mental health and welfare. Regarding these topics, the promise of better health in the not-so-near future is rarely convincing enough to maintain behavioral changes over time. That is why, he says, it is essential to understand the main motivations at an individual level. Smith clarifies that Alpha Health doesn't intend to replace doctors, but rather complement them by focusing on the aspect of behavioral changes.

"75% of chronic diseases come as a result of our daily lifestyle: what we eat, how we sleep, what we drink." Alpha Health believes that it is possible to combine this with increasing knowledge on how people behave, why they make certain decisions, what percentage of those are made unconsciously, etc. By understanding the neuroscience, psychology and psychiatry behind these decisions, Smith believes we can create a personal medical assistant that can truly help people change their behavior. The missing factor is what Smith calls "the digital phenotype," the data that allows us to understand someone's personality so we can individually advise how to proceed in specific scenarios and under specific circumstances, based on who the person is and who they are with.

> The digital phenotype data enable us to understand somebody's personality

Smith believes that personal medical assistants "explain quite well where AI is going in terms of healthcare." He also thinks that predictions 'have a lot of potential.' Alpha Health and the UK's National Healthcare Service (NHS) conducted a joint study of anonymized medical records from a psychiatric hospital. They were able to predict anxiety crises, and other types of crises that patients could experience, two weeks before they happened with 80% accuracy. According to Smith, the results will be published soon.

Another area that is seeing significant work is telemedicine, which does not use AI but can become "a Trojan horse" for it (in the good sense of the word). "Companies that encourage the patient to communicate

with their doctor via telephone, or to be diagnosed through a videoconference call, are really after the patient's data to do other things that actually use AI. And that's really where we can change the rules of the game."

Lastly, he addresses what might be considered the broadest scope of application of AI to healthcare: diagnoses and predictions. They're just pure machine learning: extract a vast volume of data, develop an algorithm and instruct "this is what you should do." It sounds easy, but it's not. Take Watson, IBM's cognitive system that has come up against the huge technical and human challenges of using AI in clinical practice. This was discussed in **a profuse article** in *Spectrum magazine*, published by the US Institute of Electrical and Electronic Engineers (IEEE).

Researcher **Eric Topol**, who's also a professor and the director and founder of the Scripps Research Translational Institute, states in an **article** in Nature Medicine that virtually all doctors that practice medicine will use AI in the future. However, he warns, even though the field is surely promising, there's no strong evidence of those promises. The risk of running into faulty algorithms is also exponentially higher than in a single doctor-patient interaction, but there's a significant reduction in error making, inefficiencies and costs. He concludes, "AI in medicine cannot be an exception: it requires rigorous studies, the publication of results in magazines reviewed by experts, and clinical validation in a real-world setting, before it can be deployed and implemented to assist patients."

1.2.4
Mobility

First we saw the GPS, that essential application on our smartphone that is so basic and easy to use, but works thanks to AI. It not only gives us directions via text or voice, but it also checks traffic and lets us know if there's an accident or any other problem, to which it adapts in order to configure new routes. On its own, it is extremely useful, but it can be even more so for general traffic management in cities, with transportation systems that can be adapted to real-time demand.

Now let's add into the equation all autonomous vehicles that want to maximize the efficiency of this process. This is a promising application of AI to mobility that is still emerging, barely taking its first steps into the implementation of these vehicles across cities. "When I started working on this 13 years ago, at Stanford in 2006, there were two groups, Sebastian Thrun's group and mine. Each one would build a car and we focused on different issues. Sebastian addressed safety, I addressed the idea of car sharing," says **Raúl Rojas**. He has spent 20 years working on mobile robotics, 13 years on autonomous vehicles and 30 years on Artificial Intelligence. He is now a researcher and professor at the Freie Universität Berlin, Germany.

Raúl Rojas
Professor of AI at Freie Universität Berlin

Rojas believes that the future–already present–of transportation will be mobile phones. "You don't need to own a car, wash it or park it in the garage; you call a taxi instead. That's the best definition of autonomous vehicles I have. It's just like calling a taxi whose driver is a computer," he states. Rojas mentions several simulations on how many cars we can save if we share them. For instance, the one conducted in Berlin showed a statistical reduction of 90% in cars from the city's traffic through car sharing and vehicles that cooperate with public transportation.

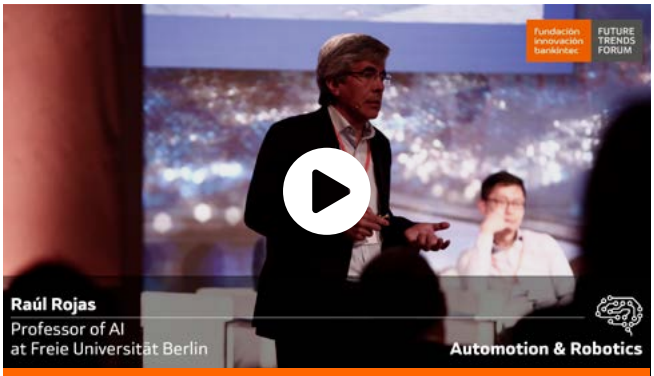

Duration: 16:15 min

Rojas talks about how the focus used to be on the driver, but soon it will be on the computer, which is currently the copilot. "Not many people know this idea of the computer being a copilot, they always think they're great drivers and they never have accidents, but that's only because the computer helps them." He's referring to curve stability programs on highways, among others. He believes that the idea of a computational copilot will keep being developed before seeing the full deployment of autonomous vehicles.

The concept of computational co-pilot will make its entrance before the self-driving car does.

Autonomous vehicle designers and manufacturers have complex challenges ahead of them. "Driving in Arizona is easy, but try to drive in Mexico. It's completely different. Passing a car on Mexican highways is very hard because there's no error margin, there's no room to the left or to the right. And there are other things that are different as well, like roads with no lanes where the car has to estimate as best as it can which is the right way to proceed." And that's just a small part of the problem.

How can we solve this? "Just like we do in many other aspects of Artificial Intelligence. Through brute force," he states. One example is his prototype, the autoNOMOS car, which has many sensors, four-times redundancy, nine video cameras, eight radars, six laser scanners at the bottom of the car and one on top. "All of this provides us with lots of information, that will be supplemented in the future with telecommunications."

However, this brute force approach crumbles, as Rojas acknowledges, when it is used in a context in which the autonomous car is surrounded by other cars with human drivers. Rojas remembers an autonomous driving experience in Berlin the day after a snowstorm. "It was completely white. AutoNOMOS knew exactly where the lanes were because it has 3D-positioning based on the environment's characteristics and a very precise GPS. But people didn't know where the lanes were. The first person at the front got in one supposed lane and the rest of the drivers followed. At the end, autoNOMOS was about two meters away from the rest of the cars. AutoNOMOS was doing the right thing, positioned in the middle of the lane. But since everyone else was further away, we were the atypical ones, the ones that were going to get into an accident."

AutoNOMOS checks the behavior of the other drivers, on top of the navigation map and the sensors.

After running into this problem, they started to develop what they call "swarm behavior." Now, AutoNOMOS not only checks the map and the sensors, but also the behavior of the rest of the drivers in order to adapt the way birds or fish do instinctually. 'There ARE rules for this swarm behavior, and that means that in this specific case, we had to get rid of valid highway rules and develop something new that takes external behavior into account. I think it proves that, at some point, we'll need to complement this brute-force approach with something smarter."

These problems are ground-related, but there are different challenges for sea and air mobility. We see an increasing number of drones, different types with different functions, ranging from rescue to the transportation of goods and all kinds of products. We see the beginnings of the development of driverless, private transport flying vehicles; companies such as Uber or Airbus already have prototypes that are being tested in controlled settings. Amsterdam wants to fill its channels with autonomous boats. All these applications are already here, and they would not be possible without AI and big data. Notwithstanding, its large-scale implementation and deployment will still take a few more years.

1.2.5

Fintech

2018 was a stellar year for the fintech industry. Over 1,700 agreements worth almost 40 billion dollars were signed and 39 unicorns were backed by venture capital firms, according to a CB Insights report. The industry will keep growing. According to a KPMG analysis, 2019 is the consolidation year for mature areas such as payments and loans, and also emerging areas such as blockchain. There will also be clearer regulation. And, according to Capgemini's World Fintech Report 2019, data will increasingly become a key asset and advantage.

Gary Ang
Managing Director of Strategy at Temasek

To KPMG, 2019 is also the year in which AI will become mature within the industry. How can AI add value to the industry? "Growth in fintech will be driven, on the one hand, by the regulation of Artificial Intelligence and data, and on the other hand by strategic competition," says **Gary Ang**. He is the Strategy Director at Temasek, a global investment company owned by the government of Singapore, whose portfolio is about 200 billion euros.

> From an AI perspective, a fintech company is just as good as the data it has

Duration: 07:59 min

Ang states that any sustainable and successful fintech company needs to have a business model based on constant data acquisition. Generally, data can be obtained in two ways: through the acquisition of a client or, once acquired, by selling the client products and services that generate more data. "From an AI perspective, a fintech company is just as good as the data it has."

Ang point outs that the most valuable asset is data acquired in the process, not the commissions associated with transactions. He mentions Ant Financial, Alibaba's fintech branch, the largest company in the industry. 90% of their added value comes from data acquisition (which means that barely 10% come from transactions).

How can we create added value with it? We can use AI to create brand new services that traditional banking does not provide. "50% of Chinese citizens have an Alipay or a WeChat account; these are mobile payment systems whose value reaches 17 trillion US dollars annually." He also points out that these fintech giants have become significant financial institutions for the Chinese government in the span of 10 years.

Another innovative fintech company mentioned by the Temasek executive is Square (USA). It started as a tool for small businesses that couldn't process credit card payments, by offering them a mobile payment system. It quickly evolved and started basing its business on the provision of software and loan services to the unbanked sector. Additionally, the company is in the process of obtaining a banking license. Through their data analytics, they have been able to create products that didn't exist before, making clients more dependent on Square and its ecosystem.

Ang mentions that the second main added value in the fintech industry is, in fact, the creation of an ecosystem that acts as a fully digital bank. Whether they are in China or in the US, the most successful fintech companies such as Ant and Square, want to escalate in order to increase their scope. By moving large sums of money and challenging the domains of traditional banking, they also become digital banks. Then they're subject to regulation. This regulation might be based on traditional data metrics, but also on new AI-based metrics.

Ang explains that Singapore is testing how this would play out through a sandbox system (creating a safe space to test out innovation and technology). He states that regulation will have a strong impact on winners and losers because, at the national and strategic level, it is driven by nation states that compete against each other. "Fintech companies on the rise will need to watch out for open markets and also be aware of how to navigate regulations to address underlying national and strategic concerns."

According to The Ang, there are two key components for the regulation of fintech companies. The first one is traditional regulation, where global standards have already been defined. The second one is trust, which is key for users to share their data. Traditional financial services that face the threat of disruption, including the incumbent retail banks and insurance banking, are opting for the strategic resource of trust. "They are starting to focus not on what we can do with data, but on what we should do with data, and all the odds are in their favor."

If there's no trust, it is very unlikely that consumers will be thrilled about this open banking model in which traditional banking customers' data is shared with fintech companies. "These companies have done a good job in terms of providing low-cost services to the unbanked, but we still don't know how sustainable this model is, because there are plenty of data leaks and cyber attacks", the expert points out.

This is where the ethical use of data and AI plays an essential role, as we will see in the next section.

1.2.6
Other Areas

During the FTF, experts also mentioned other application areas for Artificial Intelligence. The professor and researcher **Christopher Markou**, from the University of Cambridge, UK, specializes in Law and Artificial Intelligence. He analyzes how AI is changing the legal system, which has traditionally resisted technological change. "In Law, there was never a real motivation to adopt new technologies. The system has always been based on a bill-per-hour model. Therefore, working faster meant making less money, so big large law firms did not have incentives to adopt new technologies and accelerate processes," he explains.

Christopher Markou
Researcher & Lecturer on AI at University of Cambridge and King's College London

There are multiple areas of AI application in the legal and police areas

AI is changing the game. First, according to Markou, the incentives structure has changed: now it's hard for companies to resist using AI to help with legal matters, including in the court room. Some countries, such as Estonia, are testing out robots that can replace human judges and attorneys. "Artificial Intelligence is going to change the concept and the practice of Law. It already has. But what I think is going to the happen in the future, what people like me are concerned about, is that human attorneys and judges might be replaced by computational systems that might not share any of the values of our legislation and legal system."

Lauren Dyson, a data scientist at Civis Analytics, addresses AI applications in police work. "Most of the police stations in the US use predictive algorithms to decide where police officers should be sent in order to prevent as many criminal offenses as possible," she says. This wissue is controversial, since it has been proven that these systems are detrimental to people

of color, to whom are attributed a higher criminal risk (we will address this topic in the next chapter, together with aspects regarding the computerization of the legal system, among other topics).

Lauren Dyson
Senior Data Scientist
at Civic Analytics

Andrés Torrubia is the founder and executive director of Fixr. He's a serial entrepreneur and an expert at e-commerce. Regarding logistics, he states that AI is impacting online commerce (and offline commerce as well) in different ways. Some of the changes are obvious and visible, such as better product recommendations or suggestions about when to buy recurring products. AI also has an impact on logistics optimization, from delivery-route planning to delivery by drone or even autonomous robots in warehouses. Creative departments use AI to help them come up with future products or even let the programs generate products or designs on their own. "And there are plenty of examples. E-commerce has been doing many things in the most rudimentary way, but that is changing thanks to AI," he states.

Andrés Torrubia
CEO of Fixr.com

Torrubia mentioned the automation of work management in warehouses. This work automation is happening across the board, not just in manufacturing and industrial areas, but in many other sectors as well. For example, journalism and mass media. The Washington Post, Bloomberg and many other news outlets use bots to generate factual news in finance or sports or just to follow trends. Thanks to AI, Deutsche Welle, the German public media corporation, has created an automated transcription, translation and voice tool for video content and other simultaneous transcription, content distribution and fact checking solutions.

Without a doubt, all of this allows processes to be more efficient, even though editors emphasize that AI will never be able to replace human judgment and that human beings will always be a part of the news production chain. Additionally, if machines take up repetitive tasks, people will be able to handle the most important, value added work. That's the promise of automation: putting an end to aligning tasks and serving as a tool to increase workers' capabilities. However, the promise entails the loss of jobs to machines. That is another issue that we will address in the next chapter on the social and ethical impact of AI.

02

The Social Context of AI

Technology has made it possible for us to hunt, form groups, communicate, settle in cities, increase our life expectancy and connect our ideas. Yet it has also driven warfare, ideological battles and inequality. It has brought us to the brink of destruction, pushing humanity almost to the verge of extinction in order to acquire more powerful technologies. Essentially, this is the balance between technology and risks.

This was the opening statement from **Jade Leung** at the FTF. She is Head of Research & Partnerships at the Center for the Governance of Artificial Intelligence (GovAI), a team within the Future of Humanity Institute at Oxford University. Before we implement and use a given technology, we must weigh its benefits and drawbacks. ArtificialIntelligence is no different.

In the previous chapter we discussed current AI applications and how technology can improve our health, business management and optimization, mobility, finance, logistics, legal services and decisions. Certain drawbacks come with the benefits, or, at the very least, there are risks and challenges surrounding this technology's implementation. We will address this topic in the following section.

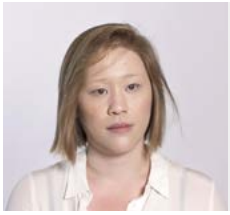

Jade Leung
Head of Research & Partnerships at GovAI

2.1 The Social Impact of Artificial Intelligence. Risks and Challenges

2.1.1 Impact on Human Behavior

Ethics researcher **Fiona McEvoy**, founder of YouTheData.com, concurs and points out the wide range of ethical and social challenges and concerns that we face. She

points to algorithmic biases, unfair systems or those that insufficiently respect privacy and security. She highlights systems that try to influence our opinions or decisions for the benefit of others as well as systems that promote our addiction to being connected.

Fiona McEvoy
Founder of
YouTheData.com

McEvoy maintains that asymmetric information is at the root of all of this. There is an imbalance between users (everyday people) on one side and big technology companies (or governments, in some cases) on the other. The latter is in the driver's seat when it comes to data- and AI-driven technologies.

Their area of expertise is influence. They use AI to modify our behavior, based on all of our information (our likes, our preferences or our location) and predictions concerning our likes, location, purchases, etc. According to McEvoy, "In some cases, companies exploit our cognitive biases, the irrational part of our decision-making architecture, to nudge us towards buying their products and services."

Duration: 07:23 min

This worries her because we seek out increasingly immersive environments. McEvoy cites virtual assistants, extended reality (virtual, augmented and mixed) where complete environments are created for the user. As an example of its potential, she mentions a 2009 study from Stanford University, when these technologies were much less sophisticated. The study determined that a few weeks after a group of children were exposed to a virtual reality experience in which they swam with orcas, they remembered that experience as if it was real.

McEvoy maintains that with these technologies and greater personalization through AI, entities can significantly increase their ability to influence users. "All of the data that defines our habits, our clicks, will be combined in real-time with dynamic, biometric information on our bodies, including movements, such as where our eyes look, or our vital signs, like our heart rate."

She asserts that part of the reason why these systems are able to manipulate people is because they are designed with human characteristics: appearance, voice or behavior that seems human. McEvoy believes that this stems from the tech community's fascination with trying to replicate humans. "I understand that it may be necessary to an extent, but maybe we should reexamine how necessary it is." People should feel that the app, assistant, robot or program with which they are talking is simply that: a device. This comes down to transparency.

Dor Skuler, CEO & Co-Founder of Intuition Robotics, states that he agrees with technology's transparency and honesty in not trying to look human. "That's why we designed ElliQ, to make an 'objectoid' as opposed to a 'humanoid.' ElliQ does not look like a human. It neither has eyes nor hands, and the voice is clearly mechanical." However, he believes that certain design rules should not me broken. The design of these devices should be based on the needs and goals of the product that is being developed.

Dor Skuler
CEO & Co-Founder
of Intuition Robotics

According to McEvoy, AI can also impact our behavior through our belief that we are being watched in public. "On the streets of San Francisco (USA), where I live, facial recognition has just been banned. In the United Kingdom, where I'm from, the authorities are using technology to an increasingly aggressive extent in order to collect very personal, biometric facial data." That's what the economist and emeritus professor at Harvard **Shoshana Zuboff** describes in her book *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (Public Affairs, 2019). This is an Orwellian image of a Big Brother that uses increasingly sophisticated technology to control people, either as citizens or users.

McEvoy's final argument concerning influence over our behavior relates to authenticity in two senses of the term. The first concerns our appearance. "It seems that the whole world knows everything about us. There is more information out there about me than I know about myself or that I can remember. But more and more, as users we know less, and we're deceived regularly," she

states. The second aspect of authenticity is related to misinformation, fake news or so-called *deep fakes*, which use deep-learning systems to create sophisticated fake content (especially videos or images) that do not seem obvious to the naked eye. She concludes, "It is becoming harder and harder to distinguish real from fake, authentic from false."

2.1.2
Impact on Employment

In the previous chapter, we discussed how the automation that drives AI coincides with the loss of jobs to machines.

"I want to emphasize that, like any technological change of this magnitude, production automation and the manufacturing of autonomous vehicles will affect employment in the automotive industry, which is of vital importance," said Raúl Rojas. He is a researcher and professor at the Free University of Berlin (Germany). In the previous chapter, he spoke about his experience in developing autonomous vehicles.

Self-driving vehicles will reduce the total number of vehicles on the road and will naturally affect sales as well as driver employment for TNCs, taxis and trucks. Rojas believes that we must find a solution. "Sometimes I feel a bit guilty that I work in this field because AI is a deal with the devil. On the one hand, we're trying to do good, but there are also negative effects. We need to find a balance," he states. "We need to be aware that what we're doing in AI negatively impacts many people."

That seems like an obvious statement, but plenty of people make the mistake of denying it, according to **Calum Chace**, writer and expert lecturer on AI. "They reject the idea of job loss as if it were a fallacy," says the founder of the Economic Singularity Club. The organization defines itself as "a group of technologists, academics and writers who think the threat of mass technological unemployment is worth taking seriously." According to Chase, skeptics tend to refer to the agricultural sector as an example; it went from accounting for 80% of US workers to 1% currently.

Calum Chace
Writer and Speaker on AI

"It was a very difficult transition. And, frankly, many people died, but at least it didn't cause widespread, irreversible technological unemployment," he jokes. People generally associate automation with mechanization. When machines replaced manual labor, we were still able to leverage our cognitive abilities. Yet the latter are now becoming automated. According to Chase, this is what the future has in store for us.

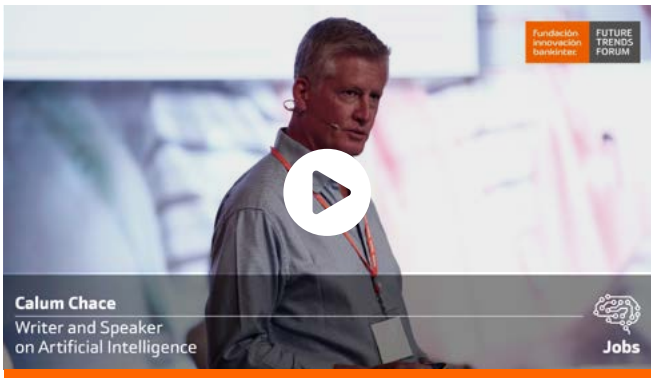

Duration: 05:11 min

He believes that Moore's Law (which states that the number of transistors on integrated circuits doubles every year) is changing. Chase says, "Our machines are becoming much more powerful." According to this rule, within ten years, machines will be 128 times as powerful as they are today and one million times as powerful within 30 years. Thus, Chace believes that "saying technological unemployment will not happen because it has never been the case is akin to complacency."

What are we going to automate? "Rojas told us that probably by 2030 there won't be many taxi drivers in some cities. And I think he's right." He projects that taxi drivers and truck drivers will switch to jobs in warehouses, factories or call centers, which are also becoming increasingly automated. "Automation will not just affect cheap and repetitive manual labor. All professions will be affected."

The report ***How robots change the world. What automation really means for jobs and productivity***, from Oxford Economics (2019), points out how automation is also starting to affect the services sector. According to the report, innovations in Artificial Intelligence, machine learning and computing power suggest that robots are being implemented at a significantly increased rate in the sector, especially in term of logistics. Robot implementation is led, in part, by the global expansion of Amazon and other multinational, e-commerce companies. However, the study highlights that it will be difficult for machines to replace humans in professions that require compassion, creativity and social intelligence, such as physical therapists or social workers.

On the other hand, the book ***The Technology Trap Capital, Labor, and Power in the Age of Automation*** (Princeton University Press, 2019) debunks one of the most quoted figures regarding work automation. In a 2013 article, two researchers from Oxford University stated that 47% of US jobs face a high risk of automation by midway through the 2030s. Now, one of the authors, Carl Benedikt Frey, wants to clarify the misunderstanding in his new publication. "It is not true that half of all jobs will be automated within one or two decades," he states in an interview with ***The Economist***. Frey explains that the study predicted the percentage of occupations that are most vulnerable to automation, which does not necessarily mean that they will become automated. According to the Economist article, "That, they underscore, will depend on many other things, such as cost, regulatory concerns, political pressure and social resistance."

Frey believes that the biased use of the data from his report reflects the polarized debate surrounding the nature of automation and the future of employment. He is not worried about the destruction of jobs. Instead, he is worried about employment instability. "This will lead to unrest and opposition that could hinder automation and growth in productivity." He is not troubled by a world of too many robots; he is concerned that there may be very few. Frey believes that lawmakers must manage the transition by expanding wage insurance and other forms of compensation as well as instituting education reform and other measures.

Mary L. Gray, an anthropologist and researcher at Microsoft, also refutes the claim that automation will destroy jobs. She points to "the paradox of automation's last mile", in which the desire to eliminate human jobs always leads to new tasks for humans. She talks about this in her recent book ***Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*** (Houghton Mifflin, 2019), in which she and social computer specialist Siddharth Suri examine the impact of automation across workers' experiences in the online, on-demand economy.

The researchers discuss the notion of *human computation*. These are jobs that start and end online, comprising any type of task that can be managed, processed, carried out and remunerated online. For instance, the ones requested through platforms such as Amazon Mechanical Turk. According to a study from Cogniltyica, by the end of 2023, the AI-related, tagging tasks alone will be at least a billion-dollar, global market.

These jobs (tagging, classifying, identifying hate speech, etc.) fuel the systems, websites and AI applications that we all use and see as part of our everyday lives. TripAdvisor, Match.com, Google, Twitter, Facebook and even Microsoft are just a few of the most famous companies that create on-demand tasks on these platforms. According to the

authors, "Everyday there are new companies with business models that depend on these tasks. Not only is this type of work growing, but it in fact represents a broader, systemic restructuring of employment." The authors say that the problem with these types of jobs is that they could obscure the work of hundreds of millions of people. That is why they call these types of jobs "ghost work."

There are multiple areas where AI can be applied, from the law to the labor market

Just like Frey, Gray points out that part of the solution is to increase social coverage for workers, among many other potential measures. We will discuss the aforementioned measures and other possible solutions in the next section.

2.1.3
Impact on Justice

In the previous chapter, **Christopher Markou** talked about how AI is impacting the legal system, specifically incentive systems, bureaucracy, management and administrative tasks, even attempting to replace lawyers and judges with robots. What are the limits of automating the law? If the law is nothing more than an elaborate system of rules that may be reinforced by a sanction or threat, then perhaps a large portion of the law, to a certain extent, may be quantifiable, according to the researcher and professor at the University of Cambridge.

Duration: 08:59 min

Markou believes that is the wrong way to approach this. "First and foremost, the law is a social institution with rules, categories and socially constructed frameworks that we use to understand the world." Laws have evolved over long periods of time. Markou believes that if we try to formalize them through logical or mathematical methods, we run the risk of oversimplifying the world to make it fit into a quantifiable model. We also restrict one of the law's most important, social and anthropological functions: protecting society and the individual from the adverse and "potentially dehumanizing" effects of technology.

Codifying laws and legal knowledge into a machine-readable system is incredibly difficult

According to the expert, codifying laws and legal knowledge into a machine-readable system is also incredibly difficult. This is because, among other things, legal counsel and legal decisions are not provided in a vacuum. Markou explains, "A competent lawyer must be able to understand and intuitively know why something like a snail in a bottle of ginger beer is a bad thing. It is natural for human intelligence to make these kinds of complex inferences based on facts, but it is extremely difficult for even the most advanced computer system."

Markou notes that another obstacle to automating laws is that the development of machine learning algorithms requires that an "objective function" be set. This is the empirical measure of success or precision of a given task. It is impossible to establish an objective function for legal decisions, where belief in what is "good" or "optimal" varies depending on the client, cultural context and changes over time. Furthermore, according to Markou, many legal questions require a very precise model of the human mind. For example, the "reasonable person" standard requires that a lawyer interpret a contract by creating a mental framework of how a hypothetical person would act. "It's not enough for a computer to emulate aspects of human intelligence. It must be able to think like a human."

Markou concludes by stating that the law is ultimately language, which makes it "extremely difficult for a computer to decode and understand because language is incredibly vague." As an example, he cites the problem of common sense: knowing things and applying them to new scenarios or solving problems with incomplete information. "Even though AI is improving at completing these tasks, it is far from the intelligence that an average child possesses."

This is why Markou is skeptical of Estonia's recent attempts to implement "robo-judges." The real danger is not that these systems work, which would be great. The problem is if they fail but we believe they work, and then we use them under that false assumption", stated Markou in an interview with *Innovadores*. According to Markou, even though there are reasons to believe that Artificial Intelligence can help in relatively simple cases, the key issue is determining the other legal contexts in which we should allow algorithmic judges to make legal decisions.

Philosopher Lorena Jaume-Palasí, founder of Ethical Tech Society, greatly agrees with Markou's analysis. "AI cannot put things in context. That is why a software program cannot perform legal interpretation through automation. This is the job of a human judge and is a task that requires putting things in context. Nor does she think it's possible or even desirable in the near future. "I believe that this is one of the tasks in the legal field that should never be automated."

However, Jaume-Palasí mentions that these systems can help detect inconsistencies in judges' legal decisions because "external factors like time limits, hunger, or time of day have an impact on decision making." But the essential parts of a judge's job (contemplating the law, searching for aggravating and mitigating factors, determining whether the verdict conflicts with the ethical inklings of society, etc.) cannot be automated.

Jaume-Palasí finds the efforts to automate most alarming. Are people trying to make the justice system fairer and more efficient, or do they want to create machines to replace people just because it is possible? Markou shares this concern. "Just because we have technologies that are able to automate the justice system doesn't mean that we should do it. Think about the decision to use the atomic bomb in World War II. First, we should think about the consequences of using them, exercising caution and foresight."

Therefore, the efficiency and speed associated with technology are not priorities when addressing the legal system. According to Markou, justice is the priority and, ideally, it is delivered swiftly. Evan Selinger, a philosopher, and Clive Thompson, a journalist, echo this sentiment in *The Efficiency Delusion*. Striving for efficiency is not always innocuous. Placing efficiency at the top of the list can be a lapse in ethical, political, personal or professional judgement. Some personal or civic interactions flourish when they are conducted deliberately. They fall apart when accelerated.

Whether or not the changes occur rapidly, Markou firmly believes that the legal system will transform substantially. "There's no reason to exclude technology as a whole from improving the legal system. However, governments and the general public must be critical and skeptical regarding the nature, degree and pace of this transformation." Markou concludes, "We must have robust and effective deliberative measures to help guide development of the technology and ensure that it is used fairly in service of all.

2.1.4

Impact on Governance

Jade Leung spoke of her experience in governance. "We have quite the long-term view regarding the meaning of governance in this context and the type of consequences we need to focus on." Leung holds a PhD in AI Geopolitics from Oxford University and her work centers around shaping relationships among business, government and the scientific community as well as identifying the dynamics that lead to cooperation and conflict.

The essence of Leung's speech is encapsulated in the following reflection: if we look back to the origins of any technology, we can see that it was clearly and utterly wrong to think that we could predict and anticipate how the technology should be governed from the outset. "Sometimes we are cautious and lucky enough to choose a path that benefits society, and sometimes we're not so lucky. Sometimes we are tempted by the power of nuclear weapons, or we use genetic engineering irresponsibly. Occasionally, we implement technologies even when we are uncertain of the consequences. Then we have a technology that causes great uncertainty, which is an occupational risk for the governance of today's technologies."

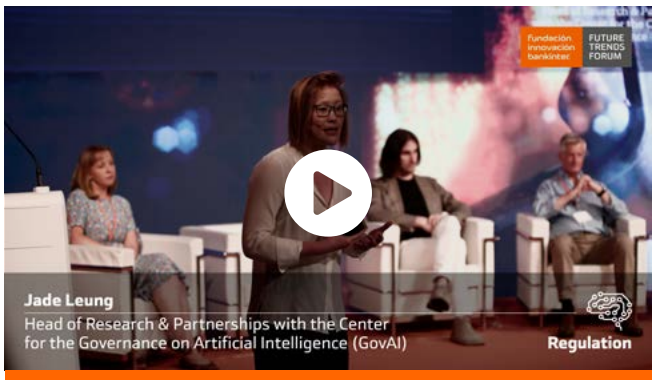

Duration: 08:08 min

These governance challenges require great effort when traversing, with limited resources, the mountains of opportunities to avoid falling into the valleys of risk. AI is a current navigation challenge. According to Leung, it is one of the most complicated challenges for two key reasons. First, AI is a general-

purpose technology (GPT). "This type of technology transforms the key functions of civilization, such as energy production or communication. AI transforms intelligent processing."

AI is a general-purpose technology

GPT come with strong incentives for the speedy extension of their functions throughout society. "It's exciting to see their economic potential, but governance is more complicated because, naturally, there is neither centralized control nor centralized governance objectives. That makes things complicated." Leung adds, "We haven't seen good examples of management of these technologies, and we've made many mistakes along the way."

This technology is also difficult to regulate because we are trapped in an environment of strategic competition within the industry. We've gone from the space race to the AI race. The clear frontrunners vying for the lead are the US and China. These two superpowers are supplying massive amounts of political and financial capital, both public and private, to lead the field. Seven out of the ten largest corporations in the world by market value -according to data from the Kantar's 2019 report ***BrandZ Top Global Brands***- are companies that place AI at their very core (Amazon, Apple, Alphabet, Microsoft, Facebook, Alibaba and Tencent), and they invest large sums of money in R+D. "Both in the public and private sectors there is strong competition. This isn't necessarily a bad thing, unless it exacerbates structural risks", says the expert.

These are not the only dangers. There are even greater ones. These companies wield significant economic power, market penetration power and technology. Therefore they indirectly impact governance to a significant degree. All of them -except for Tencent- have joined forces in a consortium called **Partnership on AI** that has been accused of acting as a lobbying group that supports legislation -or lack thereof- that benefits their interests. This allows them to self-regulate without subjecting themselves to the scrutiny of shareholders. According to Markou, "Self-regulation is not an option. Governments and citizens cannot trust these companies just because they say they are committed to ethical practices, responsible practices or whatever word is in vogue and helps their business."

It is the right–and responsibility–of these companies to participate in the debate on promoting innovation or fighting hate speech, manipulation or piracy. But the line between having a voice and making the rules is becoming increasingly blurred. PD66 progressive member of the European Parliament Marietje Schaake writes about this in ***Bloomberg*** "I am entirely in favor of companies acting responsibly.... But when the most powerful tech companies take on the responsibility of global rule-making and cross-border governance... that is deeply problematic for democracy and the rule of law."

The risk that the private sector will capture the public interest is called privatized governance

Schaake warns about the risk that the private sector will catch the public's eye and that regulations are built with no transparency, accountability and human governance. That's what she calls the "privatization of governance". "Technology empires are making decisions about our lives and imposing their laws upon us. They want to connect all of us to control all of us", states in an article in *El País Retina* Renata Ávila, Human Rights lawyer and executive director of the Smart Citizen Foundation.
According to Markou, a hypothetical result of this is a future in which AI has consumed society. Computer systems have replaced judges and lawyers, and the rule of law has become the rule of technology. The legal system has become opaque and incomprehensible for everyone except those in control. In his essay Manufacturing An Artificial Intelligence Revolution, system biologist Yarden Katz, affiliated to the Harvard Medical School and the Berkman Klein Center for Internet & Society, talks about the false impression that current systems have exceeded human capabilities to the point that we think machines can handle many areas in a better way. These statements are based on a narrow and radically empirical view of human intelligence. The French intellectual Alakn Supiot dubbed this "governance by numbers."

That "numbers don't lie" is a dangerous assumption, says **David Weinberger**, the philosopher, technologist and researcher that we introduced in the previous chapter. According to Weinberger, machine learning learns from the data that we feed it, and this data often reflects existing prejudices. Plus, we decide which data we believe is relevant for solving a problem, which is another way in which human subjectivity and bias shape AI. We also decide which results are acceptable and what is considered fair.

Weinberger explains that although we use these systems when we believe their results are better than those a human could generate, sometimes that superiority is just a matter of greater speed and lower price. "When we use AI because it is more precise, we can say that machines are superior when making decisions. This assumes that we have independently decided to act upon their results. But this doesn't make those machines highly rational, except within a narrow definition of rationality.

> Machine learning systems find statistically relevant and/or useful patterns by correlations

Machine learning systems find statistically relevant and/or useful patterns by correlations, some of which might be direct and causal, and some of which might be signs and signals, whatever makes them better at what they do. Unlike Markou, Weinberger does not believe we are dehumanized by these machines any more than we are when we use computers to make calculations or write in notebooks to remember things. "I think they put humans in a different and more realistic place. They remind us that our minds are quite limited in the amount of information that we can manage. When we realize that, it humbles us. It's absurd to think that we are the epitome of living rationality and that we are destined to understand the order of the universe in detail. I don't think that's true." Markou agrees that human decision-making is very biased and deficient. He does not believe that the solution is to remove us from the equation but rather improve the nature and quality of our decision-making.

2.1.5
Impact on Privacy

When you go into a store, you connect to the WiFi. On page 15 of the terms and conditions agreement that you must accept, there is a small clause that says, "If you connect to our network, we can trace you down even after you've left the store." Without even reading it, you click "accept," and the company starts tracing you. They know you go to the gym, to the pub, to eat out. They know the doctors you visit at the hospital, how fast you fall asleep and how many hours you get each night. They have plenty of information about your lifestyle that they will use, at least, to try and sell you things.

That is the example that **Adam West**, Marketing Director at Satalia, uses. In the previous chapter, he spoke of AI applications in marketing and advertising. It is not a real case, but it could be one of the many cases that happen everyday. Scandals are constant: investigations that reveal illegal use of data, privacy violations and abuse as well as hefty fines for the two leading tech companies, Google and Facebook.

One of the most recent scandals: over 1,000 applications for Android (Google's mobile operating system) compile data even after users deny them permission to do so; they elude security systems in a variety of ways. Just like in West's example, they take advantage, among other things, of WiFi network locations. This has been disclosed in a study with global repercussions that was conducted by researchers from the Carlos III University in Madrid, the IMDEA network and the University of Berkeley.

Several organizations from different European countries have recently filed complaints against the online advertising industry, particularly Google, through their respective data protection agencies. They accuse companies of a "massive data leak" and noncompliance with the General Data Protection Regulation (GDPR)

They specifically denounce the fact that the industry keeps using real-time bidding to share personal data with third parties, without the users knowing or being able to control it.

As explained in this article from *El País Retina*, data from your online profile ends up in an ad server when you access any website. It then goes to an ad bidding and per-impression bidding system. Next, your data goes to an ad exchange system that matches data from the bidders' demands to the supply of ad slots. Finally, one of the bids will be the one to be served on the next visited website, which will receive new data to complete the information it already has on us.

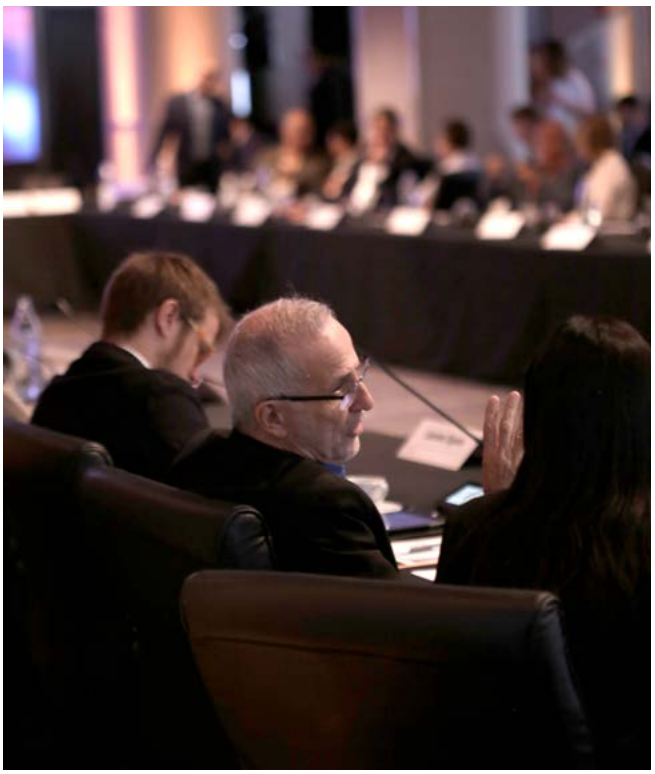

What type of information is sent? Johnny Ryan is one of the claimants and is Director of Policy and Relationships with the Industry of the Brave browser. He explains that third parties receive the following user data:

- Our browser history. What we see, watch or listen to.
- Content categories: anything that we have looked up for any reason, no matter how sensitive or private, is associated with us under categories such as incest, drug abuse, infertility, mental health, whether you vote for right-wing or left-wing parties...
- Our identifier (ID), individual pseudonym, which hides our identity but allows us to be recognized under that same ID in following browser sessions.
- This ID is associated with our profiles that advertising buyers have.
- Our geolocation.
- A description of our device.
- Our IP address (in some cases, depending on the real-time bidding system that is used).

Ryan points out the high likelihood that the different parties involved in the bidding process connect all the data they have at their disposal to create a more complete profile of you. One of the main consequences is that every single online person can be highly profiled. "The IP address is enough to obtain a single identifier that belongs to your device and can be connected to your name and last name," says Gemma Galdon, president of the Éticas Foundation. This organization conducts studies and awareness-raising campaigns on the impact on new technologies on society.

The Interactive Advertising Bureau (IAB) acknowledges that they have not adapted to GDPR yet, but they criticize the fact that "there's no distinction between companies that comply with it and those who don't. They put everyone in the same boat." IAB believes that the vision portrayed by claimant organizations "does not represent the entire sector."

Apple, which upholds privacy as its main concern, was among those companies that seemed to care. Recently, we learned about the company's lack of control regarding privacy in Apple Store applications. Technology columnist Geoffrey A. Fowler, together with privacy company Disconnect, published a study in *The Washington Post* revealing that around 5,400 companies traced his iPhone every week. They shared 1.5 gigabytes of data with third parties in just one month. Among the trackers that sent his information were applications such as Spotify, Microsoft OneDrive, Nike, Weather Channel and even The Washington Post app.

According to Fowler, the App Store requires developers to have clear, officially published privacy policies; but these policies don't necessarily provide protection. It also requires apps to ask for permission from users before collecting their data. However, many do not. The App Store also does not disclose the names of companies that track users or specify how they protect personal data. The more points of transfer of personal data there are, the harder it will be to hold companies accountable for misconduct.

These are obviously not the only ways in which data leaks occur. Apart from cell phones, there are plenty of devices that can be connected and share their information with third parties.

Next, we will turn to the Internet of Things (IoT). According to the *Cisco Visual Networking Index* (VNI), in 2022 there will be 3.9 trillion IoT mobile connections, four times more than in 2017. The *Ericsson Mobility Report* estimates there will be 4.1 trillion connections in 2024.

As the number of connected devices increases, so does the risk of cyberattack. Device use at an industrial level, at connected homes or on our own bodies (wearable devices) poses a threat. The wide range of devices that can be connected (TVs, thermostats, locks, alarms...) generates a large amount of potential access points for hackers, who could even penetrate the deepest layers of the systems.

Data scientist **Lauren Dyson** warns about this. "I think the greatest potential, and also the greatest challenges, lie at the intersection of big data and AI-driven, IoT-enabled optimization." She mentions use cases for these systems in urban infrastructure, where sensors were incorporated. These sensors facilitate the compilation of vast amounts of data in real time, and they can provide information about everything, from the dynamic optimization of the traffic flow to the ads shown in public transit. According to Dyson, the downside is that this raises important concerns on data security and privacy, particularly regarding individual information and personal identification. "We need solid, external oversight to guarantee responsible data management."

2.1.6

Related technical challenges

In the previous chapter, **Gary Ang**, Strategy Director at Temasek, discussed how social and regulatory pressure on the ethical use of data and AI plays -and will increasingly do so- an essential role in technological competition among countries; access to data is key to the development of these technologies. He also mentioned trust as a key motivator for users to share their data. The ethical use of data and trust are directly related. If data is used incorrectly, there will be no trust.

This is very clear to **Oliver Smith**, Strategy Director at Alpha Health. He previously explained their goals regarding their personal healthcare assistant. This area is particularly sensitive because of the type of data involved. What are the main ethical concerns and challenges when developing this type of application? He explains there are two central challenges. The first is maybe the most obvious: in order to offer a highly effective, personal healthcare assistant, the company needs to access a large amount of user data. "We don't think companies should have that kind of access unless they can build a relationship based on trust, and working in an ethical manner is an important component of that."

The second challenge has to do with the fact that personal healthcare assistants make recommendations and enable tools to determine what's 'good' or 'bad' for someone at a specific moment in time. "These recommendations are sometimes based on judgments about who would benefit from a product: is it present you or future you? These judgments show that we need to take into account our values," he states.

Smith is convinced that ethics need to exist in healthcare settings in order for them to be successful, because people have different goals in those settings. "We need to be able to understand what the right goals are and when. For instance, knowing when you should not have another drink or when you can say 'I've had a horrible day, I'm going to go grab a drink with my friends. It will make me feel good'. Understanding this is essentially an ethical and moral decision. But there is another ethical component: being perceived as an ethical person who can be trusted, and earn that trust."

To face these challenges, Alpha Health has developed a set of ethical promises that define their values. Smith says, "We'd like to contribute and help others that are joining us on this important issue. But we don't want to go ahead and set any goals until we make sure that we can achieve them." So far, they're testing how they can implement their ethical vision both in their processes and their products. On the one hand, they have started doing external audits of their work "to better understand what we need to improve." On the other hand, they are focusing on trying to find out which design templates work better, how to make privacy policies more comprehensible and how to develop intrinsically explainable algorithms.

The last issue, regarding explainability, is one of the main points. For Smith, the relationship between obtaining benefits from AI in general and this challenge is not fully understood. In healthcare and medicine, this disconnect happens precisely because their most advanced algorithms are too opaque, turning medical regulation into a curse. "Regulatory agencies obviously require consistent results, obtained under the same circumstances, for medication and devices; any deviation from these results must be explained and understood. However, this is impossible to do with the most advanced machine learning techniques, such as deep learning."

In these types of systems, algorithms are essentially black boxes; no human being can explain why a specific result is obtained or why it may vary depending on the circumstances. This lack of explainability is unacceptable for regulators, healthcare professionals and patients. As a result, medicine won't be able to fully benefit from AI. Fortunately, Alpha Health, among many other companies, is already trying to solve this by creating methods that explain deep learning algorithms. These efforts are still preliminary, and their development must involve all interested parties in order to guarantee the right degree of explainability. These efforts must also be encouraged and promoted by regulators.

2.2

Potential Solutions

In the previous section, we saw some of the negative effects, challenges and risks that developing and implementing AI and similar technologies can cause. Now, we will examine potential ways to address them.

2.2.1

Ethical Framework

AI's derivative and potentially illegal uses have resulted in endless initiatives, declarations, manifestos, principles, guidelines, analysis, etc. These focus on observing, measuring and analyzing the effects that AI systems might have on people, as well as measuring the potential negative impact and the preventive development of an ethical AI.

> There are initiatives that study the negative impact and the preventive development of ethical AI

Different institutions, governments, tech companies, consulting firms and various types of organizations have come on board. Large tech companies like Google and Facebook have created their codes of ethics and principles. Microsoft has created the AI Now Research Institute "to ensure that AI systems are sensitive and respond to the complex social domains where they are implemented." The British government, among others, has stated it wants to support the cause.

At a greater scale, the European Commission has several related initiatives. Recently, the EC has published multiple documents on this issue. Among them are their Ethics Guidelines for Trustworthy AI, which aim to provide guidance on the promotion and guarantee of an "ethical and solid" AI. The document's approach is based on the fundamental rights to identify ethical principles and the values that should be respected in the development, deployment and use of AI systems. These values include the following: respect of human autonomy, damage prevention, equity and explainability, special attention to situations that affect the most vulnerable groups, situations characterized by power or information asymmetries, acknowledgment of potential risks and the implementation of adequate measures to mitigate them.

The EC's document offers a number of recommendations to implement these measures, as well as indications to assess the reliability of AI-driven systems. The document could provide initial orientation, but putting into practice the codes of ethics and guidelines in every specific case is complicated. This difficulty can even be used as an excuse for not implementing them. The problem is, then, that the principles for AI's ethical development are just a mere declaration of intention. "For many companies, public efforts to address ethical concerns are just displays. There is excellent motivation to develop principles and guidelines, but they will be meaningless if companies cannot show the relevant work they're doing to support them," McEvoy states.

The researcher is aware that we are all quite new to this era of accelerated technological development, in which "we have never had to deal with moral consequences so urgently." The easiest way to do this, according to her, is to make promises and public statements that benefit the companies' image. Markou calls this the "ethics theater" in the industry of AI: an attempt to improve their reputation through ethics; this is also known as ethics washing. So far, it has backfired for Google in 2019, with the controversial dissolution of their Ethics Committee a week after its creation due to the identity of some of its members.

Christopher Markou believes that the ethics framework is a brainteaser. To comply with due diligence, it's not enough for companies to have a team that thinks about these things. They must take specific steps and there must be accountability. "If you do something wrong, there must be consequences," he states. According to Smith, this is an important discussion because it leads us to a new type of relationship (both contractual and in terms of rights); this is equivalent to the workers' movements in the 19th century that led to an improvement in working conditions. Rules and limits were put in place that resulted, in turn, in increased productivity.

For **Weinberger**, the most interesting question is how AI could be changing some of our ideas on morality. "The need to provide extremely explicit instructions to these machines leads us to formulate several models of impartiality." Weinberger claims that conversations on different types of impartiality and the results that we want from AI applications are leading us to talk less about moral principles and more about values. These values are often not consistent with each other, so they need to be compensated. "We tend to classify conversations on value compensation as a political issue, instead of a moral one. AI can lead us to think that morality is essentially political."

He points out five essential considerations to ensure ethical AI development. They are as follows:

1. Facing the biases, "the original sin of AI." It learns from data that inevitably reflects our existing biases. Weinberger states that they're working hard on ways to identify and minimize these biases. But as long as society is biased, it's possible that biases are present to some degree.
2. People at the core. "The communities that develop these technologies increasingly acknowledge that the groups of people affected by them must participate in the entire process. Their voices need to be heard, directly and indirectly."
3. Secondary and tertiary effects of these systems. "They might work perfectly according to our specifications and still have damaging effects on the community," Weinberger states. For instance, the system could be successful when making bus routes more efficient, but that might cause the non-inclusion of the outskirts or isolated areas during rush hours.
4. Transparency. "We shouldn't assume that transparency is the right answer to all moral problems cause by AI. And, if it were, we must be specific about the need to be transparent." Data? Processes to inspect clean or certify data? A diverse group of citizens that have been involved in the process? What are the desired results? What is the system's performance regarding these results? Are there problem resolution processes under way?
5. Public discussion. According to Weinberger, the desired outcome of AI systems should be discussed, decided and supervised by the public. Do we want autonomous vehicle makers to freely decide which values they want to prioritize?

Aligned with the philosopher's vision, **McEvoy** highlights the importance of considering the types of tasks that must be automatic and need AI, and the ones that must be carried out by humans. Like Weinberger, he also talks about matching products to people's interests. He believes that a priority for the business must be consulting with a wide, diverse range of interested parties and sources in order to ensure that the most pressing concerns are addressed.

We must assess what tasks should be AI or human-driven

McEvoy insists that the use of AI "must be open to a true, public debate" before the development of the technology. "As a society, it is mandatory that we have the opportunity and capacity to assess and reassess the moral action parameters of AI. Sometimes there will be some obvious cases of misuse and unsatisfactory results, but other times we'll have to deeply think in order to anticipate the potential damage that emerging technologies could cause." He also adds that ethical disagreements don't mean there is no 'right' ethical solution to reach if we have the right approach.

2.2.2 Education

Just like the speakers at the FTF, the EC includes in their guidelines the need for public debate and the involvement of interested parties in the entire lifecycle of AI systems. Yet they add one more element to the equation: education. "In order to guarantee AI's reliability, it is necessary to create and maintain an ethical culture and mentality through public debate, education and practical learning, well beyond the development of a set of rules."

The EC discusses promoting training and education in a way that all interested parties know reliable AI and receive information on the topics. Entrepreneur **Andrés Torrubia** agrees. "AI is a technology that is going to impact virtually all social industries and sectors. Software has transformed plenty of industries in the last 20 years, but AI is going to be a tsunami with even greater power."

■ Software has transformed plenty of industries in the last 20 years, but AI is going to be a tsunami with even greater power

He encourages people to learn about AI for two main reasons. First, because of social responsibility. "The technological industry is ruled by companies that, so far, haven't really had to worry about the indirect impact of their actions. If a company in another industry, for instance, pollutes a river while manufacturing a product, we don't think it's acceptable. As a society, we have implemented mechanisms to minimize what we call external costs (the cost of a product of service that ends up being the responsibility of the society, not the producer)."

The software industry, and AI to a greater extent, are immature in terms of dealing with external costs or externalities. "We are starting to see quite frequently in technology abuse (addiction, depression, loss of attention, lack of information, large scale manipulation, etc.) that AI amplifies both the good and the bad."

Torrubia thinks he is a tech optimist, but a political pessimist. "Lots of people talk about ethics, biases, governance and accountability regarding AI, but very few know what they are actually talking about, with a few exceptions." That is why he would like to see more diversity in discussions about AI, which would require other professionals (lawyers, philosophers, economists, etc.) to undergo significant training. "I know it's hard, but the alternative of having them talk, and even draft legislation, without enough knowledge on the topic is unacceptable."

The second reason why he insists on encouraging others to learn about AI has to do with the modernization of the Spanish industrial structure. "Our IBEX 35 looks more like the IBEX in the 90s (construction and energy companies, banks, etc.) than the transforming driver in the 21st century: software-driven technology *development* companies that are now AI-driven. These companies create quality jobs, and there is a lack of talent." That is why we need more technical talent.

Torrubia, who has a degree in engineering, studied AI and machine learning through an online learning platform. "This type of platform has three main advantages: first,

the speed at which the content is adapted to the market; second, the availability of professors and excellent material well beyond what you could find in the area where you live; and third, flexibility to learn at your own pace."

Executives also need to educate themselves on these technologies. At least that is what **Michael Li**, founder and president of The Data Incubator, believes. That was his goal when he created a company focused on training and staffing in the data science industry. "We help companies develop a data-driven economy by hiring data scientists and training not only their technical teams but also their management teams and other professionals on issues regarding the impact of big data and AI on their business."

Michael Li
Founder and CEO
at The Data Incubator

Companies like Google, Amazon and Netflix know how to capture the value of data; they lead the race. That is why being data-savvy is mandatory in order to be competitive, according to Li–both for companies and their employees. "Before companies can monetize data, they first need to understand it. Many of our clients come to us after they find out middle management is crushing data-driven initiatives because they don't understand, prioritize or value them."

Duration: 09:30 min

Before companies can monetize data, they first need to understand it

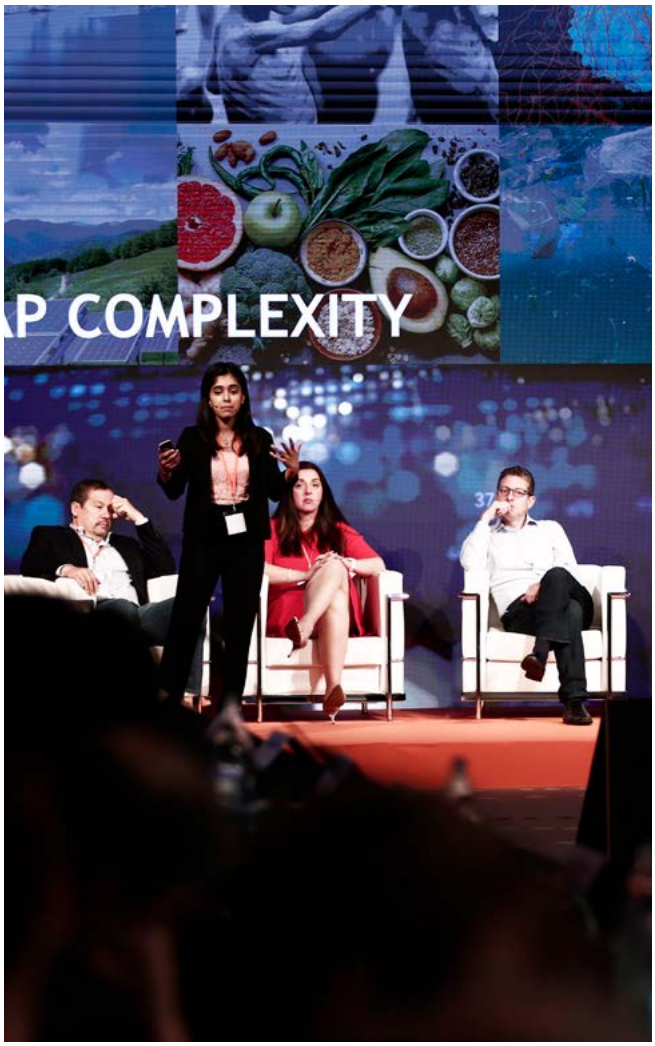

Li points out that The Data Incubator has observed an increase in demand for their *Business of Data Science* course as more companies and professionals realize they need it. This course teaches about the potential impact of AI on a business, project prioritization, end-to-end management... "There are important issues that executives and managers need to understand if they want to lead a data, AI-driven company."

As for the technical component, Li says that the goal is to teach experts what he calls "data analytics 1.0." In the previous technological wave, these experts worked with quite small, well-formulated data amounts. Now, they can work on data analytics 2.0 with the new tools and skills required to make AI work.

Li believes that thinking about automation and unemployment as a war between humans and AI, between people and machines, is a mistake. "The future is going to be a combination of both; we need to let each of them work on those areas where they have a comparative advantage, economically speaking." According to him, computers and AI will probably be better at things like

persistence and statistics. Humans will be there to provide context and empathy.

Data science training might also be differential in emerging regions. In the case of Africa, technologist and entrepreneur **Juliana Rotich**, associated with the MIT Media Lab, highlights the importance of investing in AI's local capabilities, implementing training plans on data science in schools and universities. "Training plans in most African universities must meet the country's needs. Africa is trying to catch up."

Juliana Rotich
Technologist, MIT Media Lab
Director's Fellow

She points out that data science training can help channel a flow of solutions for problems in Africa, but not all problems would benefit from it. Therefore, it would also be necessary to determine the aspects that could contribute value and those that may cause harm. The motivations and goals of the development of these technologies must respond to social needs, not just to stakeholders' demands. She concludes, "Innovation ecosystems must be strengthened through continuous investment in new companies and also by improving African governments' capacity to better understand the benefits and risks of using AI."

The purpose and goals of developing AI must connect with social needs

2.2.3
Digital Rights

The risks mentioned by Juliana Rotich are related, among other things, to the violation of human rights. Several initiatives have emerged within the context of the conversation about the impact of AI and technology in general. These initiatives demand the consideration of digital rights as an extension of human rights. That is the goal of the RightsCon Conference, organized every year since 2011 by AccessNow. It is also the goal of the **Cities Coalition for Digital Rights' Declaration.** The document's third point focuses on "transparency, accountability, and non-discrimination of data." It states that everyone should have access to understandable and accurate information about the technological, algorithmic and Artificial Intelligence systems that impact their lives and the ability to question and change unfair, biased or discriminatory systems.

The Ranking of Digital Rights (RDR) goes well beyond that, producing a Corporate Responsibility Index to bring forward corporate commitments, policies and practices that affect users' freedom of expression and privacy. The *2019 RDR Corporate Accountability Index* evaluated 24 companies across 35 indicators, examining different aspects of their governance, policies and practices. The RDR index measures the potential adverse impact of companies' trade operations on users' human rights. This year the RDR went further and addressed the companies' due diligence efforts regarding the use of AI-driven automated tools for decision-making and their specific advertising policies and practices.

According to the report, only three companies – Deutsche Telekom, Microsoft and Telefónica –were assessing the human rights risks associated with AI use. Telefónica was the only one who showed that such assessment is part of their official, continuous process of estimating their impact on human rights. The company shot ahead of every other telecommunications company in the RDR Index because, according to the document, the company disclosed significantly more than its peers about policies affecting freedom of expression and privacy. It also made more improvements than all other companies "by a landslide."

Based on the results, the RDR index recommends that companies evaluate their impact on human rights; strengthen direct oversight of risks regarding users' security, privacy and freedom of expression; commit to evaluating third parties based on international human rights standards; establish effective and accessible mechanisms of appeals and solutions and get involved with the affected, interested parties in order to create new processes that identify risks and mitigate damages.

What is the role of governments in all this? For **Jade Leung**, the challenge of AI's governance is "a problem that can be solved with a lot of effort." First, she thinks we should focus on history, since it is not the first time that a technology has become politicized or strategic. "Nuclear technology, biotechnology, cryptography... we've had rights and wrongs in all these cases, and we should learn from them." Leung believes it is essential to take those analogies into account.

Secondly, Leung claims that we should strengthen the impact of institutions that people trust the most. Part of her work with GovAi entailed conducting a survey with US residents on several AI-related issues. The conclusion is that people trust Facebook very little but trust researchers and the armed forces very much. "This is the kind of data we should study in order to help research centers influence governance, and this also drives trust in institutions, which is essential for me."

Leung says the ball is in the private sector's court. "It must develop an essential role in terms of governance. It is frustrating to see what large tech companies are considered nowadays; they're essentially the bad guys. I think we need to understand that government legislation has never kept up with the pace of technological governance problems, nor will it. These technologies are developed in the private sector, and that's where we will find the technical experts that will help us make them safer, more resistant and transparent, as long as they have the right incentives in place."

A potential approach to technological governance is to bring private sector companies to the front and let them lead, but ask them to do so responsibly for everyone's sake. Markou has already mentioned that he does not think this is possible. He compares this case with the 2008 financial crisis. Markou explains, "We allowed an industry with no regulation and no control to drive the global economy into bankruptcy. That cannot happen with AI.

As for human rights, the RDR Index states that governments should not only protect them but that they can perform an essential role in ensuring that companies implement adequate oversight and management of the potential human rights risks they may cause. Its recommendations: a robust data protection law (such as the European GDPR), implementing and demanding transparency and ensuring access to effective mechanisms of appeals and solutions, among others.

2.2.4

Data Ownership

Another issue under discussion is whether companies should pay users for their data, given the fact that data is the basis for the business model that many of the major tech companies use. Economics Reporter of *The New York Times* Eduardo Porter believes it is worth exploring this model based on data ownership. He says he is unsure about the value of marginal data for current business models, driven by advertising. However, he adds that they might be much more valuable in a strongly AI-driven economy.

> People need to better understand how their data is used and need to have better control over the process

Regardless of whether companies pay for data or not, Porter reinforces what other FTF experts mentioned: people need to better understand how their data is used and need to have better control over the process. Additionally, the journalist mentions that if AI drives an increase in productivity, wealth will increase as well, leading to social and political challenges. One of which is the challenge of ensuring that technology is deployed to improve productivity, not to leverage on tax advantages. Another challenge is that of ensuring that new wealth does not end up in the hands of a bunch of plutocrats, and instead in the hands of most of the population.

Nuria Oliver states there is a social cost that companies are not paying for; doing so could be a good way to establish balance. The questions is how, given the fact that some of these companies have been accused, precisely, of tax evasion. The French parliament has approved a 3% tax on the income that companies such as Google and Facebook obtain within the country, claiming that they currently take advantage of global fiscal loopholes to reduce their contribution to the Internal Revenue Service.

There are other options on the table, such as granting legal status to the machines that replace human labor. Markou believes this idea is "silly, dangerous and degrading. It doesn't mean that we cannot create some sort of mechanism for accountability, but we do not need to use personality as the starting point to build it. It'd be like pouring new wine into old wineskins."

2.2.5

Artificial Intelligence for Social Good

In the previous chapter, Adam West discussed how companies can use data to sell us more products and services. He now gives us an example about how the very same data could predict the likelihood that someone develops a heart condition. It's just an example of how focusing on social benefits instead of merely economic gains might completely change AI's impact.

This is where AI for good or AI for social good fits in. "AI has the potential to help address some of the greatest social challenges in the world." That's how ***Applying artificial intelligence for social good***, by McKinsey, opens. In this paper, the consulting firm compiles approximately 160 use cases suggesting that existing capabilities could help achieve the United Nations' Sustainable Development Goals (SDGs), "therefore benefitting hundreds of millions of people both in developed and developing countries."

Analyzed use cases range from diagnosing cancer to helping the blind navigate their surroundings, identifying online sexual exploitation victims and contributing to emergency relief efforts (such as the floods that Hurricane Harvey caused in 2017). However, as expected, AI is not a magic wand. "It's just part of a much broader set of tools and measures. So far, issues like data access and the lack of AI-trained talent limit its application for social good."

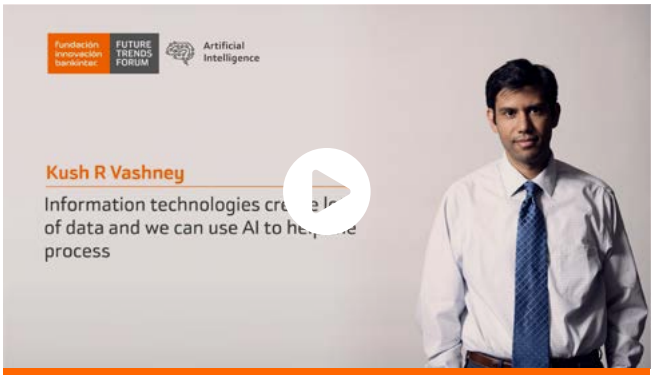

Duration: 00:28 min

Kush Varshney works from a technological development standpoint in order to contribute to the SDGs through AI and data science. He is co-director of the IBM program IBM Science for Social Good. He states, "We believe that the proliferation of technologies that are transforming our world such as AI, and efforts to improve social welfare and to address global challenges can be integrated within research and technological development and addressed in a sustainable, scalable way."

According to Varshney, in order to ensure that technology for social good has a real, measurable impact, we need to go beyond the current practice of creating tailored solutions to using an open platform paradigm that provides essential capabilities that social change organizations with similar needs can use in their projects. He believes that this requires a new philanthropic model, bringing together technologists, foundations and organizations for social change.

03

Future Scenarios and Potential of Artificial Intelligence

We discussed AI applications, threats, risks, challenges and possible solution pathways. Now **we look to the future to ask ourselves,** "What does AI have in store for us?" In this chapter, we analyze the future from two standpoints. **First, the technical aspects. Then, we examine AI's impact on life in the future and take a look at the different scenarios that could take place.** We will take into account the ideas that we presented earlier as well as the predictions made by experts at the FTF.

3.1

Technical Potential

How far will AI advance technically? There are opposing answers to this question. On the one hand, tech super optimists refer to artificial super-intelligence that could become more powerful than anything this planet has ever seen before and will pose a challenge to the very existence of the human species. Members of this camp subscribe to the beliefs of the philosopher Nick Bostrom. He is the founding director of the Future of Humanity Institute at Oxford University (FHI) and the author of *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press, 2014). Comments from big names like Stephen Hawking and Elon Musk, among others, have given traction to this branch of thinking.

> Technological optimists believe AI could become the most powerful technology on Earth

Skeptics occupy the other camp, led by another philosopher from the University of Oxford, professor Luciano Floridi. He is director of the Digital Ethics Lab. He wrote in a 2016

essay published in *Aeon*, "After so much talking about the risks of ultraintelligent machines, it is time to turn on the light, stop worrying about sci-fi scenarios…"

One such scenario concerns the fear that AI will progress to the point where it can improve itself and become difficult to control. Yoshua Bengio rules out this idea, saying, "That's not how AI is built. Machine learning involves a slow and intricate process for acquiring information from millions of examples. "Yes, the machine improves itself, but very, very slowly. And it does so in highly specialized ways. And the types of algorithms used are nothing like little viruses that can program themselves."

Floridi agrees that "the reality of AI is much more mundane than we think." That is why he criticizes movements that preach about technological singularity, which is the idea that artificially intelligent machines, cognitively superior biological intelligence–or both–will surpass human ability. According to this philosopher and engineer, what he calls the 'Singularitian Church" is "irresponsible" in "distracting" people by spreading their ideas.

Floridi also criticizes the opposing camp, the "Atheist Church," for getting wrapped up in a "nonsensical" argument with the "Singularitans." As Turing wrote in his famous 1950 article 'Can a machine think?', the very question is so useless that it does not even merit discussion. And, according to Floridi, both sides are wrong. "AI that truly thinks, generally and robustly, isn't logically impossible, but it is completely implausible. It's just highly unlikely. We have no idea how to design it, mostly because we know very little about how our own intelligence and brains work."

This is why we should not lose sleep over the possible rise of ultraintelligence. "You can disregard any catastrophic views about AI." "What really matters is that the growing presence of more-intelligent technologies is having an enormous impact on how we understand ourselves, the world and our interactions. It does not matter if machines are conscious, intelligent or able to know things like we do. They're not. No conscious entity will come from a Turing machine."

According to Floridi, the key point is that there are a growing number of technologies that can handle more tasks, better than humans, and that fact poses a threat to our anthropocentrism. Floridi dubs this process "The Fourth Revolution in our Self-Understanding." "We are not at the center of the universe (Copernicus), of the biological kingdom (Darwin), or of the realm of rationality (Freud)." And now we are no longer at the center of the infosphere (Turing). Instead we share this realm with digital technologies. Floridi defines these as "ordinary artifacts that outperform us in ever more tasks, despite being no cleverer than a toaster." "Their abilities are humbling and make us reevaluate our unique intelligence and special role in the universe."

There are increasingly more technologies that perform better than human beings

Our expert **David Weinberg**er agrees with this stance and also fails to believe that achieving superintelligence is the ultimate goal of AI. "I don't believe that current AI systems are intelligent or conscious. Nor do I think they are evolving to reach either of these two states. Computers process symbols. Brains, like stomachs, do not."

Ramón López de Mántaras also warns –as we mentioned in Chapter 1– against the risks of falling for "unrealistic promises of AI." However, he believes that perhaps, in a few years, these systems will be able to comprehend to a degree. But, for now, the first challenge is that of developing multi-area systems that can learn multiple tasks and avoid the nagging challenge of so-called "catastrophic forgetting."

Catastrophic forgetting entails teaching new skills to a system without loosing the old ones

This concept, which we briefly introduced in Chapter 2, refers to the difficulty involved in teaching the system new skills for new tasks, without losing the skills that the system learned before. For example, if a network is first trained to distinguish between pictures of arms and pictures of legs and is then re-trained to distinguish between arms and hands, the first set of information will be overwritten, and the system will lose the first skill. Mántaras believes that in 10 years we will have broader versions of AI, which will be able to overcome this obstacle.

Researches and technologists are working to overcome a different AI obstacle, Moravec's paradox. According to Moravec and his colleagues, machines can do things that we find difficult (playing chess, for example) but, unlike humans, they cannot learn the psychomotor or perception skills that even an infant possesses. That is why there are researchers and technologists, participating in initiatives like MIT's **Quest for Intelligence**, who are trying to understand how children can do what they do so effortlessly.

Moravec's paradox states that AI has intellectual skills, but lacks psychomotor or perceptive skills

The inaugural director of this ambitious project, **Antonio Torralba**, is in fact a Spaniard. He explains that one of his lines of research is to understand child development during the first eighteen months of life. "When it comes to interacting with the world, at just one year old, infants are already experts at many things: perceiving three-dimensional forms, precisely manipulating objects, agility... This is something a robot can't do or will really struggle to achieve." That is why he has taken it upon himself to understand the nuances of child development. This is the same process by which "Japanese babies can distinguish between the letter 'R' and the letter 'L,' while adults cannot."

Torralba sees this line of research as a *moonshot*, when people seek to achieve something that, at first, seems impossible, at least in the short term. However, there are other fields in which scientists believe AI will notably improve in the near future. Mántaras believes machine learning will be better at combining a growing number of techniques. We will depend less on data, and AI will be more explainable, safer with respect to privacy, and sustainable. AI will also have a greater awareness and understanding of cultural context. In order to reach this goal, people must remain part of the process. Mántaras believes not only it is the right thing to do, but people also play an essential role. "Collaboration between man and machine is key."

Explainability is a challenge for AI, as it requires explaining each outcome

Explainability, both in terms of understanding and trusting algorithms, is a key factor in this partnership. Oliver Smith, Strategy Director at Telefónica Alpha, discussed this topic in the last chapter. Smith explains that when Alpha Health worked with the UK National Health Service (NHS), the latter insisted that the former not use deep learning based on neural networks. Alpha had to use simpler technology, decision trees that NHS staff could understand.

"It worked at the time, but it won't always work," Smith emphasizes. That is why his team is researching how to convert neural networks into decision trees to avoid the rejection of more advanced technology. "We've turned to, TREPAN, a very promising technique from the 90s." The problem is that this only works with unidirectional neural networks, as opposed to networks that allow for arbitrary connections or loop structures (recurrent neural networks).

IBM has taken a different track in an effort to improve the field of AI explainability. As recently as August 2019, the company released a set of open-source programming resources called "AI Explainability 360." It contains a set of eight algorithms, as well as guides/tutorials to help people understand when they should use a specific algorithm and identify the target audience of the explanation.

Tiernan Ray writes in an article on *ZDnet* that this field will require greater research and exploration. "Getting a toolkit from IBM is just the beginning of the hard work." Smith believes that we can solve this problem by 2030 and that surpassing this hurdle will catapult AI's development as well as its application in the field of health. He warns, "But that will mean nothing if we fail to remember the key role that humans play in all of this."

Other FTF experts spoke at length about their hopes for other aspects of technological development related to AI. Read on.

3.1.1

Data Science

What is data science? **Nuria Oliver** explains it for us. She is a member of many other associations, including RAI and the Data-Pop Alliance. Oliver has been working on computational models of individual and aggregated human behaviors for 25 years. She explains, "Data science is a discipline that has emerged in the last ten to fifteen years out of the intersection of math and statistics, computing, business management and domain knowledge." She says that data scientists are in such high demand that in August 2018 there were over 151,000 unfilled positions in the US (Linkedin report).

Duration: 10:09 min

Oliver says a new field within data science, computational social science, leverages large-scale, aggregated behavioral data in order to verify social science theories. For the last ten years, Oliver has also been working in an area of computational social science: data science for social good. The goal of this area is to leverage the analysis and processing of massive data sets in order to make better decisions, which "could save millions of lives," in fields like public healthcare, financial inclusion, climate change or emergency relief in natural disasters.

As the name suggests, data science requires data and science. Regarding the former, Oliver claims that reports estimate that every year we generate more data than in the previous 5,000 years combined. "The capacity to interpret it, additionally, is an extremely valuable asset in today's economy." According to a European Commission report, the European data economy will be worth over 700 billion next year, "as long as favorable policy and legislation is implemented in time and investment in ICT is promoted."

Data science also needs science. It's good to have the data, but if you don't know what to do with it, the data becomes digital waste. We need to know how to interpret data– large volumes of non-structured data– which are invisible and incomprehensible without science and its disciplines, particularly data-driven machine learning.

"We've been fortunate in that regard because in the last eight years, we've seen a boom in several factors that contributed to great progress in data-driven machine learning; we've also had large volumes of data available. Now, we have high-performance, affordable computers and a very sophisticated machine learning architecture

that allows us to ingest data and understand it." Oliver's explanation echoes aspects of Chapter 1.

According to Oliver, these are very relevant factors. "We are seeing a democratization of data-driven machine learning. There are plenty of tools and libraries that create such a level of abstraction that there's no need to program all the algorithm equations in order to train the machine. These tools and libraries are essentially Lego bricks", she says. But the expert also warns that we need to know what we are doing. "Learning something that makes no sense can be dangerous," she states.

Oliver believes that progress over the last eight years has resulted in amazing achievements, such as Google's AlphaGo victory against the Go world champion. However, algorithms also impact our daily lives, as previously explained. "Data-driven algorithms are going to have an increasingly significant impact on our lives, they already do. I would say, then, that data science is not just the combination of data and science; we need much more than that. Data science is multidisciplinary: we need domain knowledge, multidisciplinary teams and ethics."

Data Science is working on six areas of improvement

Oliver is convinced that, unless all these pillars are included, data science as a whole will not be good for humanity. She lists six areas of improvement that the computing and data science community is working on:

1. Computational privacy violation
2. Biased social exclusion or discrimination
3. Asymmetry in computational skills–plenty of data, but mostly in the hands of private actors
4. Opacity and lack of transparency
5. Veracity
6. Ethics

So, what is the future of data science? Oliver believes that, inevitably, it will address the inclusion of FATEN algorithms. F stands for *fairness*, focusing on cooperation and non-discrimination; A, for *autonomy, accountability* and *augmented*, preserving the human value of sovereignty, clear accountability and models that augment, not replace, human intelligence; T for *trust* and *transparency;* E, for *education* and *beneficence*, and N, for *non-maleficence*, minimizing the negative side effects, ensuring a minimum level of reliability, security, reproducibility, caution... always protecting people's privacy.

"Not all innovation mean progress, and we need to focus on progress, not on innovation in and of itself," Oliver states. We need to make sure that we are sustainable, diverse and truthful. We need to ensure that education is available at every level: for children, teenagers, professionals, politicians and, above all, for people who make decisions. That is because "they often decide on issues they don't really understand." To sum it up, the future of data science means we need "to trust algorithms, always focus them on people and invest in education."

3.1.2

Natural Language Comprehension

The field of computational linguistics–the study of language from a computational perspective–is over 70 years old and still has a long way to go. From the first translations done by a machine to the current boom in personal assistants, chatbots, etc., the increase in the number of natural language processing and language analysis tools in the 2000s has resulted in an obvious evolution of this field, whose limitations are still blurred.

Pilar Manchón
Vice President of AI
at Roku Inc.

Pilar Manchón, studies, develops and applies these technologies. She is a Spanish philologist, with a passion for computational linguistics. In 2013, she sold her pioneering chatbot company, Indisys, to Intel, and became Head of Cognitive Interfaces at Amazon. She now works as Head of AI at Roku, a TV streaming content platform and player.

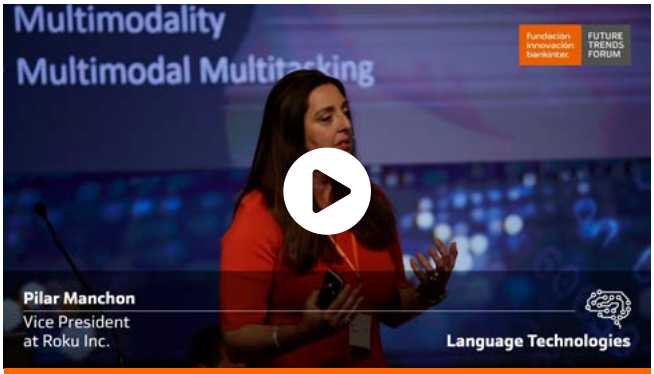

Duration: 09:27 min

Natural language enables machine-to-human communication

Manchón describes computational linguistics tools as "a package of technologies that have come together under the same category, but that are actually very different; they have different requirements in terms of performance, data, etc." Among these technologies, we can find:

- **Voice recognition** is the transcription of audio content, someone's voice and what they say, into words. Then we have natural language comprehension, which involves language comprehension and correlating what we see and what we believe. Manchón believes that these systems exercise "competence without comprehension," a statement that is widely shared and questioned in the scientific community. Philosopher of the mind and cognitive scientist Daniel Dennett proposed it; it essentially means that, at a functional level, a system can reach a performance level (competence) that would be attributed to comprehension in human settings (intelligence); however, the system would not comprehend what it is doing.
- **Dialogue management** is essential to having a conversation. "One needs to understand when it's their turn to speak, how to associate what someone just said and the information shared with what they have said and where the conversation is going depending on what both parties know."
- **Natural language generation** allows the machine to communicate in a way that a human can understand, by having an idea of what the interaction will be about and how it will take place. It's about generating a task, a message and information... anything—in natural language. "To put it mildly, it's about building a grammatically correct sentence."
- **Speech synthesis** makes it possible to generate natural language in the form of a voice.

Virtual assistants who can hear, see and feel will be developed

And what is next? "Multimodality and multimodal multitasking. We are going to leave behind voice and text; think about a more complete vision of a virtual assistant that can see, feel and know other things apart from what you tell them. If I say something like 'Show me that' or 'Open up that window over there', the virtual assistant would literally need to know what I am pointing at in order to understand what I'm saying." So, apart from language, there's the interaction among all contextual modalities that we need to put in common in order to know what the real intention is, depending on the context in which these modalities take place.

"Apart from all that functional intelligence, there's a social and emotional component. That is because humans are very needy creatures. In order to trust a semi-intelligent entity that performs a role in our lives, we need to have some sort of emotional connection. If we don't have it, we'll still develop it." As an example, she mentions María, a virtual assistant that Indisys developed for the Community of Andalucía's healthcare system. "Users said they were comfortable when consulting the machine about sexual health, because they don't feel judged; however, they talked about it as if it were a person."

AI will detect personality through language

Another component, according to the expert, is detecting personality. This is the ability to detect the type of personality that someone has through language. This has to do with the idea of trust: who you are, where you are from, what you know... Inferred intelligence not only has to do with that you say. It involves several sources of information and adding one, two or three more things to what you're actually conveying in a specific message; it's also about how to put all the information together in order to take one more step forward in a smart way. And then there is synchronization, which has to do more with coherent behavior, so the user's expectations about the interaction are not altered.

In the future, the field of language comprehension will put together all these pieces and make them match in a fair, ethical way, avoiding negative effects, as pointed out in the previous chapter. The development of these technologies has–occasionally sensitive–technical needs. Pilar Manchon states that "Virtual assistants, for instance, are present in all aspects of our lives. Alexa, Google Home, Cortana, Siri, etc. They are at home, in the car, at the office... They could be everywhere, so they know a lot about us."

Data ownership, data management, who owns those data, the amount of data we're sharing, if we're sharing more than we actually want to, how you say what you say, when you say it, who you say it to, which tone you use to say it... "When we ask Alexa for a recommendation, will it tell us what's best for us or for the companies that paid to be there? When a virtual assistant offers recommendations, who are they really helping?" Since machine learning learns from observation, how can we explain to the algorithm that we don't like what we observed, that we don't want it or that it is incorrect?

Manchón states that some companies use our conversations with telecommunications companies, for instance, to determine our profile: what type of people we are, if they can offer us their products and services, if they can sell you something or do cross-selling, if you get angry easily... "They create a profile about you, and you don't know where that profile goes. It could affect a job application, a mortgage, your university education or your insurance, even your financial decisions. Therefore, companies, entities and services that you cannot see judge you as a user based on information you cannot control. They're not qualified to do so, they don't have permission and cannot validate that information."

In addition to this, these systems can be used for noble purposes, such as detecting online sexual predatory behavior, particularly with children, or detect bullying and how people are manipulated through social media. Linguistic technologies can also be applied to document analysis and classification and help in the field of education, making it more accessible and easier for those that are now not able to study. Manchón thinks it is essential to "extract the complexity of artificial intelligence through language and empower others, so they can use it and grow better and faster."

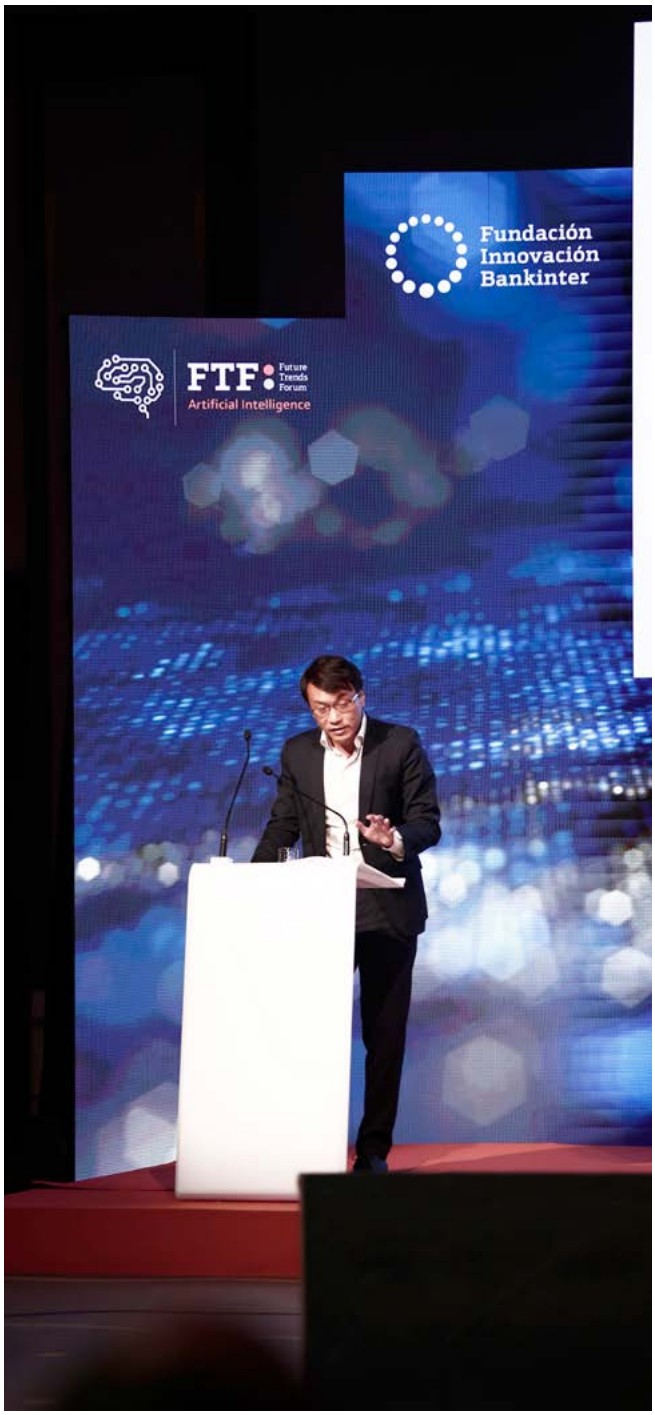

3.1.3

Autonomous Vehicles

Autonomous vehicles are no longer science fiction. There are real prototypes: ground vehicles, aircraft, trucks, bikes, drones... The main automobile and aerospace manufacturers, newer companies such as Uber or Didi, major tech companies like Google–and obviously plenty of universities and research centers–include autonomous vehicles in their R+D plans. According to estimates by Allied Market Research, this new, 54-billion-dollar market will be worth 556 billion dollars in 2036.

> The autonomous car already offers pilot services in freight, package and human transportation

There are several tests and pilots under way with these types of vehicles, for the transportation of goods, packages, people, and even rescue transportation. The big question remains, when will these vehicles be ready, and when will the infrastructure necessary for their functioning be fully developed and deployed?

Raúl Rojas is a professor and researcher who attended the FTF. He introduced us, in Chapter 1, to state-of-the-art of these vehicles and, in Chapter 2 he presented their ethical implications. According to Rojas, autonomous vehicles could take up to ten years to be ready, the time needed to see the changes in paradigms driven by a technological leap. "The transition between servers to microcomputers promoted the introduction of integrated circuits; microprocessors motivated the switch from minicomputers to personal computers. Then, we got the chips that allowed us to implement wireless networks. We have been talking about the IoT since 2010, and the question is, what can we expect in the next ten years?"

In his opinion, the next paradigm will be on-demand intelligence. "There will be smart devices that will have artificial intelligence, not on the cloud, but in a chip, as Lewis explained. So the intelligence is going to be there, in the devices themselves. That will create plenty of new applications, it will be the Internet of Smart Things."

Rojas assures that the process is under way and believes that the integration of telecommunications in cars, which according to him will begin in 2020, is an important step. "Next year, most of the large automotive companies–like Mercedes, Volkswagen, and Volvo–are going to bring to market what they call 'high-level automated driving' for highways. You just press a button, and the car will stay in its lane." he says. But this is just on the highway; the technology cannot be applied to cities yet.

> To develop the autonomous car we require more accurate maps of cities

One of the main challenges for the implementation of autonomous vehicles in cities concerns the availability of precise city maps. There are several companies working on this, mapping in 3D and therefore obtaining much more information than through Google Maps. "We also need more precise and affordable sensors than the ones we have nowadays, high-speed telecommunication between cars and better behavioral predictions", he adds.

So, what can we expect from a technical point of view? "High-level automated driving is on the agenda of most companies, but we still need millions of miles of experience using these systems. I think experience on the highway will be the basis for complete automation in some areas of large cities, which might happen by 2030."

In terms of the economics, he points out that there needs to be some justification. "Waymo says they want to be a taxi company; they want to be just like Uber, but with computers. But that's the same thing Uber and Lyft want." And then there is social acceptance, which Rojas believes to be extremely impotent. No matter how much autonomous vehicles (AV) are promoted as safer alternatives than human-driven vehicles, people will not use them unless they trust them. According to a survey by AAA, women feel less comfortable (79%) than men (62%). Likewise, an MIT study points out that 53% of women (compared to 32% of men) would prefer to have a driving assistant instead of a completely autonomous vehicle.

3.1.4

Edge Computing

We live in the era of cloud computing. We not only use our computers, tablets and cell phones as storage devices, we also use them to access online storage and file-sharing

services such as Google Drive, Dropbox, One Drive or Slack. People and companies all over the world trust the infrastructure of a few providers, mainly Google, IBM, Amazon and Microsoft, although there are more. In this environment, we are starting to hear about perimeter or edge computing, which is computing that happens on the data source or on the edge of it, instead of trusting the cloud to do all the work. It does not mean that the cloud is gone; instead, it essentially gets closer to us.

How does this relate to the future of AI and, more specifically, of machine learning? According to **Anthony Lewis**, vice-president and Head of Artificial Intelligence in the Emerging Computing Lab for disruptive technologies at Hewlett-Packard, "Machine learning looks for patterns in the data, and its biggest problem is how to provide an answer to data we have never seen before." He is referring to the need to generalize well beyond what has been observed in a data pattern. For instance, an algorithm that has been exposed to pictures of cats could identify a cat it hasn't seen before.

Anthony Lewis
Head of Artificial Intelligence and Emerging Compute Lab at HP Labs

Machine learning looks for patterns in the data, and its biggest problem is how to provide an answer to data we have never seen before

For this to happen, there needs to be some sort of representation or understanding of the data that we receive. Lewis used the following example: if we have a picture with five megapixels, there could be a wide range of combinations on how to map those five million pixels to have "dog" or "cat" for an answer. It is some sort of internal representation that can be done manually or it can be learned. According to the expert, this is one of the main milestones of machine learning: the use of neural networks to learn that internal representation. "It entails the involvement of all economic sectors, and that somehow helps us understand ourselves and demystify what being human means.", this HP executive explained.

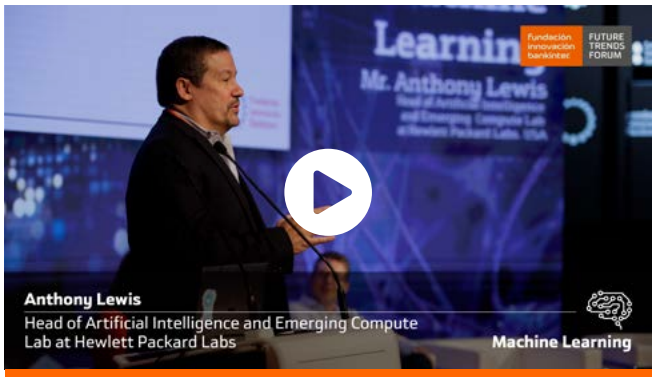

Duration: 07:23 min

What is new? Machine learning has been evolving for quite some time, and it has become a central component of artificial intelligence. As we have said in previous chapters, the large amount of available tagged data has significantly contributed to its progress. But that increase in data, according to Lewis, has placed us at a crossroads: machine learning largely happens in the cloud, but the increase in bandwidth is minimal.

"Data is trapped at the limit, and a lot of data has to be processed up until that limit. That is why there is an increasing need for a specific hardware, a neuronal accelerator, which is not too expensive and is extremely powerful." This could lead us to distribute our capacity to the limit; we would then have a processing capability of terabytes (thousands of GB) right in our pockets, in our cars... There would be large amounts of data processing happening around us.

Lewis mentions a side effect of these processors that could emulate up to 10 billion neurons. "The cortex of the human brain has a little over 20 billion neurons, so we are reaching a scale that corresponds to that of humans." That is, at least in terms of the figures because the neural simulations "are not nearly as realistic as human ones."

He also points out that these advances are based on an "unprecedented, open-source movement. The minute you publish something, it is reproduced by other research groups that can enhance it. So, this field evolves incredibly fast." He says that people are learning how to build self-refining systems, in a way that the skills previously required to master machine learning are no longer needed. He believes this will involve a drastic revolution of AI. "Big companies will not have any power, and college students will be able to do amazing things."

HP also focuses on "symbiotic machine learning." This involves copying and improving people's performance in a symbiotic way. For instance, HP monitors how call center employees interact with people in order to extract adaptable models and make interactions more efficient. Lewis believes this will be one of the main trends that we will see in the next few years.

3.1.5

Quantum Computing

Tanisha Bassan, a young, quantum programming developer who was named a 'Young Innovator to Watch' in CES 2019, attended the FTF to talk about quantum computing and its association with AI. Her work revolves, around, among other things, how machine learning and quantum computing are going to overlap. In order to understand this, we need to understand what quantum computing actually is. "Quantum computing is leveraging quantum mechanics properties, the rules and laws that we use to understand how nanoparticles work and interact with each other, to create a computer that will help us expand our computational power.", explains Bassan.

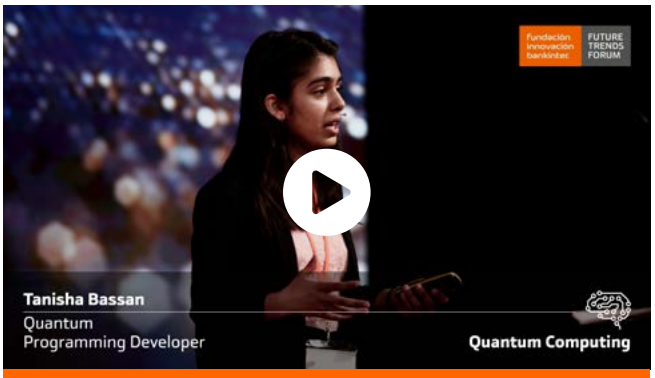

Duration: 07:30 min

- Quantum computing can simulate physics and help us understand the complexity of nature

According to this researcher and scholar, quantum computing has paved the way for a classic, computing revolution, whose possibilities are limited by hardware. "We now see a significant amount of technology in the field of quantum computing that will help us go well beyond the meaning of computing and what can be achieved." She talks about a technology that can simulate physics in its entirety or help us understand the complex way in which nature works by leveraging nature itself.

"The basis for quantum computing is leveraging on quantum mechanics, the rules and laws that we use to understand how things at a very small scale work and interact with each other." She mentions "two very useful properties." The first one is superposition, which allows particles to be in different states simultaneously. The second one is entanglement, the idea that it is possible to have two particles or two points of information that are correlated. Whatever affects one of them, directly affects the other one.

She explains, "These two properties are the basis for quantum mechanics and the reason why computation can, at that scale, evolve into exponential, computational power that will allows us to resolve some of the most complex problems in the world." As an example, she presents an analogy. "Let's think about a library with millions and millions of books, accounting for a huge

amount of knowledge. I grab a book and choose a page with information I want to retrieve. I put a mark on it. If I ask a traditional computer to find that very same mark for me, it could take forever to find it because it would need to go over every single one of the pages in each book, one by one. A quantum computer, however, could access the same library and check all the pages of all the books at the same time, finding that mark a lot faster."

Bassan believes that is particularly interesting because most of the data that already exists, as well the data that has not been created yet, is not structured. We can neither analyze nor understand it. There are complex problems that are hard to emulate mathematically with our current programs and hardware; she believes that quantum computing will help us address them. "For instance, climate change, hunger, the food crisis, the housing market crisis, cancer... thanks to quantum computing, we would be able to access their underlying ideas and concepts and map how complex these problems actually are; then, we could create hardware that would understand them in ways that we cannot even imagine yet."

To provide a more specific example, she mentions the discovery of medical drugs. She starts with the premise that current computers lack the necessary computational power to emulate real environments with precision, whether that is for a drug or for any physical system. "If we could do that and understand how the drug that has been developed interacts under specific environmental circumstances, we could dramatically reduce the time it takes to research and develop that drug. It would go down from 10 or 15 years to just a couple of weeks. We would then be able to address many other problems that we are currently facing in terms of disease control."

What does this have to do with machine learning? She uses the analogy of a car. Quantum machine learning would allow us to define how to address complexity, using different techniques yet to be researched and developed and allow us to offer a context to understand it and make decisions. Quantum computing will be, as she states, the fuel for the engine and will provide the necessary power to understand complexity very quickly.

Quantum computing will help solve optimization problems

"If quantum machine learning becomes the tool of the future, it will probably be able to solve, in the long term, some of the most complex problems in the world." Bassan believes this revolution will be similar to that of the Internet. "Just like back in the 50s, we didn't know where computers would lead us or that the Internet was possible; there is no way for us to know what quantum computers will do 50 years from now, but it will surely be revolutionary." In the shorter term, 10 or 15 years from now, she believes that this technology will change how we collect and understand data. Quantum computing will also help solve optimization problems, helping find the best possible solution, and it will increase computational power in most industries.

3.2
Future Scenarios

We have asked our experts to think about the situation 10 years from now (2030) and predict how AI and its technological developments might affect our daily lives. In this potential glimpse into the future, we will talk about the following areas:

3.2.1
Employment and Society

AI will play a key role in the future of employment. It will drastically change the way we work and our productivity.

The welfare state will become very relevant because of the implementation of a basic income system driven by machines' productivity. It will be calibrated by a cumulative social credit system that will classify and rate employees to measure quality of life and access to medication, among other factors. Whoever receives this basic income and all its benefits must give up their privacy and allow access to all their data.

This system will manage housing and give access to information and on-demand services (in a system where the product is offered as a service), leisure, videogames... making people satisfied and happy overall.

This system will face the following challenges:

- Making the Basic Income system work and achieving social welfare
- Managing housing as a key asset
- Promoting education (since we would no longer need it to have a job)

3.2.2
Education

AI will be broadly applied to the field of education. This will result in a more customized education system, which will integrate AI training across all stages of education to different extents, regardless of what is being taught.
Teachers will learn AI-related education techniques that will allow them to identify students who won't be fit for the technological industry. These students will receive training to develop emotional, care-oriented skills and careers. Teachers will teach each student individually, and not as a group.

Students will not only have teachers, but customized AI tutors. Education will not focus on the lowest common denominator, but rather on the individual development of each student. Students will be prepared to become productive members of society and reach their full potential thanks to complex, creative and emotional jobs.

The main challenge in terms of education will be teaching social skills, since customized education minimizes socialization.

3.2.3
Health

Thanks to AI and patient data (including connected, wearable devices that will send health information constantly), health professionals will devote more time to each patient. They will spend fewer hours writing reports and will offer more precise, customized diagnoses.

Doctors will not be replaced. Instead of having AI as a personal physician, we will have health assistants that will help people change their behavior every day; these will be driven by wearable devices that will send information to doctors. As a result, people will change their daily habits in ways that we cannot even imagine today.

This will cause a paradigm shift in the general healthcare system, particularly in healthcare services. We will no longer treat diseases, but predict them and prevent them instead of reacting after the fact. Therefore, emergency assistance costs will decrease, and we will have more equitable health systems and better healthcare in the long run.

This new health system will face several challenges:

- We will need to cohabit with technology and learn to cooperate with it; this requires training both doctors and patients.
- We will need to design technology from the perspective of helping and improving healthcare professionals, not replacing them.
- Logically, the cost of assisting chronic patients and taking care of them will be more expensive, especially given the increase in life expectancy and population ageing.

3.2.4
Public Administration

Technology will make plenty of data (information on cities, roads, hospitals, schools, etc.) available to society and to AI, thus facilitating decision-making. In turn, this will improve efficiency in public services, access to housing, food safety and healthcare as well as digital rights and data protection, which will reduce corruption and establish more precise property rights.

People will have more visibility; therefore, communities will be more empowered. They will be able to decide where they want to live, and how. Cities will then be more fit to inhabit.

The challenges are:

- The availability of public information. Governments will have to deal with higher social demand to resolve problems, without prioritizing needs.

- We will see how citizens make decisions, particularly about the self-management of communities and neighborhoods.

3.2.5
Management

A complete integration of AI into company management and monitoring will enable the automation of some of the corporate decision-making processes. AI will change the way in which companies and institutions position themselves in the market, as well as the role they will play in the value chain of the corporate ecosystem. Additionally, we will need regulations regarding automation, job destruction, digital taxes, etc.

AI will pose the following challenges for business management:

- Employers will have to acquire new skills to apply AI.
- There might be sales and mergers and acquisitions with increasing business aggregation.
- New companies might stagnate.

3.2.6
Citizen Security

Improper use of AI might facilitate criminals, hackers and white-collar criminals the right tools to steal, intimidate and manipulate anyone, including those with the most power (governments, mass media or people).

This is why it is essential to develop the necessary security, privacy and digital rights mechanisms.

3.2.7
Extreme Scenarios

Overall, predictions by our FTF experts have been quite optimistic and moderate. Below, we present the most radical predictions about the future, which are mostly pessimistic:

- There will be a new AI winter due to social rejection of this set of technologies.
- AI and human co-dependence will occur, humans being emotionally dependent on AI.
- Frustrated AI: its development will be hindered due to arbitrary legislation that responds to society's complaints.
- Cyberrivalry among countries: competition among governments to develop AI will turn into a military race.
- AI's development will lead to a fractured world with a disjointed society.
- Cyber disasters due to cyber attacks will occur.
- Algorithmic colonialism: colonial domination of technology powers over the rest of the countries; Europe will lag behind.
- Surveillance based on a social credit system will become the norm; it will control citizens' behavior and knowledge.
- 'AIstocracy': AI will become a luxury item.
- AI and the Internet will converge in a new network: Net-AI.
- Countries' power will be overshadowed by large, AI and technology providers.
- Integration with extended realities and other technologies will give rise to trans-human species
- Economy of abundance: due to AI, productive, manufacturing, logistical and commercial efficiency (among others) will reduce the cost of all goods and services that are necessary for ensuring good quality of life.

04

Recomendations

The analysis of this trend **has led us to underscore 5 general recommendations for optimal**, AI development in the future.

- Humans must always be the priority and exist at the center of AI, from inception to design, development, and implementation. AI must augment human ability, not replace it.
- In order to leverage AI and integrate the technology in any entity, company, organization or system, humans must collaborate with machines or software. Therefore, we need to develop integration and training strategies that allow people to see the mutual benefits and feel that their abilities are enhanced.
- In the corporate world, we must understand how AI development and integration can transform business capacities, optimize, help make predictions and create brand value.
- In government and society, we must take advantage of the possibilities that AI offers in different areas and disciplines for the improvement of resource management, mobility, education, health and quality of life. AI can help streamline process and provide services more efficiently. It can help us take on great global challenges and could prove especially useful in meeting the Sustainable Development Goals (SDGs).
- While taking advantage of the opportunities that AI affords us, we must anticipate and prevent risks. These are our recommendations for ensuring ethical AI development:
  - **a.** Tackle biases.
  - **b.** Place people at the center. Try to enhance their abilities, not replace them entirely.
  - **c.** Strive to prevent unwanted, secondary and tertiary effects.
  - **d.** Be transparent.
  - **e.** Ensure accountability through external audits.
  - **f.** Improve designs.
  - **g.** Formulate privacy policies that are more coherent.
  - **h.** Develop algorithms that are intrinsically explainable; they need to allow people to understand how an AI system reached a certain conclusion.
  - **i.** Promote debate, decision making and supervision by the people.
  - **j.** Use AI to teach people about AI.
  - **k.** Expand and defend digital rights.
  - **l.** Hold companies accountable for the appropriate supervision and management of the risks that stem from these technologies and ensure that companies address any damages that these technologies may provoke. They must pay the social costs that stem from their business.
  - **m.** Explore new forms of wealth redistribution and data ownership.
  - **n.** Promote AI development for social good.
  - **o.** Push for the development of FATEN algorithms (Fairness; Autonomy, Accountability, Augmented; Trust, Transparency; Education, Beneficence; Non-maleficence).

Acknowledgements

We would like to express our gratitude to the journalist Esther Paniagua, who wrote this report that captures the trend analysis and recommendations made by the experts who attended the Future Trends Forum.

We would like to thank the members of the FTF, the attendees of the 32nd conference and our collaborators for organizing and defining the methodology of the FTF conference:

Christopher Meyer

Garrick Jones
Clemens Hackl
Georg Seiler

Fernando de Pablo

Lastly, we would like to express our gratitude to the Bankinter Innovation Foundation for their commitment to our mission of promoting innovation in our society:

Bankinter Innovation Foundation
Juan Moreno Bau
Marce Cancho
María Teresa Jiménez
Raquel Puente
Carmen Nestares
Nerea Igoa
Jose Carlos Huerta

The opinions expressed in this report are those of the author. They do not reflect the opinions of the experts that participated in the **Future Trends Forum conference**.

Speakers and Assistants

Gary Ang
Managing Director of Strategy at Temasek

Tanisha Bassan
Quantum Programming Developer

Charles Bolden
Former NASA Administrator. President of The Bolden Consulting Group LLC and Trustee of Bankinter Innovation Foundation

Calum Chace
Writer and Speaker on AI

Antonio Damasio
University Professor and David Dornsife Professor of Neuroscience at the University of Southern California, Trustee of Bankinter Innovation Foundation

Hanna Damasio
Dana Dornsife Professor of Neuroscience and Director of the Dornsife Neuroscience Imaging Center University of Southern California

Soumitra Dutta
Professor of Management at Cornell University and Trustee at Bankinter Innovation Foundation

Esther Dyson
Executive Founder of Wellville

Lauren Dyson
Senior Data Scientist at Civic Analytics

Stephen Johnson
Senior Fellow of the Mossavar-Rahmani Center for Business and Government at Harvard Kennedy School

Richard Kivel
Director General de GrayBella Capital y Patrono de la Fundación Innovación Bankinter

Philip Lader
Senior Advisor of Morgan Stanley and Trustee of Bankinter Innovation Foundation

Jade Leung
Head of Research & Partnerships with the Center for the Governance of Artificial Intelligence (GovAI)

Anthony Lewis
Head of Artificial Intelligence and Emerging Compute Lab at HP Labs

Julia Li
Founder and CEO at HCD Learning

Michael Li
Founder and CEO at The Data Incubator

Qing Liu
President of Jiangsu Industrial Technology Research Institute (JITRI)

Ramón López de Mántaras
Director of the IIIA (Artificial Intelligence Research Institute) of the CSIC

Pilar Manchón
Vice President of AI at Roku Inc.

Christopher Markou
Leverhulme Fellow and Lecturer, Faculty of Law, University of Cambridge & Course Director, 'AI, Law & Society', School of Law, King's College London

Fiona J McEvoy
Founder of YouTheData.com

Emilio Méndez
Director of Energy Science and Technology Department at Brookhaven National Laboratory and Trustee of Bankinter Innovation Foundation

Nuria Oliver
Chief Data Scientist at Data Pop Alliance & Fellow at the Spanish Royal Academy of Engineering

Eduardo Porter
Economics Reporter of The New York Times

Pablo Rodríguez
Founder of Telefonica Innovation Alpha

Raúl Rojas
Professor of AI at Freie Universität Berlin

Juliana Rotich
Technologist, MIT Media Lab Director's Fellow and Trustee of Bankinter Innovation Foundation

Michael Schrage
Research Fellow with MIT Initiative on Digital Economy

Eden Shochat
Partner at Aleph and Trustee of Bankinter Innovation Foundation

Dor Skuler
CEO & Co-Founder of Intuition Robotics

Oliver Smith
Strategy Director, Health Moonshot at Telefonica Innovation Alpha

Tan Chin Nam
Former Permanent Secretary, Singapore Public Service and Senior Corporate Adviser. Trustee of Bankinter Innovation Foundation

Andres Torrubia
CEO of Fixr.com

Wilfried Vanhonacker
Co-founder and Ex-Dean, CEIBS (Shanghai) and MSM Skolkovo (Moscow). Trustee of Bankinter Innovation Foundation

Kush R Varshney
Principal Research Staff Member and Manager at IBM Research

David Weinberger
Senior Researcher, Harvard Berkman Klein Center for Internet & Society
Writer in Residence at Google PAIR

Adam West
Head of Marketing at Satalia

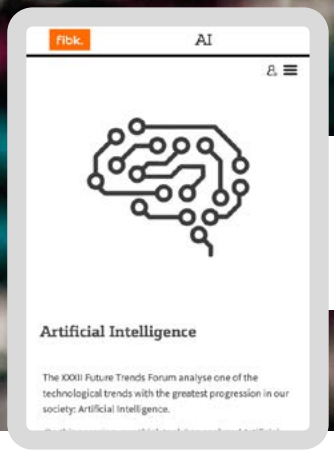

To obtain more information about this trend, click on:
fundacionbankinter.org/en

fundación innovación bankinter. FUTURE TRENDS FORUM

Fundación Innovación Bankinter
Paseo de la Castellana 29
28046 Madrid, Spain

FundacionBKT
Fundación Innovación Bankinter
Fundación Innovación Bankinter
Fundación Innovación Bankinter
fibankinter