viaSarfatti25
BOCCONI UNIVERSITY, KNOWLEDGE THAT MATTERS
Issue 2 / 2023
ISSN 1828-6313

Artificial Intelligence is a revolution transforming the economy and upending society. It can be ethical and equitable if issues of fairness and transparency are addressed. Without fear, because this is the future of the interaction between humans and machines

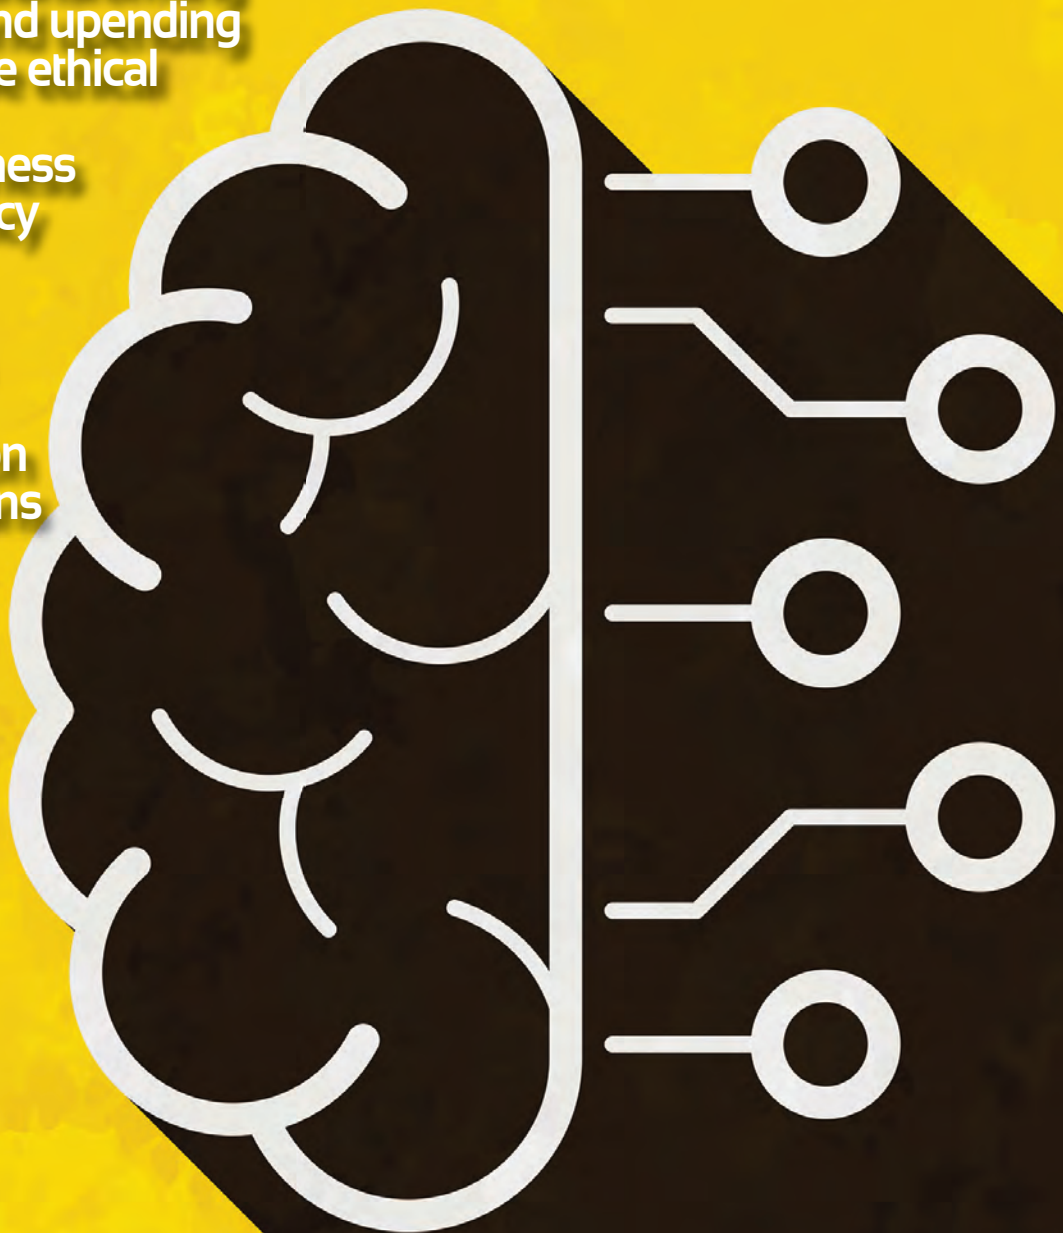

WHAT'S NEXT FOR AI


Bocconi
follow us
@unibocconi

Artificial Intelligence in the life of Bocconi

On March 27, the new Master of Science in Artificial Intelligence started receiving applications. This is the newest component of the investment that Bocconi has decided to make in the area of Artificial Intelligence, computational sciences and data science. It adds to the previously opened Bachelor of Science in Mathematical and Computing Sciences (BEMACS) and Bachelor of Science in Artificial Intelligence (BAI), the Master of Science in Data Science and Business Analytics, and the Department of Computing Sciences inaugurated last Spring. This is an indispensable investment for the growth of Bocconi and the development of impact research for society, the economy and institutions. The exploration continues with cognitive and behavioral sciences: in this field we celebrate the winning in March of an ERC Advanced Grant for the MemDec (Memory, Beliefs, and Economic Decisions) project directed by Nicola Gennaioli, working at the intersection between cognitive sciences, economy and finance.
Bocconi's choice, as demonstrated by this issue of our magazine, is to encourage the use of Artificial Intelligence and machine learning in all departments, to guarantee an innovative approach to research which is increasingly linked to Big Data in all domains. AI therefore is not only an object of study but also a tool of knowledge both in research laboratories and university classrooms. ChatGPRT has also been tested in various lessons in class: because the issue is that you should never stop innovation, but understand how to use it in a positive way while regulating its use. Without being afraid of it.

Francesco Billari
Rector

United in s

A community that recognizes itself in sports values, team spirit and positive competitiveness, experiencing sports not only on the field but also in the classroom, always aiming for innovation and research. A journey to discover the Bocconi Sport Team and all those in Via Sarfatti *who live sports as passion or profession*

Photo gallery

The Bocconi Sport Team

The focus

ports

CONTENTS

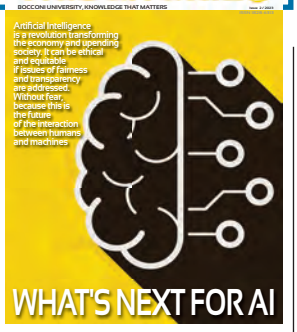

Issue 2 / 2023
Publisher: Egea Via Sarfatti, 25 Milano

Editor-in-Chief
Barbara Orlando
(barbara.orlando@unibocconi.it)

Editorial Office
Andrea Celauro
(andrea.celauro@unibocconi.it)
Susanna Della Vedova
(susanna.dellavedova@unibocconi.it)
Tomaso Eridani
(tomaso.eridani@unibocconi.it)
Davide Ripamonti
(davide.ripamonti@unibocconi.it)

Translation and revision
Richard Greenslade
(richard.greenslade@unibocconi.it)
Alex Foti
(alex.foti@unibocconi.it)
Jenna Walker
(jenna.walker@unibocconi.it)

Contributors
Paolo Tonato (photographer)
Michele Chicco, Pietro Masotti, Camillo Papini

Secretariat
Nicoletta Mastromauro
Tel. 02/58362328
(nicoletta.mastromauro@unibocconi.it)

Layout project: Luca Mafechi
(mafechi@dgtprint.it)

Production
Luca Mafechi

Registered at Court of Milan
numero 844 del 31/10/05

The AI Revolution

If the energy that fueled the industrial revolution was fossil energy, the basis of the epochal change we are experiencing today is information and its processing. But we are only at the beginning and the challenge on the horizon is to understand and control machine learning algorithms and the very nature of intelligent behavior through ever greater integration between cognitive sciences and computational neurosciences. Without ever renouncing, however, to being transparent, fair and responsible

by Marc Mézard and Riccardo Zecchina @

You have probably heard about ChatGPT. And maybe you have tried it, and you were surprised. You are likely to be intrigued, worried, or irritated about its possible uses or misuses… This software is actually the latest in a series of stunning breakthroughs in Artificial Intelligence (AI) that took place in the past decade. Machine learning algorithms, which learn statistical rules by themselves through their analysis of extremely large database, have reached human-level performance in many tasks. A few landmarks: image analysis and image-based diagnosis, playing complex games and strategical planning, predicting protein shapes and designing new molecules for medicine, generating images or texts.

No one predicted these recent discoveries, so our present predictions are likely to be wrong! Anyhow, let us take the risk and try to sketch at least some broad scenarios.

While the industrial revolution was an energy revolution, using machines that consume fossil energy as a substitute for the physical work of humans, the AI revolution is one based on information processing. New tools of artificial intelligence are being integrated into human activities as auxiliary agents and this trend will accelerate. To give just a few examples, let us mention medical-diagnosis aids, planning, generating standard correspondence, customer profiling, or AI-assisted scientific research. Recent dramatic improvements in large language models will certainly expand and help this integration, but they will also put pressure on the economic structure in order to incorporate new tools efficiently.

As we continue to develop AI technologies, it is vital to consider the ethical implications of their use. We must ensure that the development and deployment of AI are transparent, fair, and accountable; developing the right regulatory framework, and educating people to the correct use of AI devices, are priorities. The present developments are also strongly dependent on the access to very large amounts of data and computing power, as a consequence the new technologies are developed mostly by a handful of very big companies. We must ensure that the benefits of AI are shared equitably across society and are not

MARC MÉZARD
Full Professor of Theoretical physics, Bocconi University

RICCARDO ZECCHINA
Full Professor, Vodafone Chair in Machine Learning, Head of Department in Computing Sciences, Bocconi University

THE BOOK

The idea that inspired *Intelligenze artificiali e aumentate*, by Marco Pironti (Egea, 2022, 288 pp., in Italian, €38), arises from a debate between academia, entrepreneurs, managers and representatives of institutions on the need to address the challenge that artificial intelligence represents not only from a technological point of view. The need to increase the available intelligence requires an overall strategy based on shared policies and agile and adaptive governance mechanisms that this technology imposes.

kept in the hands of a few companies.
One can expect the research on AI to develop in two main directions. On the one hand, the understanding and control of machine learning algorithms used to train modern AI devices is crucial. At present, we have no complete explanation for how these machines achieve their astonishing performance on new tasks, they are used as black boxes for which there exists no guarantees. Secondly, recent discoveries have brought to questioning of the very nature of 'intelligent behavior'.
While the success of large language or image-making models is largely due to training bigger and bigger devices, we believe that scaling up their size is not going to close the gap with human intelligence. Cognitive science shows that a baby learns a concept by guessing and generalizing a rule from just a few examples, very far from the database with millions of items used in machine learning. While purely statistical approaches have shown their power, intermeshing them with logical rules is a big challenge.
On this road towards more 'intelligent' behavior, another critical development for AI will be the capability of generating internal representations of the world, building on unsupervised data. One direction could be reinforcement learning, which rewards specific actions and leads to learning by trial and error. This technique has already led to great discoveries but in limited scopes like games. In the next years it will be used in more realistic setups where actions can help build representations.
As neurosciences are also making remarkable progress in collecting data on neural activity and correlating them with various tasks, a new convergence between computational neuroscience and AI can take place. For instance, some recent experimental results on the non-stationary nature of the internal representations in working brains have stimulated new discussions in both communities. Research at the interface of these two disciplines should provide deep new insights on the nature of intelligence, with an impact both for designing new AI devices and for a better understanding of the brain.
No doubt: what we are witnessing is an AI revolution. Its impact on the organization of the societies will be major and needs to be controlled and managed. As any technological innovation, it can be misused, either deliberately or by mistake: the ethical consequences need to be considered and regulated. At the same time, this AI revolution opens major and fascinating new scientific perspectives. As scientists, let us remain humble and acknowledge the limitations of current AI systems while striving towards a better understanding of intelligent behavior and developing ethical frameworks for its use. ■

PODCAST
In an episode of the Talent Show podcast, the series devised by Financial Times and Bocconi for the younger generation with analysis on trends and career advice in the most coveted sectors, Marc Mezard, Professor of the Department of Computing Sciences at Bocconi University, talks about developments in Artificial Intelligence and building a career around it. He discusses the research around AI that is still required, pragmatic applications of AI and how it will redefine the job market.

THE DEPARTMENT

The Department of Computing Sciences aims at addressing fundamental problems in computation, artificial intelligence and their applications in computational neuroscience, bio-medicine, socio-economics and complex systems modelling. This goal is supported by a diverse and ambitious faculty that mobilizes expertise in several disciplines of computer science, mathematics and physics.

THE BOOK

We live in the era of computational simulation: to grasp this passage of civilisation and return to understanding the world, we must address the (epochal) change underway with the depth of vision proper to philosophy. In *Il mondo in sintesi* (Egea, 2022, 192 pp., in Italian, €19) Cosimo Accoto accompanies us in the surprising tale of his journeys into new unknown lands, constructing more conscious and up-to-date cultural maps of the transformations taking place.

Bocconi

Bocconi Podcasts
MUCH MORE THAN WORDS

Bocconi podcasts enable listeners to go deeper and learn more about topics that really matter, in order to get inspired and gain new perspectives on relevant issues. Another way to fulfill our mission of spreading knowledge around the world.

Welcome to our podcasts.

What are the causes behind the most puzzling issues of our time? This series of talks shines a light on solutions for this chaotic world, in order to gain clarity thanks to the experience of guest experts interviewed by David W. Callahan.

Everybody tells a story. But nobody tells it as we do. This podcast series collects facts and reports from ViaSarfatti25, Bocconi University's magazine. Listen to what professors and researchers have to say about the big topics in the fields of business, management, data science and politics.

Hosted by under 30s, for under 30s around the world, The Talent Show delve deep into the topics that really matter to today's younger generation. Find inspiring tips, analyse trends and bridge generational gaps with Bocconi professors and alumni

THINK DIVERSE is a fresh look at the issues surrounding diversity and inclusion, because we all need to understand how diversity manifests itself in our daily lives, our society and our economy. Here at Bocconi we can do so by taking a deep dive into our faculty's research on D&I.

www.unibocconi.eu/podcast/

Listen on Apple Podcasts
Listen on Google Podcasts
Listen on Spotify
Listen on Spreaker

Would you be afraid

THE PAPER

Viewpoint: Artificial Intelligence Accidents Waiting to Happen?
By Federico Bianchi, Amanda Cercas Curry, Dirk Hovy

# of a three-year old?

AI cannot decide to perform tasks that are beyond its capabilities and above all it has no theory of mind, no feelings, no conscience, no malice. If anything, AI uncovers the naivety and mistakes of its designers. Try it yourself to gain knowledge and find out

by Dirk Hovy @

AI news comes in two varieties: hype and doom. Recently, the hype cycle has been prevalent. AlphaGo defeated the best human player, machine translation became a useful tool, and chatGPT now provides witty responses to all questions. But doom is never far away: warnings of job losses, killer robots, and sentient AI. The fear is understandable, but it reflects our perceptions of AI more than its actual capabilities. The main cause is humanizing those models and assuming they have drives, motives, and emotions that would cause them to act evil when they really just do a task.

To be clear, current AI technology has numerous flaws, so we must use and develop it with caution. These problems are caused by bias and user discrimination. The automated grading system in the United Kingdom unfairly punished some students, the machine translation system that incorrectly translated "Good morning" as "Attack them," resulting in legal trouble for an innocent person, and the speed camera that issued a ticket to an innocent driver because it mistook a knitted jumper for a license plate are all examples of how AI can cause havoc. The worst example is the Indian man who starved after an automated decision system denied him food rations.

All those tools, though, acted out of design flaws, not malice. The consequences are still dire, but it identifies the problem.

Despite these concerning reports, I do not anticipate any lethal AI threats. I'm not alone. According to Andrew Ng, a pioneer in neural networks, "Worrying about AI evil superintelligence today is like worrying about overpopulation on the planet Mars. We haven't even landed on the planet yet!". What evil actions would a sentient machine translation even take? Produce poor translations to irritate you?

The concerns are understandable, though, given that we have machines with human characteristics. They play games, answer questions, translate sentences, and identify people in photographs. If they can do all of that, they must be like us, right? So they most likely also have hopes, dreams, and aspirations. But complex artificial intelligence systems must be tailored to each task, such as Go, sentence analysis, or photo recoloring. These three are each unable to complete the other two tasks. Despite the efforts of many intelligent people, AI tools frequently perform like three-year-olds. AI cannot yet decide to perform tasks that are beyond its capabilities.

DIRK HOVY
Associate Professor of Computer Science, Bocconi University

Even AI researchers are affected by humanization bias, though. An external Google employee, Blake Lemoine, claimed after extensive discussions with Lambda that the model possessed self-awareness and consciousness. A journalist posed slightly different questions to Lambda from Lemoine. And the model denied being conscious.

Maarten Sap, a University of Washington researcher, investigated language models' Theory of Mind (the ability to imagine and understand the thoughts and feelings of others). A patient's Theory of Mind can be determined using a variety of question-based psychological tests. They can also be administered to language models. But this logic is flawed. People answer questions based on their complex inner workings. Language models just generate a list of likely words in response. While their behaviors are similar, their motivations and paths to the same goal are not. As a result, asking whether models have these psychological abilities is pointless. They do not have a Theory of Mind, feelings, consciousness – or malice.

Why and how would they develop this capability in the absence of explicit programming? Each AI task must be meticulously defined and trained. That never includes giving sentience, emotions, or aspirations. AI models may reflect the naiveté and lack of checks and balances of their designers, but they do not act out of evilness.

So should you fear AI becoming evil? No. Should you keep an eye on their design? Definitely. Should you try using it on a daily basis? I invite you to try AI. What we understand cannot scare us. ■

by Luca Trevisan @

The capabilities of Artificial Intelligence (AI) software systems are rapidly progressing, and they are expected to have an increasingly disruptive impact on several industries, with substantial positive effects on the economy and society.

As with any technological innovation, the adoption of AI technologies carries certain risks, that we will have to learn to recognize and prevent. One such risk is that AI systems have the potential to perpetuate biases and discriminations against certain groups, for example on the basis of gender, ethnicity or age.

How can an AI algorithm have biases or create discrimination?

Typically, one trains an AI system to perform a certain task by using data that shows examples of correct solutions to the task. A Machine Learning (ML) training algorithm then comes up with a model of the data and of the task that works well on the given examples, and usually this generalizes well on subsequent examples.

For example, one may train a system to provide a description of a given photograph by first showing it millions of examples of photographs together with their descriptions. The Chat GPT-3 system was trained for the task of predicting the next word in a sentence, by giving it billions of examples of fragments of sentences together with the subsequent word in the sentence. Given this next-word-predictor, Chat GPT-3 basically applies it over and over to create longer texts, and it is able to create surprising coherent essays, stories, poems, and so on, as anybody who has played with it will have seen.

Back to how discrimination can arise, suppose that a company wants to automate parts of its hiring process, and to do so it creates an AI system to score the likelihood that a given job applicant will be hired, so that the lowest scoring applicants can be automatically rejected and only the top scoring applicants are evaluated by a human. In this case the company would give the ML algorithm data about past applicants and whether they were hired, so that the system can build a model of what features in an applicant's CV make him/her likely to be hired. What if, however, the company has been discriminating against women applicants in the past? Then, in building a model that is faithful to the data, the ML algorithm would learn to give women lower scores than men. The same would happen if the company had discriminated against older applicants in the past, or against applicants with certain other characteristics.

LUCA TREVISAN
Full Professor of Computer science and director of Master of science in Artificial Intelligence, Bocconi University

A similar problem could arise if one wants to build a credit rating system to predict how likely is a given loan applicant to default on the loan, or a system to predict how likely is a convict applying for parole to commit crimes again if released. One would train such systems on past data, but past data could reflect past discrimination, for example against people of certain ethnicities, and the risk is to train a model that would also make biased choices.

Unfortunately, one cannot protect from this risk by simply erasing sensitive characteristics such as gender or ethnicity from the training data, because such characteristics might strongly correlate with other data that is present, and the AI system might learn to infer those characteristics from the available data, and then discriminate on such basis!

A concrete example of this risk occurred a few years ago in the US: the COMPAS system was a proprietary algorithm developed to predict the likelihood that a given convict would become a repeat offender, and it

It is not machines that perpetuate bias and prejudice because their predictive responses rely on what they have stored in them. To address issues of fairness and responsibility, it is therefore necessary to build so-called explainable AI systems

It's all about data and whoever trains the alg

was used by courts to decide on sentencing and on parole. The algorithm took into account 137 parameters, and the race of the person was not one of the parameters, but a 2016 investigation by Pro Publica demonstrated that the system had a consistent bias against Black convicts, assigning them a risk of repeat offending that was higher than what was the actual rate of repeat offenses seen in subsequent years, while white convicts were assigned a risk that was lower than the effective rate at which they committed crimes again in subsequent years. The model underlying COMPAS, evidently, had effectively learned from its training data how to predict the ethnicity of people and then it discriminated on such basis.

Another source of biases in AI can come from a similar more benign issue, namely the representation of different groups in the training data. In ML, one needs large amount of data to achieve good precision, so if a certain subgroup is under-represented in the training data it is possible that the ML algorithm will create a model that is less precise for that subgroup. For example, facial recognition systems trained on pictures that show mostly faces of white people might have a high error rate when trying to recognize the faces of people of other ethnicities, a problem that has come up multiple times in the area of face recognition.

These issues are concerning but there are reasons to be optimistic about them. The fact that biased data leads to biased predictions is a classical problem in statistics, and there are approaches to audit training data before using it and to audit a model after the training phase to detect possible biases. Auditing a model created by a modern ML algorithm is difficult because its predictions are based on complicated calculations involving millions or billions of numerical parameters, so, partly motivated by fairness and accountability issues, the problem of building so-called explainable AI systems is being actively investigated. An explainable AI system makes a prediction and also provides a human-readable explanation for its prediction. Finally, there is ongoing work to enact regulations around the use of AI, such as the proposed AI Act currently being discussed by the European Parliament. ■

orithm

ChatGPT and generative AI risk amplifying phenomena like hate speech and online defamation, opening up ethical questions as well as legal ones. There is also the issue of energy consumption

The fallibility of Artifi

by Oreste Pollicino @

Having a dialogue with an artificial intelligence is the fad of the moment. Colossal investments, surprising results, and also growing concerns have placed technologies such as ChatGPT - generative artificial intelligence - on the crest of a wave of digital enthusiasm.

One of the most significant problems that arise in this regard is the possibility of online disinformation gets further amplified, because access to knowledge could be distorted by texts that are often convincing and persuasive, but not entirely truthful.

An essential reference in legal reasoning is that constituted by fundamental rights which are protected by constitutional charters, which are analog but still very current documents. In particular, reference is made to the right to be informed, if not truthfully, at least reliably; to the principle of habeas data in its evident digital projection and to the right to fair remuneration for copyright holders.

These initial reflections lead us to argue that, beyond legislation in the pipeline, and in particular the proposed Artificial Intelligence Act, the humanist principle that characterizes the whole structure of Italy's Constitutional Charter and the Charter of Fundamental Rights of the European Union already offers clear orientation, with particular reference to the protection of the dignity of the person, to be protected as an individual and in his/her participation in intermediate social communities, which risk being subjected to an incessant process of fragmentation and disassembly due to the explosion of artificial intelligence.

In addition to the legal questions which, such as those relating to copyright protection mentioned above, are likely fill the days of judges, lawyers and scholars, there are at least two other questions that need to be asked. The first is of an ethical nature, the second is related to the energy problem.

As for the former, the not exceedingly sophisticated automation processes that extract an apparently original and convincing narrative text from the boundless repository of data present on the internet, often accompanied by a feel-good and very diplomatic tone, push many towards the erroneous belief that AI is "self-sufficient", and therefore the human component is completely absent. Nothing could be more false. To arrive that result, as observed in a recent article by Roberto Battiston and Massimo Sideri in the Corriere della sera, there is the work of thousands of "workers" of generative artificial intelligence, who could also be called the "new slaves", if one thinks that, for little more than a dollar a day, often in Asia or Africa, they carry out research and classification work within one of the most disturbing and alienating virtual worlds, that of the dark web, precisely to tame the verbal and visual aggressiveness of the early versions of ChatGPT. It is evident that in this regard many questions of an ethical nature can be raised, not only in connection to the process of exploitation - and, one could add, of veritable alienation to which these individuals are subjected - but also relating to possible scenarios in which, by changing the intentions of those who control the verbal and linguistic framework where the answers to our questions take shape, these could take a provocative, aggressive or even offensive tone, thus returning answers which, to the problem of disinformation, add those of hate speech and online defamation.

ORESTE POLLICINO
Full Professor of Constitutional law and director of the LLM in Law of Internet Technologies, Bocconi University

The second question, on the other hand, touches on energy profiles, which are extremely relevant in our current, well-known, historical predicament. As recently noted by CNR President Maria Chiara Carrozza, the use of generative artificial intelligence techniques is by no means free. In fact, it has a very significant consumption load and therefore, in some cases, causes energy waste, because, as we said, AI for its functioning requires the activation of a boundless quantity of data together with massive computational power.

We should resist the temptation of being overly pessimistic, however. From the point of view of the theory of language, as noted by Carrozza herself, generative AI constitutes the actual bridge between the humanities and computer science, as it uses a computational linguistics technique that builds a humanistic narrative on the basis of a probabilistic calculation typical of the hard sciences. ■

cial Intelligence

From genomics to imaging, from biomedical analyses to the ability to predict a patient's response to a particular treatment, AI can turbocharge researchers and doctors in finding answers to cancer, metabolic conditions and immune diseases. But we need to train a new generation of multidisciplinary research teams to reap its advantages

by Francesca Buffa @

Data that heal

There are great expectations that in the years ahead Artificial Intelligence (AI) will revolutionize medicine. Two key factors are behind this enthusiasm: an unprecedented increase in our ability to acquire patient data, and rapid maturation of methodology to use it. Indeed, AI is already used extensively in this area and is driving important fundamental discoveries and development of new tools. Data production across all health and life science disciplines is growing fast. Collectively, we are acquiring more data than we can make sense of, and thus the power of AI presents us with a tremendous opportunity to provide meaningful and useful analyses of such data. Importantly, not only the speed of data production but also data quality and resolution are

rapidly increasing across many disciplines, thus meeting the necessary requirements for AI application. Genomics - the study of our genes and their function - has arguably seen the greatest recent data explosion, and has caught up other Big Data areas such as astronomy and social media. Indeed, AI is already being applied to analysis of genomics data at all levels, from tackling the challenge of petabyte-scale sequence searches, to identifying genes associated with specific diseases, or using genetic, lifestyle and environment data to predict risk of disease.

By analysing comprehensive biological datasets from healthy controls or people with a specific disease, AI algorithms are likely to also identify promising treatments and guide development of new drugs. One recent AI advance which may lead to faster development of new therapies is AlphaFold, an AI algorithm predicting the structure of many proteins with high accuracy. Resolving and understanding structures can help scientists to understand protein functions, and how such functions may be altered in diseases such as cancer, metabolic or immune diseases. Prior to application of AI to this problem, establishing the structure of proteins was a difficult and time-consuming process that required years of skilled experimental work. Being able to predict such structures with high accuracy will allow us not only to inform development of new treatments for diseases, but also to speed up the drug discovery and development processes.

FRANCESCA BUFFA
Full Professor
of Computational Biology,
Bocconi University

AI can also help to build models that can predict a patient's response to a particular treatment, based on the omics data obtained from clinical biopsies, before and after treatment, and the individual medical history, leading to more personalized and effective care.

Another area of medicine where AI has shown particular promise is in imaging and diagnosis of disease. Analysis of images such as those acquired by X-ray, CT and MRI scans can improve the detection of abnormalities and identification of clinically significant patterns invisible to the human eye. In this area, deep learning AI algorithms have potential to reveal aspects of the data that would be missed by other methods. This should help radiologists and other healthcare professionals to detect disease at an early stage when treatment is likely to be most effective.

As with any new technology, application of AI to health must meet strict ethical, safety and transparency standards, and this does not come without challenges. Large quantities of high-quality processed data are required for development of useful AI. On one hand, careful primary data curation together with strict preservation of patient privacy and confidentiality are necessary. On the other, multiple levels of data processing, filtering and cleaning come with their own risks, such as making the final data less representative of daily medical practice. Furthermore, appropriate validation and testing of AI models to ensure they are unbiased is critical to avoid perpetuation of existing biases and disparities in healthcare. It follows that there is an urgent need for regulatory frameworks around the implementation of AI, and for adequate training and education of healthcare professionals.

Finally, the use of AI to analyse and interpret healthcare data necessitates parallel evolution of theory, modelling and application. While offering tremendous possibilities, high dimensionality and complexity of this data requires large computing capacities combined with development of new efficient approaches. Last but not least, we need to train a new generation of scientists able to tackle the open questions and work in multidisciplinary teams to develop AI methods that are effective and applied appropriately. ■

THE ERC GRANT

The 'microC' project aims to use multi-agent modelling and machine learning to model disease-host interactions in complex diseases such as cancer. The goal is to develop models able to predict which patient will be responding to which treatment.

Not all innovation glitters

Shared databases and a robust regulatory framework are the conditions for new AI algorithms to find real application in health care, thus translating new technology into actual business innovation

di Anna Gatti @

The convergence between the life sciences and digital technology has generated for the first time in history the real opportunity for a fruitful contamination between sectors and skills. When a convergence between distant sectors occurs, it is a harbinger of great innovation.
However, it is not enough that there are innovations in the lab to speak of service innovation. Socio-cultural, regulatory, and business implications are key for translating scientific discoveries into market innovations.
While there is no doubt that the topic of Artificial Intelligence (AI) in the health care industry is widely debated and is experiencing a moment of media hype, the applications deriving from AI have different levels of readiness for adoption by the market. In order to understand the effective applicability of AI in Italian health care, it is therefore necessary to look at individual applications and avoid speaking in a general way of machine learning in health care.
To do this, every year the LIFT Lab of the SDA Bocconi School of management produces out the LIFT Radar, a new research tool that offers a reasoned mapping of new AI applications, conceived to help decision makers acquire a clearer and more complete vision of the opportunities emerging out of the frontier of innovation born from the convergence between digital technologies and life sciences. By looking at the LIFT Radar, the reader can see an evaluation of different AI applications, based on precise variables, and gain a clearer understanding of how ready a specific application is to be successfully introduced in the Italian market. By applying this methodology, it appears evident that the domestic market is ready for the adoption of some artificial intelligence applications, while there are other areas for which normative, regulatory and organizational interventions are necessary so that technological innovation can find full-fledged business application.
The predictive models of radiogenomics, for example, which require the integration of data of different nature based on AI algorithms, find the necessary conditions of adoptability in our market. Therefore, it is definitely an area that health care organizations need to invest in today to remain competitive tomorrow. The solution of digital twins (digital aliases on which simulations can be performed), which today represent an area of great interest for venture capital investment in the United States, is limited in Italy by the presence of a few homogeneous and shared databases, as well as a limited diffusion of wearables and other sensors for dynamic data collection. In this case, although the technology is ready, the conditions (i.e. for example the need for dynamic data collection, the existence of homogeneous and shared databases) make the translation of the technology into a business solution not immediate.
The same limitations are noted for the use of AI applied to drug discovery. This solution is becoming increasingly relevant also thanks to the huge capital invested overseas by venture capitalists, the artificial intelligence algorithms are ready for the creation of the necessary models, but there is still a lack of shared databases and of a sufficiently robust regulatory framework, in order to enable the full translation of technological breakthroughs in business innovation. Not all AI is intelligent innovation. ■

THE LAB

The LIFT Lab, a new research lab at the SDA Bocconi School of management, studies the economic, financial, regulatory and socio-cultural conditions for the diffusion of innovation arising from the convergence of digital technologies and life sciences, including the study of Artificial Intelligence in healthcare.

ANNA GATTI
Associate Professor of Practice in Digital transformation and Director of Lift Lab, SDA Bocconi School of management

THE PAPER

Digital Twins in Healthcare: Is It the Beginning of a New Era of Evidence-Based Medicine? A Critical Review by Armeni, Polat, De Rossi, Diaferia, Meregalli, Gatti

The future lies in population health management

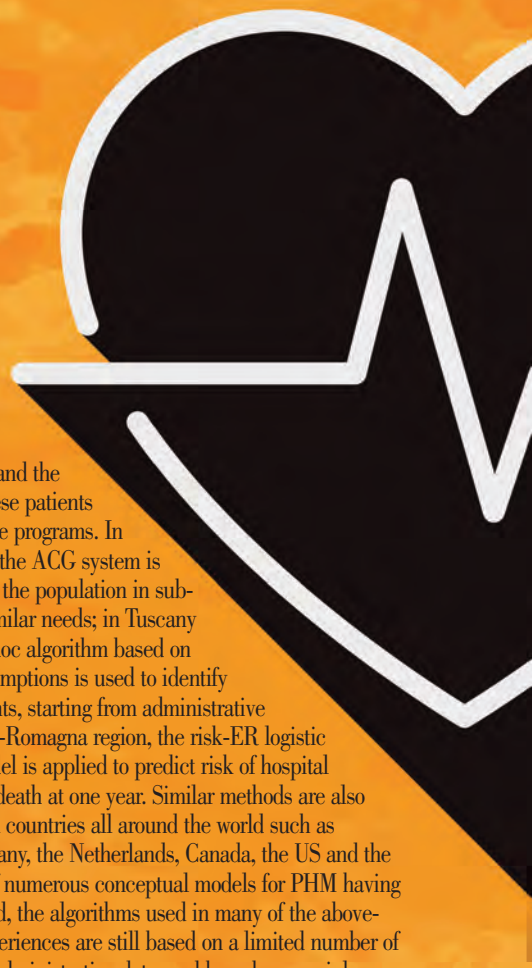

Thanks to the rapidly increasing amount of data available and the improvement of computational models, the PHM approach should be able to unlock its potential, for example by stratifying patient groups and identifying those at higher risk

di Lucia Ferrara @

LUCIA FERRARA
Associate Professor of Practice in Healthcare Management , SDA Bocconi School of management

That population is getting increasingly older and the number of people suffering from chronic conditions and multi-morbidities is increasing comes as no surprise. For health care providers, it becomes crucial to sort patients by clusters of needs and risk levels, in order to design targeted intervention according to these needs.
Great potential for personalization and targeted intervention comes from Population Health Management (PHM) approaches and analyzing huge amount of patient-level data, which maps patients' journeys across multiple sectors, and health providers, such as primary care, community care, hospital care, mental health and social care. PHM is a proactive approach to managing the health and well-being of a population that employs analytical tools (e.g. segmentation and risk stratification) to identify the patients, stratify them into groups of people with similar characteristics (for example fragile people with multiple complex health needs or healthy people) and to design targeted pathways and personalized services. PHM also enables effective prioritization, with a focus on identifying the patients at higher risk of developing health conditions or complications and define proactive models of care aimed at reducing hospitalization and death rates. There are several experiences of application of PHM approaches. For example, several regions in Italy are studying and testing algorithms and models to support the identification of patients, the segmentation, and the selection of these patients for specific care programs. In Veneto region, the ACG system is used to stratify the population in sub-groups with similar needs; in Tuscany region, an ad-hoc algorithm based on resource consumptions is used to identify complex patients, starting from administrative data; in Emilia-Romagna region, the risk-ER logistic regression model is applied to predict risk of hospital admissions or death at one year. Similar methods are also used in several countries all around the world such as Sweden, Germany, the Netherlands, Canada, the US and the UK. In spite of numerous conceptual models for PHM having been developed, the algorithms used in many of the above-mentioned experiences are still based on a limited number of variables and administrative data, and have been mainly used to describe the characteristics of populations rather than predict future needs. Indeed, predictive models in the management of health populations are relatively new compared with the area of clinical applications.
Looking at the future, thanks to the growing computational power, machine-learning methods and predictive modelling in population health offer great potential to inform proactive interventions, anticipate future disease burden and assess the impact of health policies and programs. Several studies have been recently conducted in the US, using machine-learning algorithms to predict future healthcare expenditures. Even if the Italian health system is very different from the American one, one question is if similar algorithms used to predict costs in the US could be also applied in Italy to flag patients at risks, shorten hospital stays, reduce unplanned hospitalization, and improve readmission rates. Additionally, future projects in machine learning could potentially incorporate larger datasets and additional non-traditional features, such as eHealth data, social media, and web search patterns. To unlock the potential of predictive models for public health the right combination of data, technology, and human intervention is needed. For the time being what will become reality and what will remain utopia is still a matter of debate. ■

THE PAPER

DM management in HIV patients: the adoption of population health management to transform the chronic management of HIV, by Pacileo, Morando, Banks, Ferrara, Cattelan, Luzzati, Manfrin, Tozzi

Algorithms that help people

Deep learning allows several applications in medicine, such as in the case of patients with breast cancer. For such systems to be effective, however, they must be safe, equitable and needs-focused

by Oriana Ciani @

Research on Artificial Intelligence (AI) has made great strides in recent years, and applications and positive impacts are also evident in the medical field. To date, probably the most advanced AI-based solutions are in diagnostics, for example for the analysis of radiological images. Other areas of interest with high commercial impact are projects that concern the identification of molecules with therapeutic potential. The uses of AI for the improvement of the quality of life, and more generally of the quality of care and patient experience, are also multiplying.

With respect to the latter, tools like ChatGPT, which have reached millions of users in just a few weeks, may play a role. Some initial experiences document the use of ChatGPT to support bureaucratic activities, such as the release of medical certificates or reimbursement claims to medical insurance companies. Other uses concern the use in more complex activities, such as triage, that is, establishing priorities - on the basis of numerous and predefined criteria - in accessing services in contexts where resources are limited, or the possibility of providing information about therapies in progress and at the same time collect, directly from the patient, feedback that could influence the course of care.

However, tools such as ChatGPT, and in general AI solutions designed as consumer products and not for medical purposes, raise a series of ethical and legal issues. For example, in one of the research initiatives currently underway at the Government,

ORIANA CIANI
Associate Professor of Practice, Public Management and Policy Group, Health Economics & HTA, SDA Bocconi School of management

Health and Not-for-Profit Division of SDA Bocconi School of Management and the CERGAS research center on the healthcare sector, I recently found myself evaluating the possibility of inserting ChatGPT in a medical app, so it could draft a text designed to provide all the necessary information to a woman who is preparing to face a course of treatment following the diagnosis of breast cancer. In a very short time, the chatbot generated a plausible text that we are validating with the essential contribution of clinicians and health care professionals in view of publication, because the issue is whether the information thus generated is accurate and impartial enough to be applied in a sensitive context such as that of health protection.

Other challenges related to the use of ChatGPT in healthcare concern possible and, unfortunately, probable risks of privacy violations due to the communication of sensitive data that are potentially available for future disclosures to third parties. To date, there is also a lack of informed consent forms that are suitable for the multiple possible uses of AI.

Currently at CERGAS, we are working on another solution designed for women patients who are set to undergo mastectomy and subsequent breast reconstruction. CINDERELLA is a four-year European project, which aims to improve the level

THE RESEARCH CENTER

Over the last 40 years CERGAS has been one of the leading influences on Italian National Health System helping shape the way society thinks about health and health care. SDA Bocconi's Centre for Research on Health and Social Care Management (CERGAS) contributes to relevant and timely research within the healthcare sector, while simultaneously bringing new and advanced levels of empirical and analytical knowledge to the field. Nowadays, while keeping its focus on management of healthcare organizations, it is a multidisciplinary research centre that investigates a vast array of issues and topics concerning healthcare management, health policy, health economics, social and not-for-profit organizations.

of satisfaction, and therefore the psychological well-being and quality of life of patients operated for breast cancer through the automation of the aesthetic evaluation of the results of the surgery and the prediction of results before the intervention, so as to encourage an active and conscious participation of the woman in the choice of treatment.
Faced with a multiplicity of possible surgical alternatives, today it is difficult for a woman to be able to judge which type of surgery can provide the best aesthetic results in her own case. To automate and make these evaluations objective, for some years an algorithm (BCCT.core) prepared by INESC TEC of Porto and the Breast Unit of the Champalimaud Foundation of Lisbon, has begun to classify photographs of the chests of patients operated for breast cancer according to aesthetic outcomes codified on the basis of mostly geometric metrics. For example, the software is able to autonomously establish breast volume and size of bra cup starting from a photograph. It goes without saying that the images being used are anonymized after the consent was given by the women whose chests were photographed.
To limit any unfairness in the responses provided by the software, the image database used for training AI in CINDERELLA needed to be expanded to include a variety of women who were initially underrepresented in terms of skin tone, clearness and shape. Thanks to recent advances in so-called deep learning, which allows to optimize image recognition, this advanced software will be able to provide personalized predictions for individual patients about the results of the various surgical approaches.
Underlying everything, however, there must be one element: an effective and fair implementation of AI solutions, centered on people's needs and on adequate safety and ethical standards, which is a fundamental condition for potential of these innovations to emerge, by improving the quality of healthcare assistance and the well-being of communities. ■

Everyone is looking for a

Machines are very useful for finding similar patterns in huge amounts of data, but they lack abstract reasoning, a distinctive quality of the human mind. In the future, human-machine interaction will leverage their respective strengths, explains Igor Pruenster, Director of BIDSA, the research center that combines expertise in statistics, computer science and data science

nswers from AI. But the key lies in the questions

by Pietro Masotti @

Bringing together various disciplines such as statistics, computer science, mathematics and social sciences, and applying computational analysis tools to analyze large amounts of data and create models of complex phenomena. This is the research mission the recently established Bocconi Institute of Data Science & Analytics Institute (BIDSA) is pursuing. The founding director of the research center, Professor Igor Pruenster, Full Professor of Statistics and President of the International Society for Bayesian Analysis, does not share the current fad for using the word 'AI' as much as you can. "As an academic, I care a lot about the definition of concepts, and the word artificial intelligence is anything but precise; it is a very large box filled with many different concepts and tools," he explains. "The use of the notion of intelligence is potentially misleading if we have human intelligence as a reference because an AI system works very differently. Today we are mostly dealing with systems that perform precise prediction or classification tasks by applying statistical and mathematical techniques to datasets, in order to identify patterns or recurrences. It is not difficult to imagine that if you are able to store the online browsing data of a large portion of the population, you are capable of predicting the online behavior of each of us. We are much less unique than you'd think. Just type a question on a search engine to see that it auto-completes because many others before us have already asked the same question".

➜ What then are the truly innovative aspects of AI?

The questions we are asking are not new, but the scale makes all the difference. Analyzing thousands of ultrasounds for a pattern matching is a relatively simple operation for a computer, but one for which human intelligence is extremely fallacious. Even the most experienced doctor relies above all on his own experience which, however extensive, remains anecdotal. This does not imply that the doctor's intervention is redundant, but that they will have a scientifically more solid decision-making support. It is more difficult for machines to imagine something that is not yet there. Already Turing, as he was trying to decipher the Enigma code at Bletchley Park during World War II, dealt with the simple but formidable problem of estimating the probability of discovering a new species. In their context, this corresponded to the appearance of a new grouping of letters in the intercepted encrypted dispatches. Today the interest in this type of research has grown significantly. Just think of genomics and the importance of estimating the probability of sequencing new genes or that a new variant of SARS-CoV-2 will emerge. The common view at the moment is that computers will not be able to emulate humans in the capacity of abstract reasoning in the near future. The most profitable development of artificial intelligence will therefore be in human-machine interaction by leveraging their respective strengths.

IGOR PRUENSTER Full Professor at the Department of Decision Sciences and Director of the Bocconi Institute of Data Science & Analytics (BIDSA). Established in 2016 as a research center involving seven Departments of the university, BIDSA involves over 80 faculty members and has won ERC grants on topics such as the uncertainty quantification, the robustness and interpretability of models and algorithms for complex AI problems, aspects of cryptography and networks in medicine, but also for theoretical research projects in mathematics and statistics.

➜ Today AI is used in every context. So does studying it means you have to become an expert in everything?

The digital revolution has changed research in almost every area, primarily thanks to the wide availability of data and computational potential. As a consequence, a modern researcher must have solid background in computer science, statistics and mathematics, as well as being an expert in their own field, the so-called domain-specific knowledge. Even those who develop methods for learning from data today, which are typically labeled as AI, machine learning or data science, are required to be more interdisciplinary than in the past. However, the trade-off between depth and breadth hasn't gone away and there's a constant search for balance between the two. The risk of becoming all-purpose

researchers is always around the corner. For this reason, I always recommend PhD programs in Statistics and Computer Science to students interested in AI and data science who intend to undertake research careers in universities or industry. This allows them to acquire a well-defined specialization, a natural home, and constitutes a springboard for subsequently broadening their spectrum of interests. As it has happened to many of my students who now work for Google or Amazon, in the industry they are placed in composite teams whose members have very different skills to the point that interdisciplinarity is automatically achieved at the team level. It is no coincidence that inside BIDSA we have set up four units, the Artificial Intelligence Lab (ArtLab), Bayesian Learning Lab (BayesLab), Data and Marketing Insights (DMI), and Blockchain Initiative, to give our community of researchers and students the opportunity to aggregate in smaller groups around specific subjects, thus grafting a vertical model onto the horizontal one that already existed".

➜ ***Does the interdisciplinary approach shape a new subject and a new way of doing research?***

Yes. In this sense I think Michael Jordan, who is not the basketball champion but a Berkeley professor or rather, as the magazine Science dubbed him, "the Michael Jordan of Computer Science", is right. Michael, who is often visiting professor at Bocconi and gave one of the talks at the inaugural BIDSA conference, argues that we are witnessing the birth of a new branch of engineering that is based on data and learning. In fact, AI is based on well-rooted ideas: data, uncertainty, information, algorithms, inference, optimization. These are concepts that have been studied in depth by various disciplines such as statistics, applied mathematics and computer science. In the 1990s we began to mix these components and today's novelty, thanks to current level of computational resources, lies in doing so on a large scale with a direct impact on society. This will be a new branch of data engineering and learning, in the same way that chemical engineering is founded on chemistry but analyzes its application in industrial processes. The other relevant novelty is that, for the first time, we are starting from data provided by people about people. To quote Jordan again, it will be the first branch of engineering centered around people.

➜ ***What are the frontiers in AI research?***

Looking at the fundamentals, I expect that issues such as uncertainty quantification, robustness and interpretability to gain more and more prominence. The quantification of uncertainty is a cornerstone of statistics but, as it conjoined with computer science, this was a bit lost. There is this somewhat naïve idea that, as the quantity of data increases, uncertainty vanishes. In most cases this is not the case. Probably most have heard of ChatGPT; think how much more reliable it would be if it gave replies associated with a measure of confidence about the various parts of the answer given. The second theme is the robustness of the modeling. Would the results change by changing the parameters of the model or perturbing the data? At present, the differences would be significant and this is a limitation of the models. The third factor is the interpretability of results. One of the major shortcomings of modern deep learning models is that they provide predictions that are often accurate, yet it remains very difficult, if not impossible, to understand how they arrived at their conclusion. Knowing that you have optimized an objective function is enough to decide what film to suggest a user, but it is not always enough for deciding on a medical therapy. Understanding is always important, in my opinion.

➜ ***What are the fields of application of all this that you explore within BIDSA?***

We have several ERC-funded projects, one on the quantification of uncertainty and others at the intersection of robustness and interpretability. Further projects deal with medical networks, language processing technologies incorporating demographic factors as well as pure mathematics on optimal transport, an abstract problem that has numerous applications today. And then there is a group of economists studying decision theory, a subject that does not yet play an important role in AI but which will increasingly come to the fore. Today, in fact, the focus of research and applications of AI is concentrated on data and on the best decision for a single individual, while federated learning, i.e. the interactions between individuals and between them and the context, is still little explored. ■

THE RESEARCH CENTER

The Bocconi Institute for Data Science and Analytics (BIDSA) was established in 2016 to promote and facilitate data-driven research at Bocconi University. It represents Bocconi's timely answer to the increasing interest in Data Science and the rapidly growing demand for data scientists across all industries. By acting as a central hub of data-science research and education at Bocconi, BIDSA allows for effective interdisciplinary collaboration and cross-fertilization across data-related research areas.

How to improve tax

According to the algorithmic approach employed by Bocconi researchers, measured tax evasion grows to 38%, using data from the Italian Revenue Agency. And this is just one of many examples that show promise on the interaction between machine learning and fiscal accounting

by Marco Battaglini @

Tax authorities routinely collect enormous datasets on taxpayers and should use them efficiently to audit taxpayers. Are there margins to improve the efficiency of tax auditing? Judges in the U.S. and other countries evaluate hundreds of thousands of defendants for jail or release decisions before trials. These decisions are often left at the discretion of local judges. Can Machine Learning (ML) help them keep consistent criteria and avoid mistakes and bias?

It may be tempting to see ML problems just as engineering challenges for computer scientists, and we should obviously not underestimate the difficulties associated with designing efficient algorithms. But in many applications, the technical design is the least of the problems. The real complications are in interpreting the results and translating them into appropriate policy recommendations.

Consider the case of a tax authority aiming to design an auditing plan. To be effective, the plan needs to predict the identity of likely tax evaders, a task for which ML algorithms are ideal. In short, an algorithm would use historical data to select the variables that best predict evasion and combine them in a score that could be used for the choice. If this is what the authority does, however, we have a big problem. This procedure would use only outcomes of files that have been endogenously selected for treatment by the authority in previous periods. Any potential bias in the selection process would be inherited by the data. If we can't control for all variables used to drive the selection, the result may be biased decisions, even (and indeed especially) if the algorithm is efficient. For example, imagine that the tax authority selects audits relying also on variables unobserved by the algorithm that are good markers of compliance (perhaps variables that are not stored in official datasets). If we ignore these markers, the algorithm may recommend replacing audited files with unaudited files associated with the unobserved marker. This would lead to an overestimation of the possible improvements generated by ML. Part of the problem here is that we observe the outcomes of (endogenously selected) audits, but we do not observe the outcomes of tax files that have not been audited.

MARCO BATTAGLINI
Full Professor of Political economics, Bocconi University

What can be done to address these issues? There is an apparently simple solution: do not use historical data to train the algorithm; instead, use carefully randomized data, as in randomized trials for drug testing. If this were possible, then we would be sure that no bias is present in the training data, and the predictions would be effective. But this solution is often impractical and not what is generally done.

In a recent work, we have proposed a methodology to correct potential biases in the design of tax auditing. We exploit two features of datasets generated by tax authorities: first, only a tiny fraction of files are audited, typically in the single digits; second, files can be audited for up to five years, so there are many unaudited files for which we can assess the true potential since they are eventually audited in later years. These are files that are unintentionally passed over but are then randomly selected for audits at a later stage. Simplifying a bit our approach, we can use them as counterfactuals to evaluate whether we can improve auditing using ML. Specifically, we can use ML to select historically audited files with low potential and replace them with files with good potential, restricting the replacement pool to files for which we eventually see the outcome. This is a conservative policy but one that is arguably immune from the problem that the outcome of the counterfactual is unobserved and may be overestimated. Using data from the Italian Revenue Agency, the analysis suggests that there are indeed large untapped improvements to be gained: replacing the 10% least productive audits with an equal number of taxpayers selected by our trained algorithm raises detected tax evasion by as much as 38%.

Economists have only recently started to think systematically about these issues, which blend ML design with issues more typical of economic research, such as causal identification and optimal policy design and evaluation. There are difficult and sometimes unsolvable tradeoffs in the problems; solutions, like the one described above, may only partially address concerns. Nonetheless, these are promising topics of research because even marginal improvements may be pivotal, and economics has a lot to contribute. ■

THE PAPER

Refining Public Policies with Machine Learning: The Case of Tax Auditing, by Marco Battaglini, Luigi Guiso, Chiara Lacava, Douglas L. Miller, Eleonora Patacchini

audits

Geopolitics

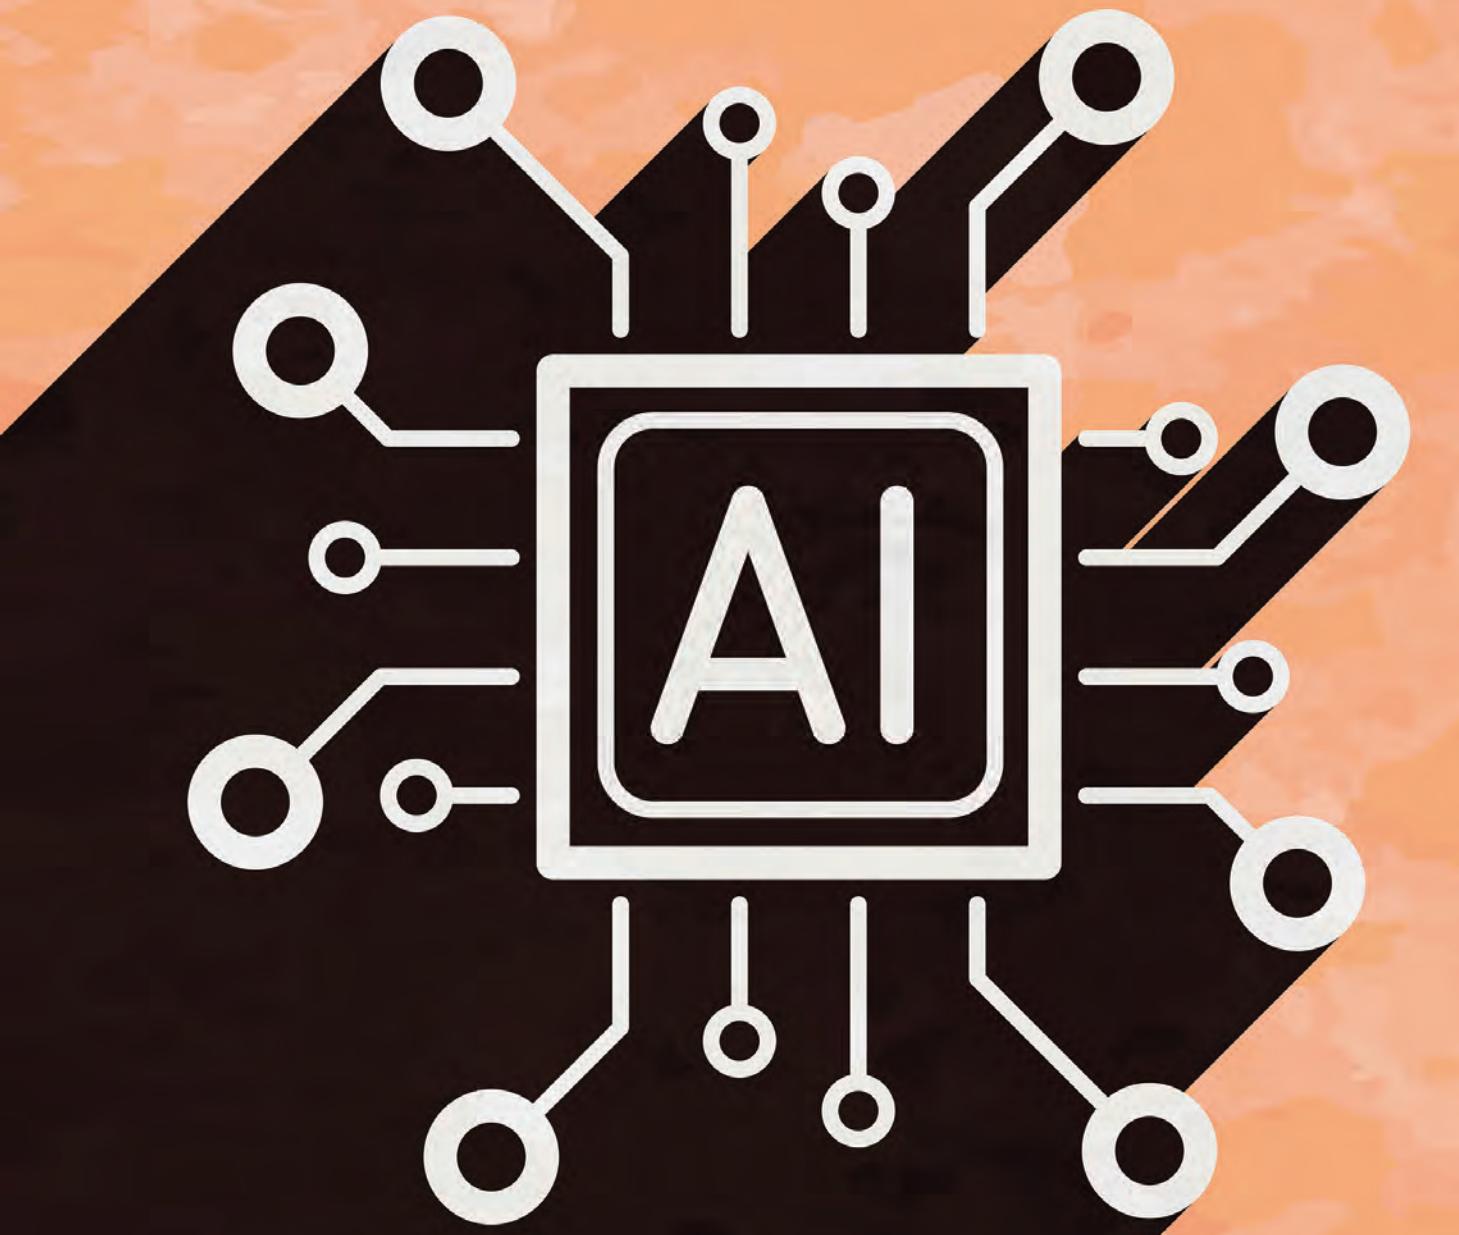

in the age of AI

The contest between America and China, won by the US for now, increasingly plays on the relative ability to govern this dual-use technology and its evolution, and control the global value chain behind the production of semiconductors. But for how long will that be the case?

by Andrea Colli @

The enormous potential in terms of applications in almost all areas of human activity is making AI one of the new "general-purpose technologies", and one of the industries which both private actors and, even more, governments, consider as a strategic priority.
AI is indeed a perfect example of the so called "dual-use technology", that is of a multi-purpose technology which, once developed, can be used both for civilian and military (included security) purposes. As the current conflict opposing Russia to Ukraine is showing, AI is deeply changing both the strategies and the tactics of modern warfare. AI provides a ready-to-use massive amount of information which allow the troops on the ground to target more easily the enemy - a sort of "Uber for the artillery", as one official effectively put it.
Not surprisingly, AI has become a key component of great power competition, and a major strategy for enhancing national competitiveness and security. According to experts in the field, at the moment China is one year behind the US in terms of AI capabilities, but the US may soon lose its leading position. In order to speed up its AI processing capabilities, for instance, China launched in 2017 the New Generation AI Development Plan for 2030. The document is clear about China's intentions to become an AI innovation center by the beginning of the next decade, being the new technology essential for strengthening the country's military and economic capabilities. These goals will be achieved by means of a close coordination between the private sector and the government, directly or by means of state-owned "champions" in the field.
In the meantime, both the EU and the US have, almost simultaneously, issued "chips acts", which aim at both re-shoring chips production and improving their domestic technological capabilities in the field.
In the jargon of contemporary geopolitics, AI is considered one of the most relevant "realms of power", a modern version of the "pivot", the control over which assigns a de facto position of enduring dominance.
Not surprisingly, in a recently published National Security Strategy document, the US administration openly mentions AI as one of the fields in which it is mandatory to ensure that "strategic competitors cannot exploit foundational American and allied technologies".

ANDREA COLLI
Full Professor
of Business History,
Bocconi University

In order to function, AI applications ultimately need machines, computers able to store and process an enormous (and growing) amount of data. Much (if not all) depends on the basic unit which fuels the machines' computational power, i.e. chips which have to be specifically designed and manufactured to be used in AI data processing. AI demand for customized semiconductors offers immense opportunities to private companies capable of producing chips at cutting-edge technological level in terms of transistor density, but also customized in order to perform the specific tasks required by AI systems.
So far, the global value chain governing chips production has been an example of borderless efficiency. Integrated circuits designed in the US and Europe are produced in South-East Asia employing US-made machine tools and sophisticated equipment. Chips for AI machines are not an exception; their technological sophistication meant that a significant portion of their production is located where there are the most sophisticated "foundries", capable of manufacturing chips with a density measured in nanometers, typically in Taiwan and South Korea.
The present concentration of technological capabilities in leading-edge AI chips in US and "US-friendly" countries gives the opportunity for the incumbent power (America) to weaponize semiconductor technology in order to keep the emergent challenger (China) at bay, through a mix of strategies ranging from export restrictions, to the so-called "foreign direct product rule" (which de facto prohibits the export of products incorporating US design or technologies), including moral suasion in order to avoid the acquisition of semiconductor companies by Chinese investors. ■

by Massimo Morelli @

It is clearly very important to have an estimate of conflict risk. In fact, especially in Africa, Asia and Middle East, the country-risk evaluated by investors is affected by the risk of conflict. Many civil conflicts, especially in Africa, can be thought of as ethnic conflicts - typically involving the government against an ethnic group. Thus, estimating conflict risk requires as an ingredient to have an estimate of the military strength of each ethnic group. This is due to the fact that, as established in Herrera et al (2022) and Morelli et al (2023), the difference between the military strength and the political power of an ethnic group (relative to the corresponding government) is a crucial trigger of civil conflict.

However, finding measures of military power at the ethnic-group level is extremely challenging. Defining relative military power of an ethnic group as the probability of winning a conflict against the corresponding government, machine learning can be used to obtain such an estimate even for groups that never experienced conflict. In Morelli et al (2023) we use an extended sample of conflicts in Asia and Africa combined with a rich set of ethnic group-level and country-level variables to infer the probability of victory for all potential conflicts between every ethnic (rebel) group and the corresponding government. The advanced learning algorithm chosen has the objective of selecting the relevant predictor variables plus an important cross-validation objective.

We use a stacked ensemble learner, which is a method that combines multiple learning algorithms. Stacking, or Super Learning, is a procedure that aims to find the optimal combination of prediction algorithms. Generally, the cross-validated error of the learner is simply the average error made on each N prediction. The various learning algorithms we use, random forest, gradient boosting machine, and generalized linear models, all have pros and cons, and the stacking procedure improves the risk forecasting performance. To give an idea, previous works on approximations of the probability of winning of an ethnic group were using relative population sizes or nightlight proxies for economic strength. With respect to such previous approximations, we estimate that the machine-learning algorithm improves precision by almost 20 percent. Twenty-percent greater accuracy is quite meaningful in almost any field.

More generally, in social sciences we are often interested in estimating the probability of success of an agent's action as a function of the agent's characteristics, and in the observations available only a fraction of the agents studied did take such actions. Hence using only individual actions and outcomes suffers from a clear selection bias problem. Machine learning allows to estimate the probability of success even for agents who never undertake a given course of action, on the basis of all available data about the set of characteristics of all agents. ■

MASSIMO MORELLI
Full Professor of Political Science and Economics, Bocconi University

THE PAPER

Power Mismatch and Civil Conflict: An Empirical Investigation, by Massimo Morelli, Laura Ogliari, Long Hong

Thanks to the stacking technique which combines predictions from multiple models, Bocconi researchers have improved by 20% the measurement of the risk of conflict between ethnic groups in Africa and the Middle East. A particularly relevant result, since the assessment of political risk guides the choices of investors, who are always wary about internal strife in a country

Calculating the risk of


conflict

by Carlo Baldassi @

"What I mean is that if you really want to understand something the best way is to try to explain it to someone else. That forces you to sort it out in your own mind. And the more slow and dim-witted your pupil, the more you have to break things down into more and more simple ideas. And that's really the essence of programming: by the time you've sorted out a complicated idea into little steps that even a stupid machine can deal with, you've certainly learned something about it yourself. The teacher usually learns more than the pupil. Isn't that true?».

For a few years now, during the first lesson of my basic programming course, I show a slide with the above quote taken from Douglas Adams (Dirk Gently's Holistic Detective Agency, 1987). I'm not entirely sure if I'll be using it much longer: the meaning of "program" could quickly diverge from that of "writing code". Machines are no longer that stupid anymore.

The first test I did once I had access to ChatGPT was to submit one of the exercises of the final exam of the course that had just ended. I had invented the exercise and ChatGPT could not therefore know it. Compared to the real exam, I initially asked it for a simplified version, but (paradoxically) I used a much more approximate and colloquial description than the one I had written down for the students, as I would have done if my interlocutor had been an expert programmer with whom certain details could be left unsaid. The machine figured out exactly what it was supposed to do and produced fully functional code in seconds (the students have about 20 minutes at their disposal). I wanted to go to the complete version of the exercise: I therefore asked him to modify the code in order to make it more efficient (continuing to express myself as if I were dealing with a human expert). This time the answer contained an error, although it went in the right direction. Having pointed this out, ChatGPT produced flawless final code.

Naturally similar experiences, of so-called "program synthesis", have been made by many programmers in recent months, and immediately raise obvious questions about the future of programming, the labor market in this sector and, as a consequence, the whole of society.

CARLO BALDASSI
Associate Professor of Computer science, Bocconi University

ChatGPT is not the only or even the best program synthesis system out there. More specialized systems (but also obtained through the training of the same type of neural architecture, called Large Language Models) had already demonstrated to be at the level of an average human expert in programming competitions for specialists. The best-known case is AlphaCode, presented by DeepMind in February 2022.

At the moment these systems are however able to write only short programs, which solve specific and well-defined problems. Most specialists believe that their role will be to allow laymen to program simple systems without knowing how to write code, and experts to drastically increase their productivity in the development of complex systems that require overview (and already exist today some programming assistants, for example Copilot from Github and Code Whisperer from Amazon, able to suggest entire portions of code). In both cases, the meaning of the term "programming" would be transformed: from writing code in the first person to so-called "prompt engineering", i.e. knowing how to ask the right requests to an AI algorithm, in the context of a collaborative dialogue. Collectively, the societal consequences will be dramatic and potentially disruptive.

This forecast could even be considered cautious: as recently as 2017, in the ponderous volume *Program Synthesis* (by Sumit Gulwani, a Microsoft researcher, and co-authors), which presented the techniques developed in the previous decades to address what the authors defined «the holy grail of Computer Science», only a few pages were devoted to AI-based approaches. Progress has therefore been very rapid and unexpected. Obstacles that seemed critical have been overcome, and there are currently no signs that progress will slow down. We may very well be just getting started.

In any case, even if coding eventually becomes a hobbyist activity, it will still be a great way to train yourself to dissect complicated ideas into their elementary constituent steps, and in so doing understand what we are trying to do, or make an artificial assistant do. ■

The future lies in prompt engineering, which is the ability to know how to ask the right questions to an AI algorithm whose task is to write code in collaboration with a more or less expert human agent

The essence of progr

amming

Ok, but is the

More and more companies, especially on the web, are adopting algorithms to fix the pricing of their products. Some empirical cases have shown that this benefits retailers (+38% for German petrol stations) and this has opened a legal and political debate on the collusive potential of this approach but also on the possible positive impact. However, the real issue remains the data and their access to correctly train the algorithms

price right?

by Francesco Decarolis @

FRANCESCO DECAROLIS
Full Professor of Economics, Bocconi University

Companies increasingly delegate their prices to algorithms trained on data on customer preferences. In some cases, these algorithms rely on a machine learning process to develop sophisticated pricing strategies that respond to both customer and competitor behavior.

Artificial intelligence algorithms are gradually coming to occupy a prominent place in the pricing methods of various companies. This is particularly true in more technological and web-operating companies, but more generally any sector where there is a large amount of detailed data available on demand and potential competition lends itself to this type of evolution. Fueled by this data, AI algorithms are designed to learn which pricing strategies will yield the greatest returns based on a trial-and-error learning process.

But what can be expected from the growing trend to delegate pricing decisions to algorithms? Empirical evidence on the impact of AI algorithms is still rather scarce. However, there are cases that have already received considerable attention. For example, in the German gasoline retail market, algorithmic pricing software has become widely available since around mid-2017. To understand what this means, we have analyzed detailed data regarding the phase before and the phase after the advent of algorithms. The surprising result was a steep increase in prices: the introductionn of algorithmic intelligence in pricing at German gas stations appears to have coincided with an increase in margins of up to 38%!

A problem with this type of study is that both the exact timing of the adoption of the algorithms by the various gas stations and the exact details about the type of algorithms used are necessarily unknown to outside observers. For this reason, in the study of the German gasoline market, the authors identified which petrol stations had adopted algorithmic pricing software and when they had done so through a statistical procedure. This is an interesting idea because it takes advantage of what should be the peculiarities of an algorithmic pricing mechanism. In particular, compared to the prices established through human intervention, the prices established by the algorithms should differ in: the number of price changes made during the course of the day, the average size of price changes and the response time of the price update of a given station compared to a competitor's price change. Precisely these measures are the ones on which pricing software companies advertise the ability of their algorithms to generate positive impacts for distributors.

But how and why does AI induce higher prices? In the case of the German gasoline market, margins have been found to rise gradually, suggesting that algorithms need time to train themselves and converge towards tacitly collusive strategies in which AI determines less aggressive pricing than would be expected under normal market competition. That is, the algorithms learn pricing strategies that are more cooperative and lead to higher price levels. If expected profits for companies are larger, economic harm for consumers is just as clear.

This kind of result explains the vigorous legal and political debate on the collusive potential of AI applied to pricing that is currently underway. In particular, the potential of using algorithms as a means to facilitate explicit or tacit collusion has been a popular talking point among antitrust authorities, business organizations and competition law experts in recent years. Above all, there is strong skepticism about the ability of current competition law to address any tendency of AIs to induce super-competitive pricing as the law is designed to punish explicit agreements between companies.

There is still no general consensus on what AI will mean for pricing systems and what the ultimate consequences will be for companies and consumers. Given the lack of empirical evidence available, current research is mostly directed towards theoretical and experimental methods to evaluate the impacts of AI. The results are interesting but full of ambiguities. If some studies underline the collusion in prices deriving from learning reciprocal pricing strategies, others indicate on the contrary that AI, by allowing companies to better predict demand, pushes them towards more aggressive pricing strategies, thus lowering prices and making collusive behavior less likely.

Finally, we must not forget the fundamental role of data, real fuel in the engine of AI algorithms: whoever controls the data controls the type and effectiveness of the AI algorithms that have to work on this data. This very aspect is at the center of a recent study I conducted with Michele Rovigatti (Bocconi) and Gabriele Rovigatti (Bank of Italy). The results, presented for the first time recently at a conference at Yale University, illustrate the problem of a digital platform (such as Google for example) which, through the type, quality and frequency of the data it releases to advertisers active in the auctions where it sells advertising space, it influences their choices in terms of which AI algorithms are used in the bidding process and with what effects on final prices.

So in conclusion, the jury is still out on how AI will impact pricing and consumer welfare. ■

Recycled, augmented

To work, AI uses human-generated data (this is why it is recycled), managing to learn from them and create new (augmented) knowledge as per Network Science, thus supporting decision-making processes. But to arrive at human intelligence it is probably necessary to arrive at the synthesis of the two

by Daniele Durante and Omiros Papaspiliopoulos @

Although Artificial Intelligence might seem a recent trend, this idea was in fact introduced in the late 1950s. Since then, public perception, research activity and funding have undergone spectacular ups and downs, including the famous AI winters. Considering the most successful AI achievements of the last decade, albeit fascinating, the original "aspiration of realizing in software and hardware an entity possessing human-level intelligence" has been progressively reframed into the more tangible design of machines that can perform, with impressive accuracy, several human tasks such as text and image generation, object recognition and motion, without necessarily requiring a full understanding of the source mechanisms that make these tasks possible in humans. This is accomplished by leveraging massive amount of data produced and labeled by human intelligence for training Statistical Machine Learning algorithms to succeed in these tasks.

Recalling the recent AI achievements (e.g. ChatGPT), the beauty of this result is that a secondary output of the source mechanisms underlying human intelligence, i.e., human-generated data, is often enough, if combined with advanced algorithms, to create an intelligence capable of mimicking the tasks underlying the creation of the original data. As such, this form of intelligence could be understood more precisely if referred to as "recycled intelligence" . This definition is even more explicative if we consider that not only the input data are a product of human intelligence, but also the powerful computer architectures currently used to process these data have recycled the invaluable and partially-unacknowledged contributions of the early 20th century human computers. Namely talented humans that, in the absence of powerful computing infrastructures, were asked to perform complex computer-type calculations. A sort of reverse-AI requiring humans to mimic computers, rather than the opposite, which however set the stage for the rise of modern computing technologies.

DANIELE DURANTE
Assistant Professor of Statistics, Bocconi University

OMIROS PAPASPILIOPOULOS
Full Professor of Data sciences, Bocconi University

Viewing the current forms of AI in terms of recycled intelligence helps to contextualize its achievements and clarify its differences from another fundamental perspective of learning from data, namely what can be defined as "augmented intelligence". Albeit sharing a common emphasis on data, models and algorithms, this perspective prioritizes the creation of new knowledge about the generative processes and causes of the phenomena at a population level from the analysis of multiple data sources. Hence, it augments human intelligence by coordinating and combining different disciplines and field expertise to answer challenging scientific questions that a single human or a single community would not be able to address autonomously. This perspective, which is a distinctive feature of Data Science, has at its core a form of hypothetical reasoning to design generative models that can produce knowledge, even in the absence of a large amount of training data.

A timely example is Network Science, which is the discipline that analyzes complex connectivity structures among a set of entities. There, the sample size is often one, since only a single network is typically observed. To cope with this problem, structured and generative statistical models, incorporating also prior information from field experts are required. Working in this direction with colleagues from the BayesLab and PhD students from the Bocconi program in Statistics and Computer Science, we have been able to generate new knowledge on the structure and function of brain networks, and on yet-undiscovered pyramidal architectures in criminal networks. In the context of criminal networks, we are now exploring new intersections among factor models from Social Sciences, phylogenetic trees from Biology, field-expert knowledge in Criminology and advanced inference methods from Statistics to further obtain unprecedented reconstructions of organized crime evolutionary histories rooted in the complex interactions among criminals.

Recycled and augmented intelligence have, and will, progressively reshape our way of intending, producing and analyzing data to support complex decision-making tasks. Their synthesis might be the solution to achieve the very original AI aspiration. ■

but not yet human

In the beginning it was Cambridge Ana

Bots and generative AI can undermine democracies by manipulating voters into believing an idea is popular and shared. Just like in product marketing campaigns

by Gaia Rubera @

In December 2012, a smiling Bono prophesied "Big Data will save politics" on the cover of a famous technology magazine. The illusion was that direct access to information through social networks would make it more difficult for politicians to manipulate voters' thinking. In 2018 the same magazine came with a concerned headline: "Why AI is a threat to democracy - and what we can do to stop it." What happened and what could happen?
In the beginning it was Cambridge Analytica which, thanks to a Machine Learning model, was able to deduce the psychological profile of an individual from his/her activity on Facebook. With this model, Cambridge Analytica influenced political elections in dozens of countries, showing different kinds of ads for the same political issue, personalized on the basis of the psychological profile of individual users. For example, with differentiated ads supporting the right to bear arms in the US, emphasizing either the need for security to the neurotic crowd and of the respect of custom and tradition to conservative voters. Cambridge Analytica was founded by Robert Mercer, billionaire and major donor to the Republican party and the far-right site Breitbart, on whose board sat Steve Bannon, Trump's strategic adviser for some time and sponsor of Italian politicians.
Anyone who studies marketing knows that to persuade an individual to buy a product, you need to make him/her believe that the product is popular. Similarly, to push a voter to embrace a political idea, it is necessary to convince him that a large proportion of the population already shares that idea, repeatedly bringing him into contact with other individuals with the same idea. On social networks, this pressure is exerted by bots: applications programmed to automatically perform certain tasks without human interaction. The purpose of bots is to amplify a certain message to slowly but systematically orient public opinion in a specific direction. Research has repeatedly identified networks of bots controlled by a single puppeteer that are activated around politically relevant events. For example, a network of bots was identified to be very active before the US election of 2016, which went silent until three days before the 2017 Macron-LePen ballot, when it suddenly became active, to amplify the #MacronLeaks (disinformation) campaign - about alleged scandals embroiling Macron - and then went again dormant for months.

GAIA RUBERA
Amplifon Chair in Customer Science, Full Professor of Marketing and Head of the Marketing Department, Bocconi University

The latest developments in the field of Artificial Intelligence have made it very difficult to identify, and therefore counteract, such bots. For example, until a few years ago, bots used faces of real people, stolen from photographs on the web, as their profile picture. A simple Google search identified the theft and the bot. Today, generative artificial intelligence makes it possible to create non-existent but completely realistic human faces, making it impossible to identify a bot from its profile photo. Also, before these bots were limited to sharing messages from their creator, since it was difficult to create original ones. It was therefore relatively easy to understand that you had stumbled upon a bot. Unfortunately, current models of artificial intelligence are now able to think strategically and express their thoughts convincingly. A recent study has shown that messages written by Artificial Intelligence are as persuasive as those written by humans. It is easy to imagine the consequences of this technology in the hands of those intent on destabilizing democracies. ■

lytica...

Does the media prop

The answer comes from computer vision and text analysis techniques used in a Bocconi research study which highlights how modern machine learning tools can broaden the scope of social science

by Carlo Rasmus Schwarz @

The media have a powerful influence on how people view and treat different groups based on gender and ethnicity, as news stories shape people's stereotypes, beliefs, and ultimately behaviors in areas such as education, family, and politics. It is, therefore, essential to understand how the media, and in particular news outlets, depict different groups and whether stereotypes drive these depictions.

Challenges arise for researchers in this regard, as stereotypes are encoded in the media in multiple ways, and not all are readily apparent or measurable. Recent research has tried to tackle this question by using natural language processing techniques that allow studying the associations between words that journalists use in their texts.

What has so far been overlooked is the importance of images when it comes to the slant and stereotypes in media outlets. In a recent research project, my co-authors Elliott Ash, Ruben Durante and I study visual stereotypes in newspaper images by making use of artificial intelligence techniques and, more specifically, recent advances in the field of computer vision.

In particular, we use a custom-trained deep-learning mode that can recognize the identity features (such as gender/race/ethnicity) of people that are shown in the images. In such a way, we can automate our analysis instead of having to code hundreds of thousands of images by hand. A further advantage of this approach stems from the consistency of the classification, which cannot be guaranteed for any human coder.

In our analysis, we focus on two major U.S. news outlets: the New York Times and Fox. In total, we analyzed over two million articles published on the web editions of the two outlets between 2000 and 2020, of which 690k are accompanied by an image.

CARLO RASMUS SCHWARZ
Assistant Professor of Economics, Bocconi University

One important finding of our paper focuses on occupational stereotypes in newspaper images. In other words, we analyze if newspapers perpetuate common stereotypes about the occupational choices of specific groups, e.g., White men working as managers.

Using our computer vision approach in combination with text analysis techniques, we can show that news article images display gender and racial stereotypes in the sense that jobs which stereotypically are considered "female" or "black" are more likely to be represented by an image showing the respective identity group.

An occupation is considered stereotypical if a higher relative share of this identity group works in an occupation. For example, even though Blacks, due to their size in the overall U.S. population, occupations like "mail processor" might nonetheless be stereotypically black as a higher share of Blacks works in this occupation compared to other identity groups. To provide another example, the most stereotypically female job is "secretary".

Strikingly, our analysis controls for the true occupation shares in each profession, e.g., the share of "mail processor" who is Black. This allows us to disentangle stereotypes from pure differences in the representation of identity groups across occupations in the U.S.

In further results, we also find that images of Black faces are associated with more negative sentiment in the associated article text, with a stronger negative association for Fox News compared to the New York Times.

As such, our results highlight the potential use cases of A.I. and computer vision tools for the study of social science research questions. Furthermore, due to the ubiquity of images in business, politics, and social media, modern computer vision tools significantly broaden the scope of the analysis in the social sciences. We believe that through the broader usage of computer vision technique, social scientists will be able to identify and study biases that have been implicitly encoded in images. ■

THE PAPER

Visual Representation and Stereotypes in News Media, by Elliott Ash, Ruben Durante, Maria Grebenshchikova, Carlo Schwarz

agate stereotypes?

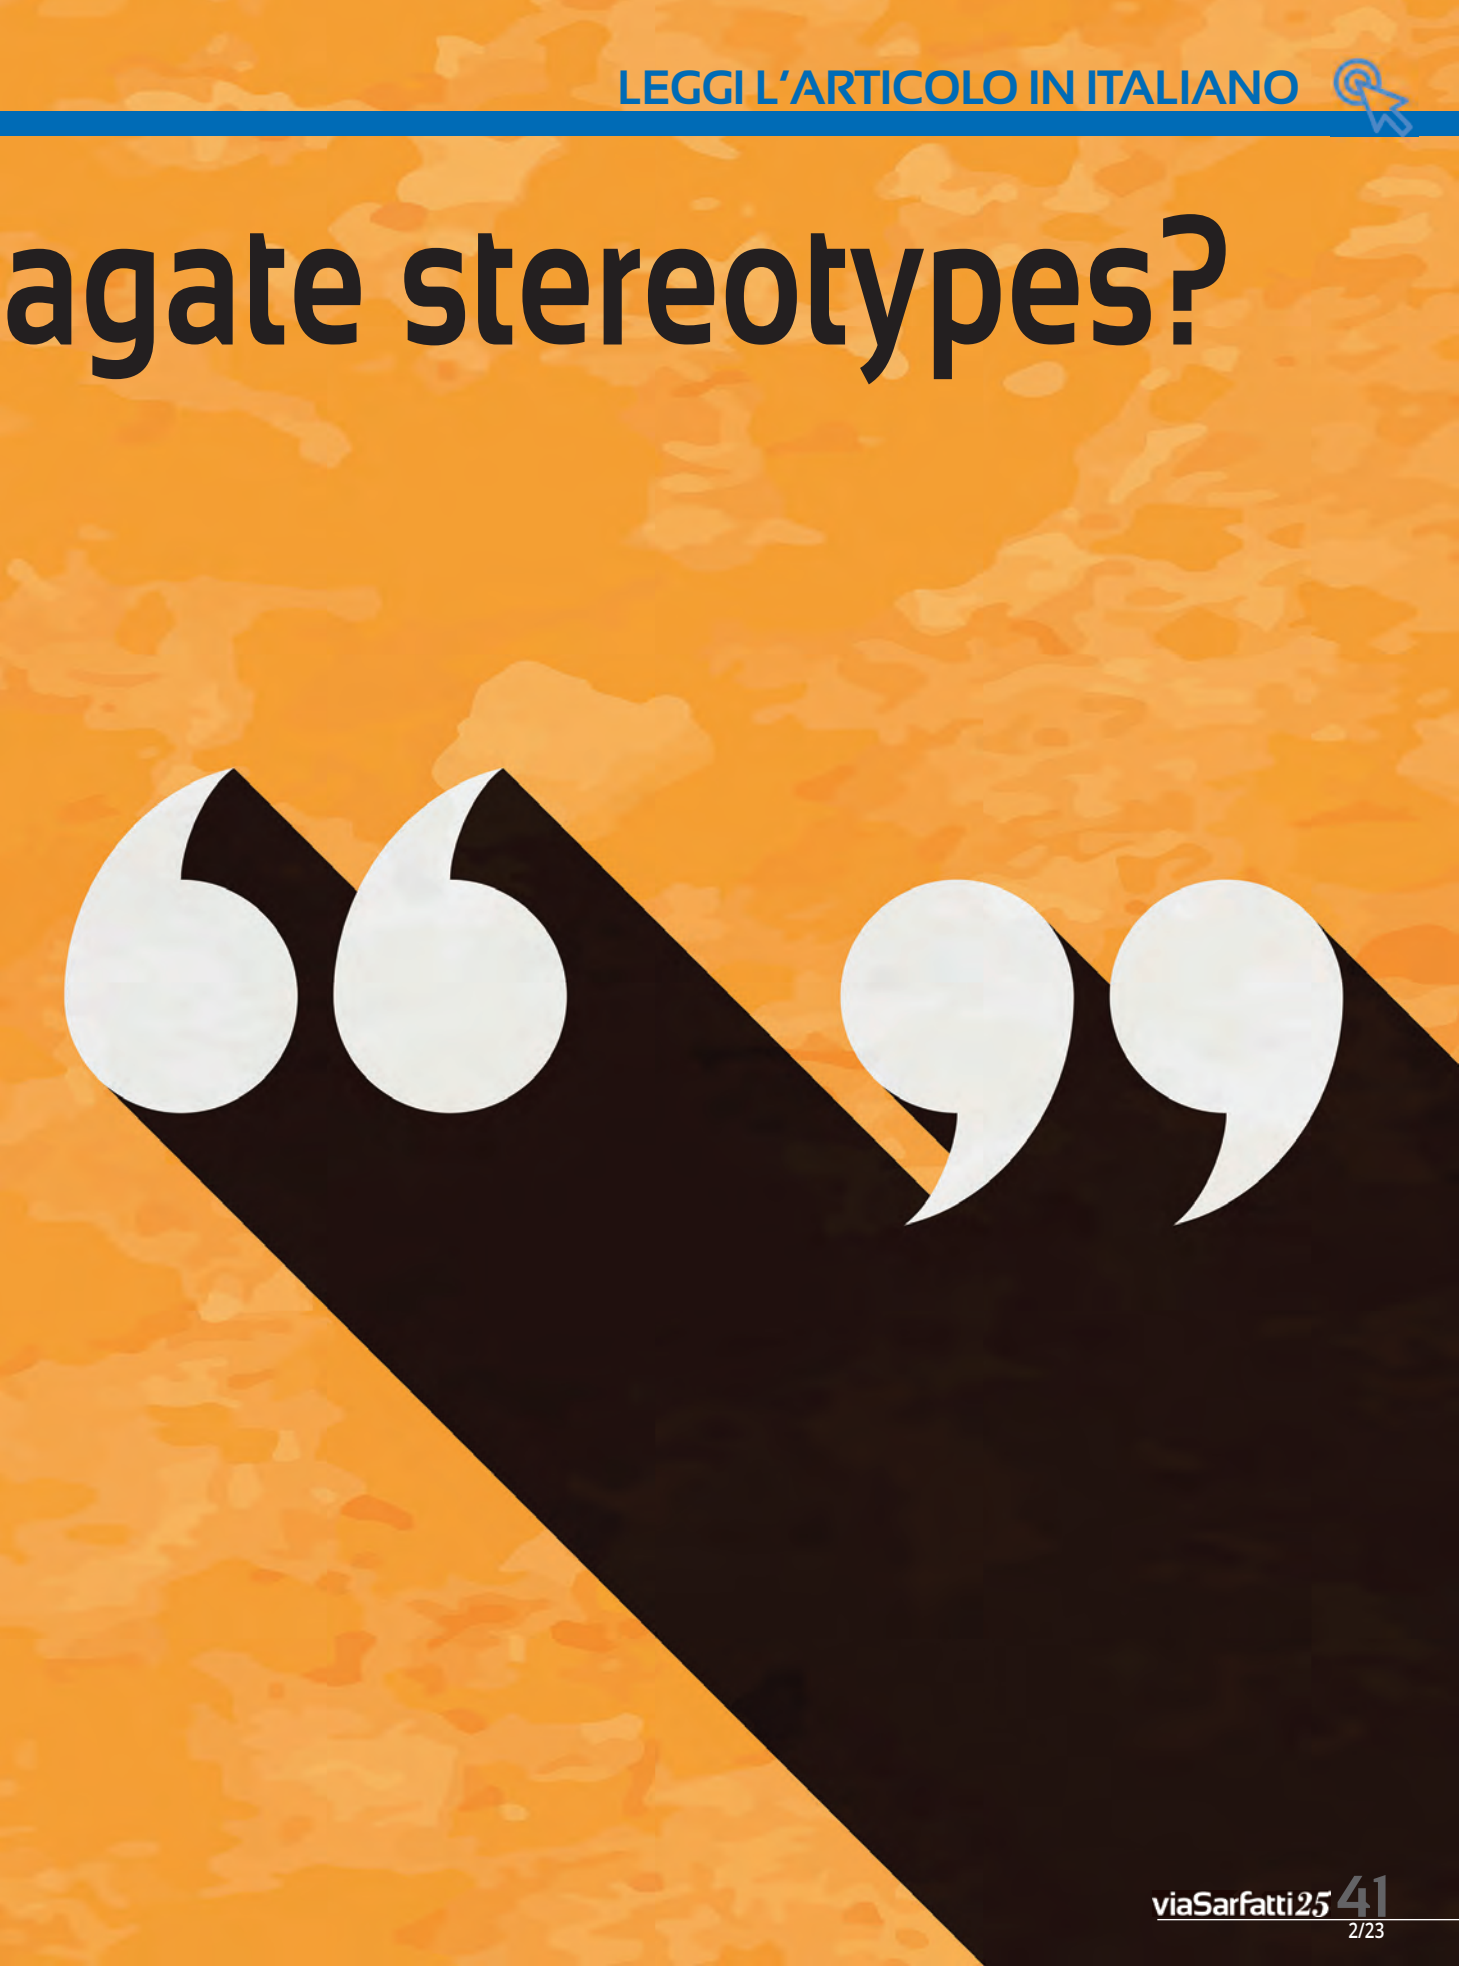

Not ju

Thanks to foundation models, i.e. basic models that are pre-trained using large datasets, it is increasingly possible to customize financial proposals according to the characteristics of individual investors. But that's not enough. To be credible, proposals must be explainable. And therefore transparent

by Claudio Tebaldi @

There is a large and established literature on the difficulties that non-professional investors encounter in orienting their financial decisions. Financial literacy initiatives aimed at disseminating the knowledge gained at the academic level to let it flow to even the most distracted investors, and those segments of the population who are most vulnerable to the cost of unaware financial decisions, are now numerous and institutionalized.

However, the experience gained in these contexts shows the existence of a gap that still needs to be filled between the context in which the quantitative allocation criteria formulated in the scientific field and the heterogeneity of the problems and constraints that affect the choices of actual investors.

Analyzing, for example, the commercial services of Robo Advisory, an automated advisory service for clients, it has been verified that the allocation rules that are implemented are often determined by a priori specific investment rules, which do not respond very well to the individual

st investment advice

characteristics of investors and the differing contexts in which these choices are made.
The possibility of covering the last mile, by aligning the actual practice of financial decisions with the prescriptions that are the outcome of research, seems to be within reach for a particular class of learning models, the so-called foundation models, i.e. basic models that are pre-trained using large datasets. Their use as decision support tools is a recent introduction and promises substantial advantages, first of all the possibility of making choices by interacting with the decision-maker and actively contributing to the constitution of the training dataset.
The basic idea is quite simple and in fact follows the logic pursued by researchers to solve limitations of the same nature. In the scientific field, we have been working for some time using anonymized, homogeneous and internationally harmonized datasets that collect the variables necessary to reconstruct the balance sheets of individuals and family households, and associate them with portfolio allocations and investment choices. This data collection expands the information assets that can be employed to zero in the factors that underly individual choices, thus helping to define investment

CLAUDIO TEBALDI
Associate Professor in Finance and Scientific Director Algorand Fintech Lab, Bocconi University

selection criteria that are more consistent with observed investor profiles.
Similarly, pre-trained models can generate scenarios and propose choices based on criteria that have a higher likelihood based on the information available in the datasets that collect investor information. The potential advantage that derives from this is the possibility of actively implementing rationality prescriptions and formulating investment proposals that are more appropriate to the specific market conditions, constraints and preferences of investors. The main obstacle to the systematic use of this approach derives from the known difficulties in maintaining the correctness check of self-produced output. In fact, the systematic use of artificial intelligence models must always be confronted with the need to operate according to adequate verifiability and correctness criteria of the contents generated. The scientific challenge is therefore linked to the possibility of formulating, together with the financial decision proposals, also an adequate explanation of the reasons that motivate them, according to the criteria associated with so-called XAI, eXplainable Artificial Intelligence. After all, the credibility of an advisor, human or cybernetic, is still built starting from the same principle, transparency. ■

The charge of Robo-Advisories

Why banks must invest in AI innovation to attract Generation Z customers and at the same time better respond to the needs of senior clients

by Gimede Gigante @

The benefits of Artificial Intelligence (AI) affect all sectors of the financial industry, from insurance to payments, from asset management to savings services (Verma, 2023). For example, thanks to AI it is possible to develop ad hoc recommendations to help consumers better manage their capital and achieve their savings goals (Aboelmaged et al., 2013). Even opening a mortgage can become a much more efficient procedure thanks to AI, eliminating the need to go to a physical branch, as well as improving risk assessment for lenders. Finally, automation represents the new frontier in the field of investments, thanks to the possibility of obtaining more information about customers and offering them ad hoc services, by considering their preferences in terms of asset allocation and the level of risk aversion (Goncharenko, 2019). In addition, the use of robo-advisories offers banks the possibility of speeding up contacts with customers, for example by sending alerts in response to changes in the markets with almost instantaneous speed (Jaksic et al., 2019). Furthermore, the ability to develop accurate analyses in terms of risk profile and market forecasts, together

with more efficient portfolio management, allows for improved rates of return. Finally, AI makes it possible to reduce manual intervention in banking processes, so as to decrease the costs of advisory services and reach a wider audience, allowing financial companies to lower commissions on individual investments.
All these elements offer a crucial competitive advantage, namely the possibility of attracting younger consumers, the so-called Generation Z, who attach great importance to technological innovation (Kaur et al., 2020). The development of personalized banking services, which manags to achieve a deep understanding of consumers' needs and preferences, would offer a win-win scenario for both banks and consumers themselves, who would end up being offered perfectly customized products and services (Gigante, 2022). To conquer this target market, it is essential for the bank to invest in new technologies to improve its digital services and pay attention to aspects such as human contact and reputation. The former represents a fundamental component to attract older consumers, for whom it is a priority to be able to easily contact their financial advisors. In this regard, some banks have introduced a "voice analyzer", which allows the automatic transition from robo-advisor to human assistant. As far as the reputational aspect is concerned, another advantage of AI is transparency, as robo-advisors employ algorithms based on financial theory (for example, they build portfolios based on Markowitz's optimization theory). In this regard, the study by D'Acunto et al. (2019) analyzed the use of the Markowitz model by a robo-advisor in the selection of investment portfolios. What has emerged is that the adoption of robo-advisors has differential effects on investors based on their levels of experience: less specialized investors have greatly improved their performance in terms of portfolio returns, thanks to greater diversification. On the other hand, institutional (and therefore already diversified) investors underperformed. This study therefore suggests that every bank must know how to use automated systems in light of the level of the financial knowledge of individual customers.
In conclusion, the growth prospects for this sector are really interesting: McKinsey has estimated that the use of AI in the banking sector can globally generate up to $1 trillion in value added every year. Similarly, Bloomberg has projected annual growth rates of 39.9% for the robo-advisory market until 2028. The use of automated systems therefore represents a necessary investment for banking institutions, and an absolute priority in a country, such as Italy, where the push for innovation has so far been sacrificed. In fact, compared to the European average, investments in deep tech by venture capital funds are very low in Italy, and the country in general ranks behind countries such as Ireland, Denmark and Spain in terms of research contributions. It is therefore a priority to build an efficient cooperation system between universities, private industry and the state, in such a way as to direct more capital towards innovation and implement the transformation of our ecosystem to facilitate the creation of spin offs and communities of high tech entrepreneurs, as well as the funding of start-ups in their initial stage of development (seed investors). ■

GIMEDE GIGANTE
Professor of Corporate Finance and Real Estate, Director of ICE-Innovation and Corporate Entrepreneurship, SDA Bocconi School of management

The key that opens

Mathematical optimization can enable a better understanding of the deep structure of neural networks, so as to arrive at an AI that is not only generative but also explainable, i.e. capable of justifying its own choices. Only in this way, AI can prove useful in assisting political, ethical, financial or commercial decisions

by Laura Sanità @

The latest breakthroughs in Artificial Intelligence (AI) build on deep neural networks, a specific type of AI systems. Today, their applications are known to the general public in several areas, in particular: the so-called "large language models" capable of producing human-resembling text and conversations (ChatGPT, Bing AI, Bard), the "text-to-image" generative models which can produce striking images from a text captions (DALL-E 2, Imagen, Stable Diffusion, Midjourney), as well as various others connected to the recognition, production and translation of speech and sound.
What is behind the success of these techniques? The domains in which those successes happened are not haphazard or random, but rather perfectly suited for the strengths of neural networks as they exist today. Indeed, neural networks are proficient at emulating complex actions, even when they are difficult to define in precise terms, as long as we can feed them enormous amounts of examples to learn from. For example, it is challenging to define a "beautiful" or even "well-formed" image, sound or text, in precise mathematical terms. We do have, however, copious amounts of each to learn from.
A big challenge that AI is currently facing in its path to broaden its range of applications, however, is that deep neural networks are inscrutable. Indeed, the underpinnings of current AI systems are essentially a gigantic table of numbers. Those are not just any random numbers: the AI systems find the numbers that allow them to best reproduce the examples they have learnt from. However, the table's size and composition become so complex that they are devoid of any real semantic structure. As a result, we understand how neural networks function only in a superficial sense: we could crunch the numbers and reproduce any output from a given output. But we know pretty much nothing beyond that. For this reason, they are called black-box systems, notoriously unable to provide a justification for their output. This can be fine for some image, text or sound generation: one ugly or malformed piece does not jeopardize the usefulness of the tool if many others are good. However, it might be decidedly inappropriate for political, ethical, financial or business decisions: here one wants to be able to explain why a certain decision is considered to be the right one. Furthermore, emulating past decision processes without having/understanding the rationale behind it might make us overlook possible biases or latent errors.
In order to widen the applicability of AI we should thus adopt algorithms that provide a justification for their answers ("explainable AI"). For neural networks, this involve complementing the massive engineering efforts of the last years (which brought on the recent stunning results) with a deeper understanding of their structure from a mathematical and theoretical perspective. Mathematical optimization approaches are a privileged route towards bridging this gap in understanding.
On its own, optimization is the key instrument for solving many operations problems. Roughly speaking, it deals with selecting the best solution for a given problem, out of a set of possible ones. The crucial point is formally proving (i.e., explaining) optimality of a given solution without enumerating all possibilities, but rather by exploiting the mathematical structure of the problem under consideration. Unsurprisingly, optimization is a pillar for the mathematical foundation of modern AI. Indeed, such systems deal with finding the algorithms and parameters that best model a given task, and this, in itself, can be seen as an optimization problem. Furthermore, optimization problems are often formalized as maximizing an objective function subject to given constraints, where both the objective and the constraints are specified in terms that can be readily understood by humans. This is why strengthening the interplay of AI systems and optimization techniques can help in breaking their inscrutable nature. It is there that the amalgamation of optimization and AI bears its most promising fruits. ■

LAURA SANITÀ
Associate Professor of Computer Science, Bocconi University

LEGGI L'ARTICOLO IN ITALIANO

the black box

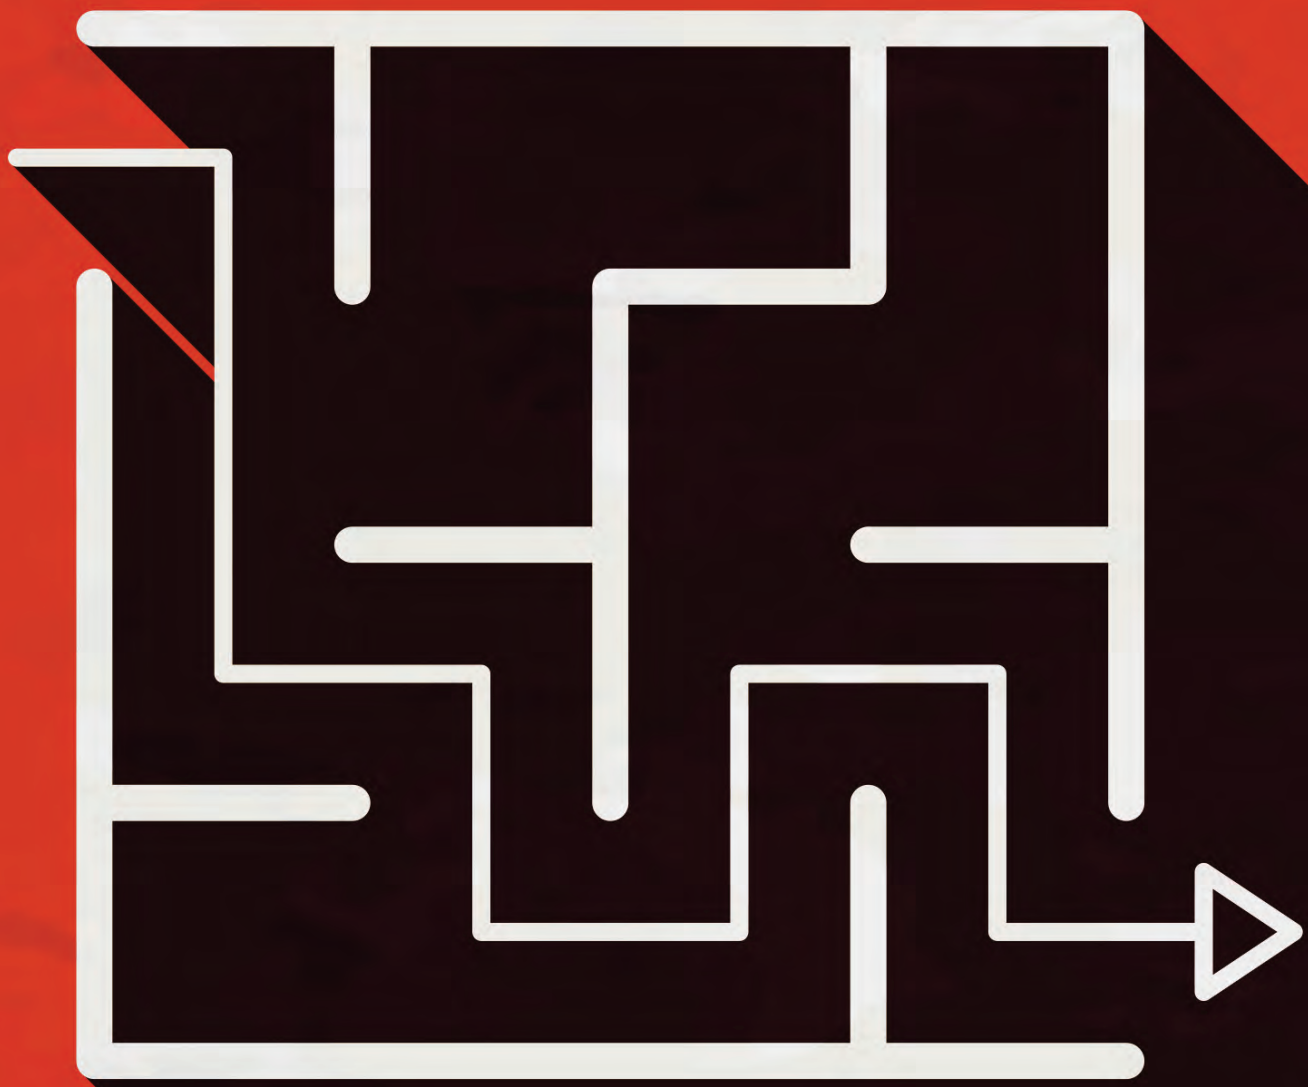

For self-driving cars, the re

At the moment the goal is not to hamper innovation with laws that block producers but for the near future the questions that need to be answered are: how to evaluate the "negligence" of a system? How to regulate no-win situations and the processing of data that autonomous learning systems require to improve driving performance? This is work for the next generation of jurists

al challenge is a legal one

by Francesco Paolo Patti @

Road traffic has always been a subject of great interest for legal practitioners. The history of cars in the 20th century was marked by constant technological progress followed by changes in social and industrial organization. With the advent of mass production and the internal combustion engine, the automobile became a mass phenomenon that had a strong impact on society's customs. Technology has also had a significant impact on the world of law. The vast increase in the number of passenger cars has caused a notable spike in risks posed to society. The law had to counter this phenomenon with detailed rules for road traffic, with innovative rules on civil liability, and the introduction of compulsory vehicle insurance. In recent times, on the basis of interventions by the European legislator, the law has adopted similar solutions in various countries, which place the legal burden on the owner of the car, who is required to insure himself against civil liability. The question that everyone is now asking is: how will the legal framework change with self-driving cars?

The technology for autonomous driving depends on automated learning systems capable of processing data and determining the movement of cars. In the final stage of technological evolution - which will soon be attained - the human driver will no longer have to pay any attention at the driving wheel and will be able to take care of other activities. Furthermore, it is expected that cars will all be connected to each other and there will no longer even be the need for traffic lights. Some sociological studies state that people will lose interest in owning cars and that the automotive world will become like a utility service available to everyone. Users of all ages will be able to "call" vehicles using their smartphones. The advantages for society are evident in terms of lower environmental impact and greater inclusion, but - above all - a drastic reduction in the number of traffic fatalities. In fact, it is common ground that more than 90% of deaths caused by road accidents are due to human error. By eliminating the human factor, the number of victims will be significantly reduced and, to use the expression of Guido Calabresi, a famous scholar of tort law at Yale, the costs for society will be reduced.

FRANCESCO PAOLO PATTI
Associate Professor of Private Law, Bocconi University

The objective of the law is to accompany this kind of technological progress without placing obstacles to innovation. At the same time, the law must safeguard the health and interests of people who come into contact with the technology. In this regard, we cannot forget the uproar caused by the first person being run over and killed by an autonomous car, a tragic event that took place in Arizona in 2018, and a piece of news that made the headlines of all major newspapers in the world.

The first studies carried out at the European level show that the reactions of the legal world have been very cautious. It seems that for a while, the system that places liability on the owner of the car with the support of compulsory insurance will still be able to offer satisfactory answers. However, on the horizon we see the need for greater involvement of auto manufacturers that oversee the development of the operating systems of autonomous driving vehicles. If today the need is not to "discourage" investments in new technology, thus shielding the manufacturer from liability towards the injured, things could change in the future. In fact, there is discussion of a special liability of the car manufacturer, which will be held accountable for damages caused by operating systems and be required to take out insurance to cover compensation payments. But the road is still long and there are many aspects that still need to be clarified. How will "negligence" of an operating system that rests on algorithmic decisions be evaluated? How will so-called "no-win situations", where the operating system has to decide whether to sacrifice human lives to save others, be regulated? How will the processing of data needed by self-learning systems to improve driving performance be regulated? Law and technology scholars will focus on these and other issues in the coming years, with only one certainty: autonomous driving will change social customs and upend existing rules. ■

Learning by playing

From the security of ports and airports to the management of road congestions and energy networks to the allocation of bids for online ads on advertising platforms, these are all non-stationary situations in which multiple agents interact and make decisions. As in a poker game

by Andrea Celli @

AI agents often face situations where they have to make decisions in highly non-stationary environments, which can arise from factors like the presence of other agents making decisions at the same time or strategic manipulations of data by an adversary. As a result, developing machine learning algorithms that can effectively learn and make decisions in such dynamic and reactive environments is crucial. In scenarios where an AI agent interacts with other agents in the same environment, the learning algorithm should consider the interactions between each agent's decisions, objectives, and the impact of decisions on the environment. In these multi-agent learning scenarios, the simple algorithms based on gradient descent often perform poorly and cannot guarantee good solutions. For instance, let's consider a scenario where two AI agents are playing a game and aiming to learn a winning strategy. If both agents use standard optimization algorithms like gradient descent, they may become trapped in cyclic patterns, resulting in the failure to converge to good strategies. This highlights the importance of developing customized algorithms for these multi-agent learning tasks.

A crucial step towards designing successful multi-agent learning algorithms is incorporating principles from game theory into the design of the algorithm. In particular, such algorithms should be able to take into account the incentives of other entities involved in the interaction implicitly. A recent example of the successful application of techniques from computational game theory is the creation of superhuman AI for two-player and multiplayer poker, developed by Noam Brown and Tuomas Sandholm at Carnegie Mellon University. These games are especially complex, as the players lack complete information about the environment (e.g., they do not know the cards in their opponents' hands). The AI agent developed for this task consists of three key components: an algorithm that learns an approximation of an equilibrium strategy by repeatedly playing against itself (i.e., learning through "self-play") and without any input from human or prior AI play, a subgame improvement algorithm that enhances the coarse equilibrium strategy for the specific subgames reached during play, and a self-improver algorithm that addresses potential weaknesses identified by opponents in the approximate equilibrium strategy.

ANDREA CELLI
Assistant Professor of Computer Science, Bocconi University

Equilibrium computation algorithms typically combine machine learning and game theory in order to compute equilibrium strategies for challenging problems that involve numerous strategic agents. Equilibrium computation algorithms have a wide range of applications beyond playing games like poker, and are used extensively to tackle other real-world issues. For example, equilibrium computation techniques are used in scenarios where defensive resources must be allocated to protect vulnerable environments such as airports or ports. Other applications range from financial markets to the management of road congestions and energy grids. Moreover, such learning algorithms are crucial for building machine learning systems that are robust to adversarial attacks and cannot be exploited by malicious agents. Finally, the equilibrium computation framework can provide valuable tools for studying and better understanding other more general machine learning problems. For instance, modern bidding algorithms for online ad allocation on Internet advertising platforms must satisfy a variety of constraints such as budget and Return On Investment (ROI). These constrained optimization problems can be modeled as two-player games where two AI agents compete against each other, and they can be solved through multi-agent learning algorithms. These are only a few of the applications in which equilibrium computation and multi-agent learning are playing a significant role. In the future, as we move towards a world that is increasingly interconnected, and where several tasks and decisions are entrusted to AI algorithms, such techniques are expected to play an even more significant role. ■


A
A

Immigration, the democr of predictive algorithms

To avoid a populist drift in algorithmic decision-making, citizens must be able to maintain their political freedoms and possess an adequate technological education. Only in this way will public decision-makers feel the responsibility to adopt solutions that share and analyze data to feed the democratic process

by Graziella Romeo @

According to the United Nations Population Division, to guarantee the "support ratio", i.e. the number of people in employment per retiree, Germany would need about 3.4 million immigrants a year until 2050. The study used an algorithm to which another question was also posed: how many immigrants does Germany need to keep the number of employed persons constant, regardless of their ratio with respect to pensioners? In this case, the answer was 460,000 people.

Reading these data without giving in to the temptation of superficial synthesis requires a preliminary clarification on the elements of the calculation. For example, it must be specified whether the AI is processing the overall migration balance, i.e. the net value coming from the difference between the number of immigrants and emigrants, or instead the total number of entries and, therefore, only immigration flows. In other words, the data of the predictive algorithm can be interpreted scientifically only if the terms of the calculation are understood. Otherwise, the algorithmic operation can be offered to the public loaded with risks linked both in terms of its functioning and its outcome.

The development of a predictive algorithm that works on big data regarding migratory flows is often motivated by a series of concerns related to immigration policy. The prediction of flows has been represented as a solution capable of inducing at least three virtuous effects: allowing public authorities to prepare adequate measures to manage inflows and reception; prevent social tensions usually associated with an increase in immigration in a certain area of the territory; understand the direction of the demographic growth of a given territory and, therefore, understand the consequences of a change in the composition of the population on welfare systems and the labor market. However, the availability of information of this type could also stimulate the adoption of policies strongly oriented towards the reduction of arrivals and the strategic use of border controls. The effect, in this case, would be to encourage preventive immigration controls, with dramatic consequences for protection of fundamental rights, first of all the right to asylum.

GRAZIELLA ROMEO
Associate Professor of Comparative Constitutional Law, Bocconi University

The problem of a "populist risk", i.e. of a knee-jerk political reaction activated by algorithmic predictions, questions constitutional law at several levels. For this reason, it is necessary to explore the opportunities for a democratic use of the predictive tool. The risk of populist use of the predictive algorithm can be averted, in fact, only if citizens maintain their political freedoms and understand, with a reasonable intellectual effort, the topics covered by civic and political debate and the reasons behing political action.

Understanding civic issues that are processed through algorithms requires an adequate level of knowledge about AI. From this point of view, an entirely algorithmic decision-making process is compatible with democracy insofar as it is assisted by guarantees of minimum levels of technological education, as it emerges, moreover, from EU secondary law. Algorithmic decision-making therefore needs to be grounded in a widespread and shared understanding of what AI-based decision-making is. Only in this way can the use of algorithms increase the sense of responsibility of public decision-makers and push them to adopt solutions that share and analyze data with a view to fueling democratic processes. In this framework, research on immigration law can play a crucial role in contributing to a process of orienting algorithmic technology towards the rights and principles of the European constitutional tradition. ■

atic use

Einstein is the example to follow

We need to build equations that, starting from theory, are capable of describing phenomena of interest in the real world: just like the Nobel Prize for Physics did. And we need to do it by choosing the least complex, most transparent and therefore understandable path. Only in this way can AI truly be integrated into our lives without being viewed with suspicion

by Emanuele Borgonovo @

The scientific world is debating intensively on how to make Artificial Intelligence (AI) more transparent. Companies are introducing AI and statistical machine-learning tools at a greater pace in all aspects of their business. While these innovations tend to benefit us to a great extent, scientists, philosophers, and writers raise serious concerns about the threats that an acritical deployment of AI tools poses to mankind. Investigators highlight the tension between automation and augmentation. With automation, we delegate tasks to machines and possibly allow humans to devote time and resources to more creative tasks. With augmentation, we promote cooperation between humans and machines, with machines that make the human intellect more powerful. However, the impact of artificial intelligence can be disruptive: the uncontrolled release of ChatGPT has caused unprecedented alarms in schools and educational institutions at all levels, worldwide. Then: should we fear this AI revolution or enthusiastically embrace it? Certainly, if AI remains a set of obscure calculations, then fear and severe criticism opposing an uncontrolled AI diffusion should be the natural reaction. Already in a 1970 paper about the interaction between managers and numerical models, John C. Little of MIT mentioned that managers refused numerical indications when these were derived from models that remained obscure to the manager. Nowadays, our mentality is much more open to the use of data and of statistical and mathematical models that extract information. Nonetheless, analysts still run the risk of seeing their efforts rejected by stakeholders. There are several famous examples of failures of algorithms, due to macroscopically erroneous forecasts to unfair or discriminatory recommendations. One action urged by the scientific world is to make AI tools as transparent as possible. How can we this goal? First, by avoiding to use of complex models when interpretable models are yielding the same or a similar accuracy for the problem at hand. Here, by interpretable, we mean that an external user has a clear grasp of how a model performs its calculations. The prejudicial view that high accuracy can be obtained in all domains only with complex numerical architectures is increasingly challenged by researchers and professionals. Spinoffs and start-ups devoted to interpretable AI find increasing success in the market. However, much is still to be done to remove the black-box menace and solve the "absence of theory" problem. In a perfect scientific construct, we have a theory, (axioms from first principles from which theorems can be derived) that yields equations (models), and these models exactly describe a phenomenon of interest in the real world. A wonderful example is Einstein's exact calculation of the sun-light bending, computed even before the actual measurements. This setup is rarely (if never) available in AI applications. In this context, we consider data (or measurements) and create a mathematical model that links a set of features (or inputs) to the target (output) of interest.

However, several potential mathematical functions can

EMANUELE BORGONOVO
Full Professor and Director of the Department of Decision Sciences, Bocconi University

play the role of the model for a given dataset. Choosing the best model is only the first step. The model must be scrutinized through proper uncertainty quantification, stability, and sensitivity analysis. Borrowing again from John C. Little, a process needs to start to find out what it was about the inputs that made the output come out as they did. This process needs a combination of tools that allow us to understand what features are important in the response of a machine learning model, whether the model behaves by an underlying physical or business intuition (or theory if possible), and, ideally, to perform an X-ray of the model. This process should become an integral part of the release process before algorithms are presented to the public; this is essential to avoid the societal damages of an acritical use of these technologies. ■

Quantifying uncertainty

It is one of the open challenges that can be addressed through the deeper integration of probabilistic thinking and generative modelling with AI

by Giacomo Zanella @

A classical success story in machine learning (ML) and artificial intelligence (AI) in the last decade is the extraordinary accuracy in prediction and classification task - such as object recognition in images - obtained with large ML models. At the same time, many recent advances and open problems in AI require a deeper integration of ML with probabilistic thinking, including the need to build probabilistic representations of data. One obvious example is quantifying the uncertainty surrounding a prediction or classification. Think, for example, at purely providing a doctor with an AI-based diagnosis of a disease versus providing an assessment of our confidence in such diagnosis.
A popular example of probabilistic modelling in AI is offered by so-called generative models, which witnessed major advances in recent years. Instances include

GIACOMO ZANELLA
Assistant Professor of Statistics, Bocconi University

models for the generation of images (including controversial so-called deepfakes), text (from automatic creating captions for images to advanced chatbots), or even art (such as music generation software). The fundamental idea there is to build probabilistic models that are able to generate data that "look like" real data, e.g., photos of people that have never existed but could have existed. They are different from "discriminative" models that learn, for example, to predict which object is shown in a picture (despite poor light, uncommon posture etc) but cannot generate new images where such an object is shown. More generally, probabilistic modelling is not only about generating new "fake" data. It is about learning generative mechanisms, i.e. building models that quantify and potentially reproduce the randomness inherent in data. Such probabilistic

representations can help to perform various ML-related tasks such as identifying "unlikely" observations that might need additional information before being able to take a reliable decision; quantifying the uncertainty about a prediction or decision produced by a ML model; detecting outliers and suspicious behaviour; and exploiting the inferred generative model to make inferences about latent structures in the data.

Broadly speaking, probabilistic and generative thinking is widely used across sciences. Although different in their aims and interpretations, fundamental concepts such as latent variables, random effects, factor models, or mixture models that are pervasive across social sciences are actually examples of generative models. A classic example is given by topic models, which allow automatic extraction of meaningful topics from large corpus of text documents or, in other words, allow understanding and characterizing what documents are talking about. These and many other natural language processing techniques have enabled researchers and companies to treat "text as data" to be fed as input for downstream tasks, thus having major impact in many application areas including research in Political Science and Economics. Another common example is probabilistic recommendation, where "products" and "customers" are assumed to possess latent unobserved features that determine the likelihood of a given customer giving a certain rating to a given product. Statistical learning is then used to infer relevant features from observed data and thus build a concise and yet informative representation of products and customers types.

Many generative models, including the examples above, build a probabilistic representation for data x by

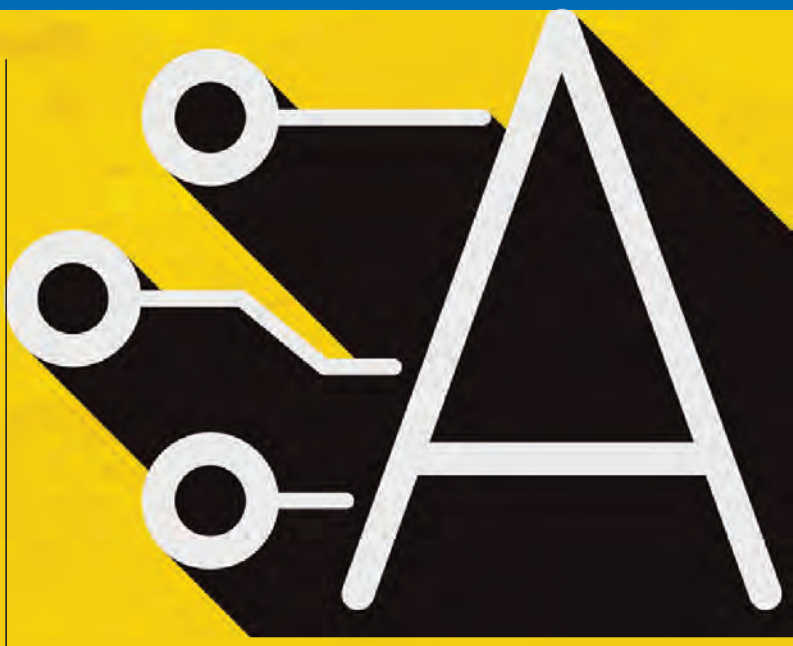

specifying a joint probability model for x and z, p(x,z), where z are latent variables aimed at modelling fundamental but unobserved sources of variation. In the above examples z would be topics and x words chosen depending on the topic; or z would be customer and product features and x the observed ratings. Learning from such models, either in order to produce new data or to make inference about underlying structures, requires to compute the marginal distribution of data, p(x), or the so-called posterior distribution of latent variables given the observed data, p(z | x). These tasks involve major computational challenges, especially in modern applications with thousands or millions of latent variables in the model. These challenges are usually tackled with one of two main classes of algorithms: variational ones, which build a deterministic and "easier-to-handle" approximation of p(z | x); and Monte Carlo ones, which build a stochastic representation of p(z | x) by appropriately drawn random samples. Providing deeper understanding of the computational and statistical workings of such algorithms in the context of large-scale probabilistic models, as well as developing better and more efficient algorithms, is the focus of my recent ERC Starting Grant for the project "Provable Scalability for high-dimensional Bayesian Learning".

Looking forward, a deeper integration between probabilistic thinking and AI can contribute to tackle key challenges in these fields, ranging from uncertainty quantification to interpretability. A fascinating aspect of current research in probabilistic and generative modelling is that similar frameworks and even algorithms are nowadays increasingly used across very diverse scientific fields. This gives a key importance and responsibility to methodological research in Statistics and ML, which can help recognizing common structures and facilitate the flow of ideas across fields.■


Bocconi
Federica
Beneficiara di esonero
La vostra generosità mi ha permesso di vivere un'esperienza che mi aiuterà a diventare davvero chi voglio essere.
La tua firma può scrivere un futuro
Il tuo
5x1000
sostiene giovani brillanti e motivati grazie ad una borsa di studio.
Find out more at
5x1000 | Bocconi - Giving
Codice fiscale: 80024610158

Why coup

Personal satisfaction, amount of paid work of the woman, some personality factors and age: these are the four main risk factors. And it was an algorithm that identified them

by Letizia Mencarini @

les break up

What does Artificial Intelligence offer which goes beyond traditional statistical models, such as regression analysis, to investigate the behavior of households, in particular the factors that cause the separation of couples and dissolution of the conjugal bond?

With Bruno Arpino (University of Florence) and Marco Le Moglie (Catholic University of Milan) we have analyzed data for over 2,000 German married or cohabiting couples, who were followed for a dozen years on average by the annual GSOEP survey (German Socio-Economic Panel), with over 900 ending in separation. By adopting a machine learning approach (specifically, Random Survival Forests) the procedure found on its own the relationship between the various factors contained in the database. In this case it considered more than 40 factors, from age to education level, from health to psychology traits: the mass of raw data was fed to ML, without making precise hypotheses, but simply indicating as event of interest the break-up of the union, and the algorithm indicated the influence of each factor contained in the data. The variables that pose the greatest threat to the stability of a union have been identified with an accuracy of 70% (a predictive ability that outperforms the 50% achieved by traditional regression methods). Not only was ML able to discover the factors behind the breakup of couples, but it was also able to use this knowledge to predict the end of a union before it happens. This is also because, instead of submitting all the data available to an ad hoc algorithm, half were used to instruct the algorithm itself and the validity of the results was verified with the other half of the dataset.

The results of the analysis are very interesting, above all because the ML methodology is able to weigh the relative importance of various factors in causing the breakup. Factors that had been particularly influential in previous studies have instead lost their relevance here, like unemployment, and the partner's high level of education and income.

The four major risk factors, that emerged from the study are in descending order: personal satisfaction, the woman's quantity of paid work, some personality factors and age.

LETIZIA MENCARINI
Full Professor of Demography, Bocconi University

The strongest predictor of separation is personal satisfaction: if both partners are dissatisfied obviously the couple won't last. Less obvious is that a strong drop in conjugal stability emerges when the woman is very satisfied with the union but the man much less so, while the reverse effect is less evident. If the woman works many hours outside the home, the risk of separation or divorce is higher, even when the man is more involved in domestic chores (but this result is nothing new and according to the existing literature it depends on the greater agency and independence of working women. As for personality traits, a high "extraversion" in men (classically linked to a higher infidelity) and a low "openness" in women, less adaptable to the changes brought about by cohabitation, are the traits that more strongly associated with the end of a couple. Also a low level of "conscientiousness" in both partners (understood as organizational capacity in daily life, and therefore - if low - as disorder and inability to respect commitments) does not help to stay together. But also a too high or too low level of "neuroticism" can be a problem. This result can be interpreted as the fact that suffering from excessive anxiety, jealousy, guilt, worry or anger clearly complicates the relationship. This is true above all for women, but, on the other hand, those who don't feel this type of emotion could lead their partners to read that personality trait as lack of interest (men, in this case). However, no pairing of personalities was determined that is more strongly associated with the breakup of the relationship. Finally, considering age, very young couples tend to be more unstable, but for women stability in relationships intensifies after the age of 40, while this is not the case for men.

ML analysis is not without limitations. In this case a major one is that it refers only to Germany and also has few details on the psychological aspects of the two partners. However, from a methodological point of view, the study demonstrates the great potential of ML techniques in demographic and sociological research in general, highlighting their ability to monitor and analyze a large number of predictive factors, to automatically find linear or non-linear relations, additive or non-additive relations between these factors and the outcome of interest, with greater precision and more robustness of estimates against collinearity than commonly used methods. ■

THE PAPER

What Tears Couples Apart: A Machine Learning Analysis of Union Dissolution in Germany,
by Bruno Arpino, Marco Le Moglie, Letizia Mencarini

Cleaning up the data to

The most modern machine learning methods use computational shortcuts and simplified models, where Bayesian methods come in handy. Stacked Penalized Logistic Regression is the result of

by Botond Szabo @

Machine and statistical learning are at the core of artificial intelligence, where the goal is to extract knowledge from the data and learn from it. Modern applications require complex models and the available real world data are never perfectly clean or accurate, often containing measurement and other errors making the problem even more difficult. Statistics is the science to analyse and interpret such noisy, imperfect data and it plays a leading role in all modern data-centric developments.
In particular, in recent years, the amount of available information has increased substantially and the models describing real world phenomena are becoming also increasingly more complex. These introduce new challenges for data scientists, since despite the ever-increasing power of computers, the computational complexity became overly large, making it impractical or even impossible to carry them out in reasonable amount of time (or memory requirement). Therefore, novel, modern statistical and machine learning methods were developed to speed up the computations using simplified models and computational shortcuts. However, these methods are often used as black box procedures without rigorous mathematical understanding. This could result in misleading and wrong answers without us even realizing it.
A particular example are neural networks, which are the state-of-the-art approaches for image classification with applications ranging from medical imaging to self-driving cars. However, it was shown that minor changes of the input images (which couldn't even be detected by human eyes) or unusual positions of the objects could result in completely inaccurate classifications leading to wrong diagnosis or incorrect detection of objects.
Therefore, it is of high importance to study their theoretical properties and derive guarantees but also limitations for these modern learning methods. One particularly important aspect is to understand how much we can rely on the derived results. In more formal terminology it is essential to correctly assess the uncertainty of the procedure, which is based on noisy, real world data, so can never be perfect. A principled way of obtaining uncertainty quantification is by using Bayesian methods. Bayesian statistics provides a natural way of incorporating expert knowledge into the model in a probabilistic way and automatically quantifies the remaining uncertainty of the procedure. Bayesian statistics is becoming increasingly popular in machine learning and artificial intelligence, for instance in Natural Language Processing for constructing chatbots often Bayesian approaches (naïve Bayes classifiers) are used to find the most likely answer.

BOTOND SZABO
Associate Professor of Statistics, Bocconi University

My ERC Starting Grant focuses in particular on the theoretical understanding of statistical and machine learning methods, including the accuracy of parallel computing methods and the information loss incurred by considering simplified models instead of the accurate, complex models. Then, based on the theoretical understanding I aim to propose new approaches which have higher accuracy. The main focus on my work is on mathematical statistics and its intersection with machine learning, information theory and numerical analysis. I am also involved occasionally in more applied projects, which build on the theoretical insights from my core research.
Working closely with scientists at the Psychology Institute of Leiden University we have developed a learning method aiming for detection of Alzheimer disease. In medical research, often different type of data is collected and combined to provide the best diagnosis. For instance, for early diagnosis of Alzheimer disease structural and functional MRI data, questionnaire data, EEG data, genetic data, metabolomics data,… etc can be collected. These data are substantially different both in overall size and quality. To achieve the most accurate early diagnosis one should find the most important features in these data sets and combine them in an optimal way. Furthermore, since these diagnostic tools can be expensive and of limited capacity, it is of important to select the most relevant ones to achieve a reliable, accurate and cost effective diagnostic method. We have developed a learning approach called Stacked Penalized Logistic Regression (StaPLR), which selects the most relevant diagnostic tools and the corresponding most relevant features for predicting to early onset of dementia. This method was successfully applied on clinical data containing patients with Alzheimer disease and a control group.

THE PAPER

Variational Bayes for High-Dimensional Linear Regression With Sparse Priors, by Kolyan Ray and Botond Szabo

reduce uncertainty

, which however require deep mathematical knowledge to eliminate background noise. This is an applied project for the early diagnosis of Alzheimer's disease

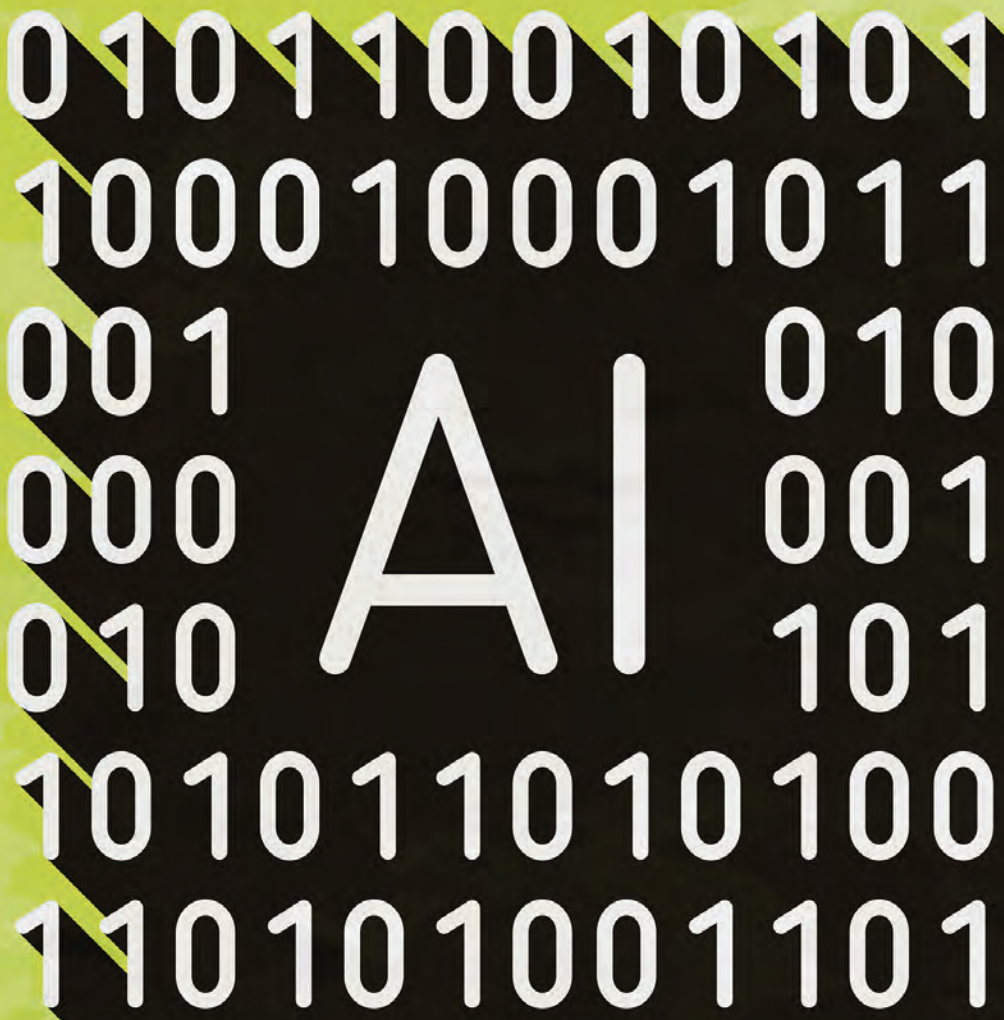

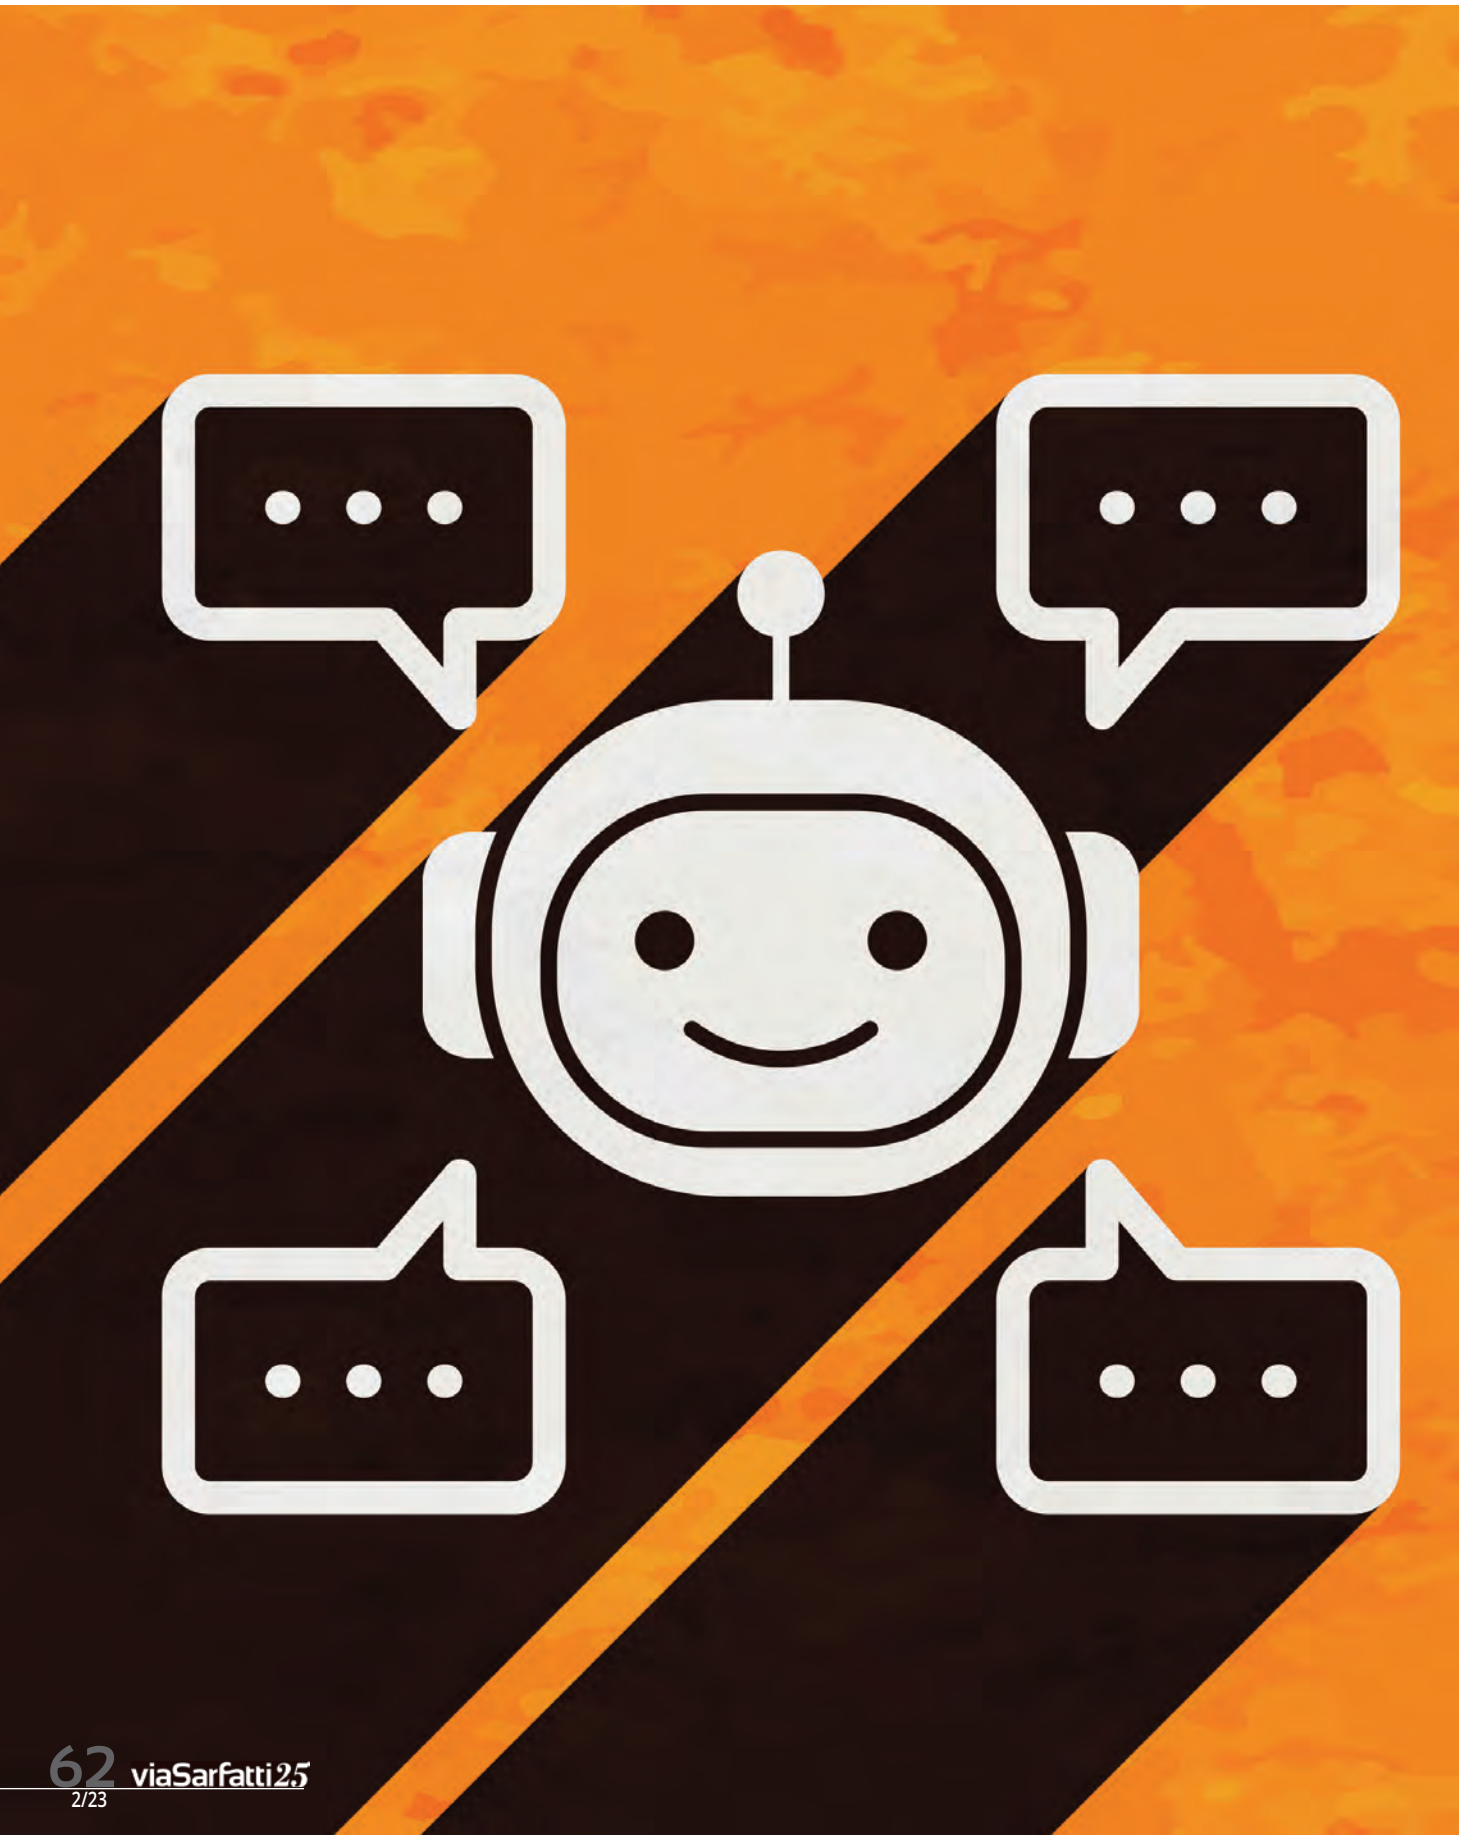

Signed by ChatGPT

NLP-based tools, through their ability to empower people, are transforming industries such as customer service and journalism. But like all new technologies, AI too must be treated with caution, in order to amplify its benefits and mitigate its risks

by Debora Nozza @

The world we live in today is one where technology is rapidly advancing and shaping the way we interact with each other and our environment. One such technology that has emerged in recent times is Natural Language Processing (NLP). NLP is a branch of Artificial Intelligence (AI) that enables computers to understand, interpret, and generate human language. The potential benefits of NLP are vast and far-reaching, particularly when it comes to empowering people.

Think about it - with NLP, we can bridge communication gaps between people who speak different languages or have different communication styles. We can make information more accessible to everyone, regardless of their background or ability. We can create new and exciting learning tools that will revolutionize education as we know it. And that's just the tip of the iceberg!

NLP is not only creating new job opportunities but also revolutionizing the way we work. The emergence of NLP-powered tools is transforming industries such as customer service and journalism. With the help of NLP, customer service representatives can automate their basic queries while focusing on more complex customer issues. Similarly, journalists can harness the power of NLP to conduct initial research, sift through large quantities of data, and generate summaries or article drafts. Giving them more time to focus on analyzing and interpreting the information gathered. As NLP continues to evolve, its potential to transform and enhance various industries is limitless, leading to a significant positive impact on current jobs.

A recent study conducted by researchers at the Massachusetts Institute of Technology (MIT) investigated the productivity effects of a generative NLP technology, ChatGPT, on mid-level professional writing tasks. The results showed a significant increase in productivity and output quality, as well as reducing inequality between workers. ChatGPT was primarily used for idea-generation and editing tasks, restructuring the writing process towards these aspects and reducing the need for rough-drafting. Additionally, exposure to ChatGPT led to heightened concerns and excitement about automation technologies, indicating a growing awareness and openness towards utilizing AI technology to enhance productivity in the workplace.

DEBORA NOZZA
Assistant Professor in Computing Sciences, Bocconi University

NLP has the potential to empower people in ways that were once unimaginable. By providing computers with the ability to understand and generate human language, NLP can help in so many aspects of our life and who know how many more in the future. However, realizing this potential will require careful consideration of the risks and challenges associated with the technology. If we can navigate these challenges, then NLP has the potential to transform society in profound and positive ways.

In closing, it's worth mentioning that ChatGPT has been an essential tool in creating this editorial piece. Its ability to generate natural-sounding and coherent text has undoubtedly empowered me to assist in conveying these ideas to you, the reader. It's a prime example of how NLP can assist and empower people to communicate more effectively and efficiently, and I hope it has made for an engaging and insightful read. So, the next time you come across an article or text that seems almost too good to be true, don't be surprised if ChatGPT played a part in its creation! ■

Make or buy? That is the question

Three reasons and three questions to understand how to integrate AI into your organization without having to reinvent the wheel every time

by Lorenzo Diaferia, Gianluca Salviotti @

The great enthusiasm of recent months regarding generative Artificial Intelligence (for example ChatGPT) and large linguistic models (such as GPT3 and now GPT4) has made even more evident a paradox that has been underway for several years. On the one hand, the increase in interest from companies and the general public for this new technology. On the other, a certain slowness in integrating complete, scalable and functional AI systems into companies, so to produce concrete impacts on business organizaitons. There are many decisions to be taken in the implementation phase. Let's try to tackle one. Remaining on the example of ChatGPT and GPT3, it is easy to see how OpenAI, the company that develops and markets ChatGPT, not only provides the ready-to-use conversational systems but also access via a series of paid cloud services (API) to the family of linguistic models underlying the solution (GPT).

In introducing these tools into the life of companies, one should therefore ask oneself: is it better to use ChatGPT as it was designed and marketed by OpenAI or is it better to start from the model that enables it (GPT3), to create customization specific to the context and types of use we want to work on? A similar question applies to virtually every area of AI, from predictive maintenance systems to quality

THE PAPER

When Standard Is Not Enough: a Conceptualization of AI Systems' Customization and its Antecedents, by Lorenzo Diaferia, Ivo Blohm, Leonardo Maria De Rossi, Gianluca Salviotti

monitoring systems, passing through forecasting and recommendation systems. On close inspection, however, this make-or-buy choice is not so different from what companies already had to do while developing previous technologies. In the current context of AI, however, some peculiarities make the problem somewhat different compared to other more consolidated technologies. At least three factors influence this dynamic.
The first is related to the very characteristics of AI. At present, in fact, the use of systems based on "machine learning" and the related "training" implies that the elements of the AI system to be customized are substantially different from other technologies. In addition to the enabling infrastructures (software and hardware), in fact, customization can involve data, extend to models or even touch the internal functioning mechanisms of algorithms.
The second factor is linked precisely to the peculiar structure of an AI system. The need for training data with certain characteristics, the computational power (and related costs) necessary to train large models, as well as the technical skills to work beyond the configuration of already existing and tested solutions, represents a barrier to entry for a large percentage of companies in terms of achieving customization.
The third factor refers to the peculiarities of the market which has structured itself in response to these characteristics. Solutions range from finished products and standardized APIs that offer little or no configuration capability, to solutions and development environments that support technically skilled professionals in developing solutions that extend more established approaches.
Although there is no single approach to determining the meeting point between these three factors, and therefore the type of product and level of customization to focus on, it is possible to trace three dimensions that can support the choice. First, design features. The availability of adequate and distinctive data to carry out the training, the resources (in terms of money and time) available, the costs related to the training, and the level of differentiation and strategic nature of the business problem addressed with AI are all considerations when thinking about the search for the balance between customization and standardization of solutions.
Second, the maturity of the organization. The availability of AI skills in the organization or the presence of adequate resources to supervise and support the work of external operators is an essential point in order to access the creation of systems with a high degree of customization.
Finally, the maturity of the market, both in terms of performance and characteristics of the AI models available to address the strategic problem at hand, and in terms of overall maturity of these systems, which include both the AI component and the connectors and software and hardware elements required for operation. To avoid reinventing the wheel, internal organizational needs must be adequately balanced wit solutions available on the market, articulated on the entire spectrum of available AI options. ■

LORENZO DIAFERIA
Junior Lecturer of Digital Transformation at SDA Bocconi School of Management

GIANLUCA SALVIOTTI
Associate Professor of Practice of Information Systems at SDA Bocconi School of Management

THE LAB

Digitization is an undeniable and irreversible trend in the business landscape. Companies find themselves deciding whether to "undergo" the effects of digitization, adapting to some "mythical" best practices, or to lead the process directly by leveraging the most advanced digital technologies and thus addressing their own business transformation. Thanks to its multidisciplinary approach and leveraging on business-oriented research initiatives and in-depth studies focused on digital, the DEVO Lab (Digital Enterprise Value and Organization) offers all the skills and capabilities needed to fully support managers in a Digital Transformation process.

THE COURSE/1

The program Web 3 e Metaverso (in Italian) provides an overview of the main current trends in digital environments, showing how companies and organisations can capture the value of the new Web.

THE COURSE/2

The Artificial Intelligence for Business program (in Italian) provides the tools to translate the possibilities of AI into action in companies, balancing strategy with technological and organisational limitations.

United in the name of

New technologies based on machine learning allow advantages in time-consuming activities such as the analysis of great quantities of logs. Since they employ vast amounts of data, however, it is absolutely necessary to ensure their security and privacy

by Greta Nasi @

GRETA NASI
Associate Professor of Public Administration, Director of the Master of Science in Cyber Risk Strategy and Governance jointly offered by Bocconi University and Politecnico of Milano

Cybersecurity requires governing several processes, including scanning, log analysis, incident response, and threat detection. These processes can be time-consuming and labor-intensive, demanding security analysts spend significant time reviewing and analyzing data. Automation can help streamline these processes by reducing the need for manual intervention, allowing security teams to respond more quickly and efficiently to security incidents. As Artificial Intelligence (AI) technologies become more prevalent and powerful and drive digital transformation through automated decision-making capabilities, their application is rapidly growing. While the benefits of this emerging technology are significant, there are also concerns to be aware of. AI techniques and systems may produce unexpected outcomes and be susceptible to tampering that could manipulate expected outcomes. Consequently, cybersecurity and artificial intelligence are increasingly intertwined. On the one hand, AI can enhance cybersecurity by identifying and mitigating threats more effectively. On the other hand, AI itself can be vulnerable to cyber-attacks which could compromise data and systems (Li, 2018).

The advantages of AI in cybersecurity are numerous (Floridi and Taddeo, 2018). AI can aid in automating tasks, identifying abnormalities in real-time, and responding to cyber threats with greater speed and efficiency. One critical process in cybersecurity is threat detection and response, which involves identifying and responding to potential cyber threats and attacks. The goal is to detect suspicious activity or anomalous behavior on a computer network, system, or application and respond quickly to mitigate the risk, minimise possible damage and prevent further harm. It involves repetitive and continuous monitoring of systems and networks to identify unusual behavior that threatens business continuity. Advanced threat detection technologies include machine learning and AI algorithms. These can help automate threat detection and behavioral analysis activities and be trained to identify patterns and anomalies in network traffic, responding to cyber threats accurately and quickly.

AI technologies are also currently employed in vulnerability management to identify technical risks in hardware and software systems allowing organizations to patch and update them before attackers can exploit them. Organisations employing complex computer networks as critical infrastructure operator that governs a power plant can significantly benefit from these technologies, as manual vulnerability scanning would not be feasible.

Overall, AI can potentially transform cybersecurity operations in the future. Thus, its implementation has challenges. They include the need for more skilled professionals, the importance of solid organizational data governance, and the potential for AI to be breached or manipulated. AI technologies offer many benefits in improving cybersecurity and other tasks and processes that typically require human intelligence in terms of automation, improved decision-making, personalized experience based on the users' preferences and increased safety and security in several areas such as food policy, transportation, and healthcare. However, they also pose some security challenges.

Firstly, AI technologies rely heavily on data to learn and make decisions. The security and privacy of data are critical (Zhang et al, 2019). Suppose sensitive data is compromised or accessed by unauthorized parties. In that case, it can be used to launch data breaches that may serve several purposes as cyber-attacks, industrial espionage, the launch of misinformation campaigns and other security incidents. In addition, it may be hard to understand how some AI systems make decisions. This needs more transparency to ensure identifying potential privacy and security risks.

Secondly, AI models can be biased, leading to unintended or even discriminatory outcomes. Bias and fairness challenges may lead AI algorithms to make incorrect decisions, which become malicious in

cyber-defense

the case of adversarial attacks that target the AI models to manipulate the data they use to learn and make decisions.
To address these challenges, it is essential to implement rigorous privacy and security measures to manage data properly and ensure the fairness and transparency of AI systems while providing human oversight and intervention. All organizations using AI technologies should adopt appropriate security policies and procedures specific to AI, invest in security training and awareness for AI professionals, and regularly audit and test their AI systems to identify vulnerabilities and address potential security risks.
As AI and cybersecurity rapidly develop, their integration is expected to offer increasing application scenarios. Future developments may include the study of AI-based situational awareness for cybersecurity, which can provide intelligent prediction and protection for cyberspace; the development of novel and specialized AI algorithms for cybersecurity, particularly for big data intelligence; and innovative security protection solutions for AI in the future. ■

Tracking terrorists on Fac

Thanks to machine learning it is possible to track down the phenomenon of terrorist recruitment. A hitherto impossible goal that Natural Language Processing has made real. A Bocconi research study shows that terrorist financing increases strongly and significantly in the recruitment phase

by Nicola Limodio @

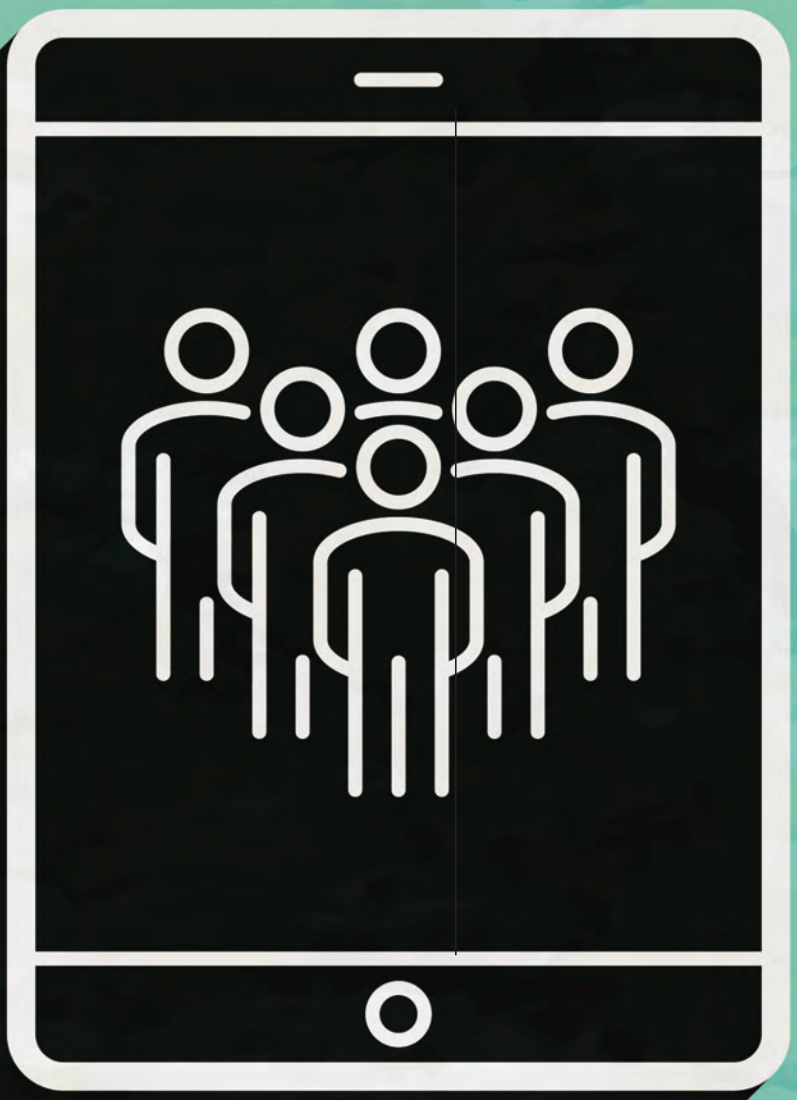

ebook and the Dark Web

Among several valuable applications, machine learning is particularly valuable for measuring previously objects that were previously unmeasurable, like the tone of a text, the colors of a picture and many more. This results in new data on phenomena that were impossible to study due to their lack of quantification. One instance that accurately describes this phenomenon is terrorism recruitment, which I study in my paper *Terrorism Financing, Recruitment and Attacks* published in Econometrica in 2022.

Measuring the recruitment of terrorists, or members of criminal organisations in general, is inherently complex because it is distinctively unobservable. At the same time, one of the many channels through which terrorist groups recruit is via online fora (like Facebook, Reddit, et cetera). For this reason, machine learning algorithms can be useful to detect recruitment and can be built to automatically evaluate each message's content, determining whether it contains elements of terrorism recruitment. To measure recruitment, I gathered data from various online fora operating in Pakistan that disseminate content supportive of jihadism. The Artificial Intelligence Lab at the University of Arizona developed a dataset containing more than four million messages, exchanged between 2003 and 2012 on six different fora operating in Pakistan. Additionally, the database includes more than 2.5 million messages transmitted on platforms in the dark web, an alternative internet network requiring a specific software for its access and navigation. Extremists and terrorist groups have routinely used these platforms to spread the concept of war against the unfaithful (Jihad). The possibility of having access to such a large number of text messages has great importance, as they can be helpful to quantify -and thus obtain a measure of - the phenomenon of terrorist recruitment.

NICOLA LIMODIO
Assistant Professor of Finance, Bocconi University

Nevertheless, determining which among those messages have the purpose of recruiting terrorists is a difficult task.

Without a specific algorithm, the determination of whether the content of a message concerns terrorist recruitment would be prohibitively expensive: judges and investigators would need to read and analyse each of the four million messages. To lower the cost of this endeavour, it is crucial to implement a machine learning algorithm that aids in categorising a message either as neutral or as having recruitment intent. This goal can be achieved through a technique called Natural Language Processing: a data science system able to understand the contents of scripts, including the contextual nuances of the language within them. The algorithm works through a method defined as supervised learning, which implies that it needs to be trained on a set of already classified data before being able to accurately extract the information contained in each message and categorise it. This is done relying on the initial work of two judges, who independently and manually reviewed a random sample of messages, highlighting and marking all those showing an intent to recruit violent extremists. This sample constitutes the training data and is used to teach the algorithm how to recognise conversations containing any recruitment material. Once the algorithm is trained, it can be applied to all the remaining messages, de facto replicating the work of several judges.

By implementing this method, I was able to characterise a measure of terrorist recruitment, which can be used to understand the determinants of terrorist attacks and can assist national security agencies. My research shows that the effect of terrorism financing on attacks increases strongly and significantly in recruitment. Additionally, this innovative way to classify written texts can have vast applications in future studies, as it can be exploited in any situation where experts are needed to assess third-party material. ■

THE PAPER

Terrorism Financing, Recruitment, and Attacks, by Nicola Limodio

The evolution of the b

A transformation in business and governance and stronger attention to fairness, transparency and privacy. Thus credit institutions are changing skin thanks to AI

by Anna Omarini @

AI-driven transformation continues to accelerate. In this phase, banks are faced with new challenges and opportunities, which force them to pay attention to a plurality of aspects, all transversally crossed by a major cultural transformation. The main aspects affected by this transformation are, on the one hand, corporate strategy, which concerns the degree of pervasiveness of AI within banking activities, the dedicated budget, the benefits encountered and the challenges that remain open, as well as of course their progress over time. On the other hand, however, there are the governance aspects, which embrace organizational issues, processes, management of skills and issues relating to risk management and ethical implications.

With reference to the strategic aspects, it should be noted that AI today represents an increasingly important investment item for Italian banks. 76% of respondents to the ABILab survey (2023) state that they have a dedicated budget and in 65% of cases this budget is more than €500,000 and 59% of those interviewed expect that budgets will further expand in 2023. There is a strong tendency to interpret AI as a strategic driver of business transformation. In fact, 88% of respondents stated that they have defined/are defining an ad hoc strategy, which will then have to be put in synergy with the company's Data Strategy. This theme also implies the strengthening of partnerships and ecosystems useful for promoting innovation. In fact, Italian banks are already using mixed sourcing models for the development of AI solutions, using make and buy levers accordingly and interfacing with a variety of players (ICT partners, startups, etc.)

In defining an AI-driven transformation path, it is also important to frame the benefits that banks intend to seek and obtain. In this regard, it is interesting to note that answers to the survey do not exclusively refer to higher economic returns from a reduction in costs (53% of the sample), but also to improvements in decision-making processes (53%) as well as strengthening the relationship with the end client (41%) and, last but not least, improving existing products/services. In fact, the majority of use cases already present on the market concern support to (assisted and self-service) channels and support to control functions.

ANNA OMARINI
Researcher in Finance, Bocconi University

On the other hand, as regards governance issues, most of the responding banks (58% of them) say they are working on implementing a framework for the governance of AI and these initiatives are often placed in continuity with the efforts that have been devoted in recent years to the definition of a company model for Data Governance, now widely operational in banks.

The main pillars on which the banks intend to build this governance system are those of transparency, explainability and traceability. They also believe it is useful to have a monitoring system that makes it possible to evaluate the results of AI from the point of view of business effectiveness, technical efficiency, governance and ethics. This kind of monitoring that is already in place in about 40% of the sample.

A further important factor concerns skills, which are necessarily heterogeneous (as they range from risk management to cybersecurity and privacy) as well as the need to hire new profiles in the AI field with reference to Data Scientists, Data Engineers and Machine Learning Engineers. In the area of governance, another relevant issue concerns the verification of ethical principles, which refer to the concepts of fairness (ensuring that AI is fair and impartial), transparency (verifying how data is used and how systems make decisions) in addition to privacy aspects (that is, aimed at ensuring that customer data are not used beyond the intended and declared use). Last but not least, banks also intend to invest in accountability, highlighting the rules, policies and models that make it possible to determine responsibility for the decisions taken by the AI system. Naturally, in all of this, a significant role will also be played by regulatory safeguards regarding the compliance of AI systems.

As a result, new challenges arise for banks. Bearing in mind that AI will not only concern technological skills, but will introduce changes throughout the organization; on the governance front, organizational processes and systems will need to be put in place to favor its conscious and responsible adoption. Furthermore, the speed of adoption of AI will be a key factor, as well as the understanding of its impacts on the ethical, social and sustainability spheres. ■

anking species

Technology in times

The impact on rights and freedoms, institutional consequences and the implications for economic rights: these are the three questions the law must deal with in the case of AI systems in emergency situations such as terrorist threats or pandemics

by Chiara Graziani @

How is advanced technology being used to stem the adverse effects of emergencies, both the political ones (as international terrorism) and the non-political ones (as the pandemic)? Is the frequent replacement of human decision-makers by "intelligent" algorithms affecting how advanced democracies deal with emergencies and ensure rights protection in times of stress?
These questions, which are at the basis of the Bocconi unit's stream of research within the PRIN excellence project "Decision-Making in the Age of Emergencies: New Paradigms in Recognition and Protection of Rights", arise from three main concerns raised by the use of artificial intelligence (AI) in the course of crises. The first is the impact on rights and freedoms per se (i.e., the risk of their disproportionate limitations); second, its institutional consequences, which indirectly have a spillover on human rights; and third, the relationship between resort to AI and the market, hence the implications on economic rights.
The first aspect is the most obvious: we are all accustomed to think that automated and self-programming systems - not just during emergencies, but in everyday life - may infringe on our rights to privacy, data protection, free speech, given the amount of data (so-called big data) that algorithms need to be fed. Nevertheless, when this happens in times of emergency (let us think of surveillance regimes based on so-called black boxes), even more sensitive issues arise: on the one side, people are prone to accept stronger limitations for the sake of a higher good (e.g., national security or public health), but this could be detrimental for the keeping of the rule of law; on the other hand, in emergencies more than in any other context, these advanced tools are often used in secret, contributing to serious troubles in terms of transparency and accountability.
The institutional consequences concern the progressive and apparently unescapable erosion of public powers that the use of refined technologies entails. Although there are some attempts to regulate these tools - also during emergency - through legally binding sources, their practical implementation is widely left to so-called Big Tech, which, in this way, is significantly involved in public functions (e.g., programming algorithms for surveillance purposes, developing apps for contact tracing, and so on).
Undoubtedly, the shift of power given to these entities is likely to impact on the guarantees for rights and freedoms, as these private bodies work according to a different logic with respect to public bodies.
The third aspect is that the more these advanced algorithmic systems are essential to effectively deal with threats (let us think of algorithms that flag potentially "dangerous" content online), the higher the risk is that only big companies will be able to afford them. Thus, smaller enterprises could end up being kicked away from the market, with potential adverse consequences on economic rights, at least if no counter-measures are taken.
The described dynamics are not necessarily evil, as they might be natural features of the development of the digital age, inevitably affecting, among others, the management of emergencies. However, for sure they need to be monitored and watched over, in order not to "scar" the main features of rule-of-law-based systems. In other words, technology runs faster than law, and this is a matter of fact; yet, this should not lead public regulators to totally pull back. ■

CHIARA GRAZIANI
Research Fellow, Department of Law, Bocconi University

# of crisis

From machines that learn, a new industrial revolution will be born

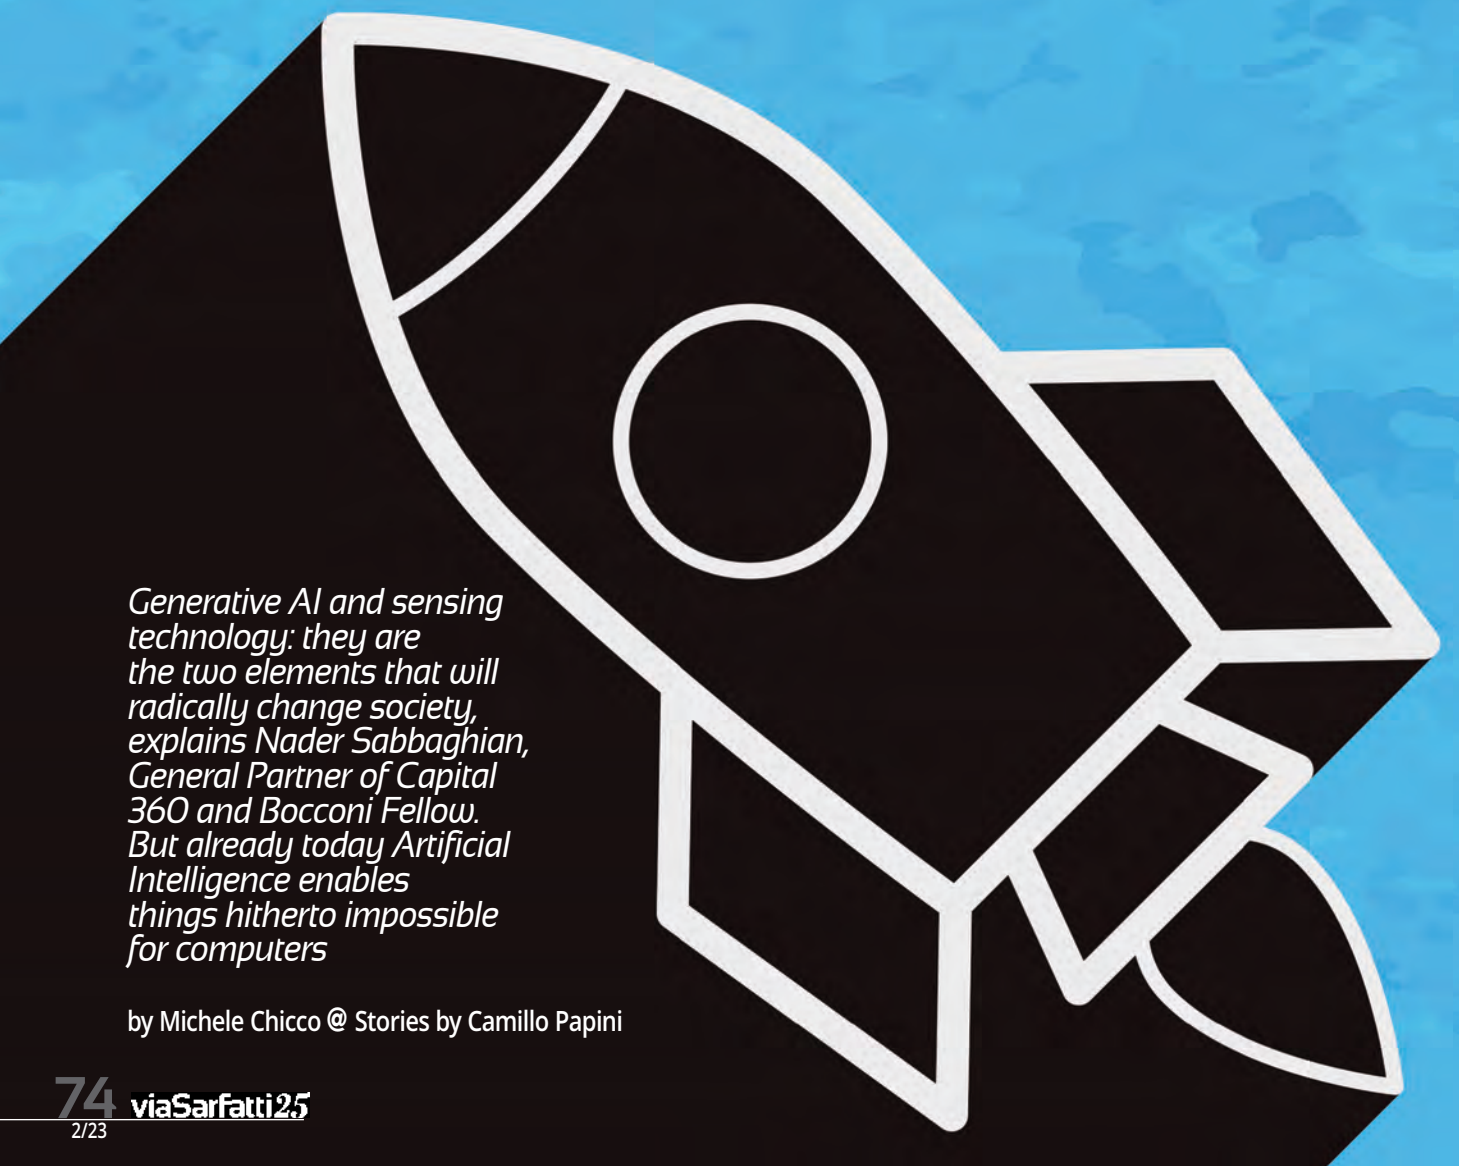

Generative AI and sensing technology: they are the two elements that will radically change society, explains Nader Sabbaghian, General Partner of Capital 360 and Bocconi Fellow. But already today Artificial Intelligence enables things hitherto impossible for computers

by Michele Chicco @ Stories by Camillo Papini

The computer science we have known until today will no longer exist. The genie is out of the lamp: "Artificial Intelligence (AI) is a new technological paradigm that will take us to the next industrial revolution, a bit like what happened in the early 2000s with the generalized diffusion of the Internet". And **Nader Sabbaghian** was there: in 2001 he was part of the founding team of Bravo Solution, today he is General Partner of Capital 360 and Academic Fellow at Bocconi University, where he teaches Technology and Innovation Management, and is faculty member of the MBA of SDA School of Management. "Classic computer science - he says - is deterministic, since it allows you to automate processes only by way of algorithms. AI is something altogether: machines acquire cognitive abilities, learn from mistakes and improve over time". VC funds are heavily betting on this: according to Dealroom data, $167 billion went to AI startups between 2021 and 2022. In Italy, it's $401 million in 59 rounds of investments.

➜What are the prospects for the market?

Huge. There is the opportunity to create new business models and new startups, but with artificial intelligence we can also improve business processes. AI works only if we have high calculating speed and vast troves of data, essential conditions for solving enormously complex equations at affordable costs. AI is now entering everyday life, and at the speed the market is going within the next decade we will see mass products and services generated by Artificial Intelligence.

➜ On what should innovators focus to attract capital?

They must identify concrete solutions for very specific problems at hand. Much of the entrepreneurial challenge is to identify the specific segments that can rely on big data to train the models. We are keener on Narrow AI, i.e. startups that are well focused on their target: two years ago we invested in Pallon, a Swiss company that automates the process of inspection of sewers. A complex activity, now made with probes that track the ducts and are then viewed by a human being. A slow activity that can be done much more precisely with AI, once it learns who to recognize leaks in the sewage network. Pallon exploited the first customers to collect data, but now after thousands of kilometers inspected, AI has learned enough to start replacing humans and do this work independently.

NADER SABBAGHIAN Accademic Fellow at Bocconi University and General Partner of Capital 360

➜ What is the lesson for founders?

Better to be the best in the world to do a super-specific thing than trying to achieve a great ambition. We know how difficult it is: the founders that come to us with an idea often do not yet have even the product prototype and if they aim too wide, they will struggle to achieve what they want.

➜ What should aspiring entrepreneurs focus on?

There are two really interesting themes that will allow you to radically change society: the development of advanced sensors and generative AI, the creative artificial intelligence that originates from ChatGPT.

➜ Sensors, therefore hardware. What is the connection with AI models?

Problems need to be solved in the physical world, not only on our computers: we will use artificial intelligence to move from a world of automation to a world of autonomy. Today we have very sophisticated robots that have no cognitive skills; to have AI rise in their stead we need a major development of sensing technology. You have to invest a lot in sensors suitable for creating machinery that can navigate this new world like we humans do.

➜ Generative AI enables a computer to reason like us.

I didn't think that I will be able to see something like that in my own lifetime. We are building the conditions for machines not only learn what we know to do, but go further that hat. Until now, AI didn't do it, it was merely a problem solver. Generative AI emulates the act of human innovation, since it combines previous knowledge to create something new. We are still at the beta stage, yet software like ChatGPT is already giving the goosebumps. What matters is the basic technique: these programs will soon be used for scientific research because they will be able to take in all knowledge, contextualize information and transform it into something new.

➜ AI will revolutionize entire economic sectors: should we be more frightened or fascinated?

The next twenty years will be very interesting from the point of view of economic development, there will be another industrial revolution. And like all past ones it will cause major upheavals. But you cannot put the brakes on technology: now it has emerged and we must understand how to adapt our way of developing human resources to the rules of the new world. Human beings know how to adapt well, and we will do so now too, and learn how to do business in the new conditions. Governments will have to make sure that this technology develops in a non-destructive way, and will be called upon to protect those who, during the transition, are likely to suffer the most from the disruption caused by AI growth. ■

THE ACCELERATOR

B4i - Bocconi for innovation, selects the most innovative ideas and the startups with the greatest potential to help them grow and succeed by investing the experience, knowledge, time and resources of the wider Bocconi community.

Tailored investment portfolios

Vittorio Carlei, CEO of Qi4M, a startup incubated by B4i, explains how AI can offer support to small investors, for example by translating financial language into natural language

The market is not yet mature, since savers are not yet educated and ready to approach fintech services directly. Thus companies that offer financial products and services through technological platforms based on Artificial Intelligence (AI must team up with large groups if they want to stay in business and find the necessary capital and a sufficient number of customers. "The market consequently tends towards concentration", states Vittorio Carlei, co-founder, CEO and Head of R&D at Qi4M, a company that has been on the market for five years and specializes in the creation of asset management robot solutions for institutional investors. "It should be noted, however, that AI will be able to involve the end saver with more accessible solutions". Among AI applications that are hitting the headlines is ChatGpt which, not surprisingly, Qi4M is researching to create a consultancy service that builds personalized investment portfolios. "This technology brings with it the advantage of being easy to use, since it translates financial jargon into natural language, and then provides solutions compliant with the European MiFID Directive", continues Carlei, who is launching a crowdfunding campaign to support the birth of the new advisory service. Final goal: to enter the retail market with services accessible to all between the second half of this year and the beginning of 2024.

At the same time, the startup, developed by Bocconi University's accelerator, B4i - Bocconi for innovation, continues to work together with an incumbent player to present the first robot investment fund on the market. "It's a pity that in Italy the legislation does not provide for this type of service", explains Carlei. "So we will launch it abroad and, if we will possibly bring it to the Italian market only at a later time, with a dedicated sandbox". But, more generally, how do you overcome distrust of latest technologies and promote them among the general public? "By explaining that we need to look at the results which can reassure us and bring us to trust them, even if we don't understand their functioning in detail", concludes the co-founder and CEO of Qi4M.

Futura, the platform that helps you with GMAT

Andrea Chirolli, Bocconi alumnus and founder of the company explains how their algorithms prepare students at a fraction of the cost of a tutor

A bespoke learning system, which the more it is used the more it captures a student's peculiar strengths and weaknesses as he/she progresses through his/her education. Futura is able to extrapolate data on different approaches thanks to artificial intelligence algorithms (supervised and deep reinforcement learning) and consequently provide personalized programs. In fact, as preparation progresses and various tests are passed, the platform further models the experience on the individual's profile. "Our company was born in 2020 and deals with preparation for passing standardized tests. Currently, the platform is aimed at young people who intend to enter medical school in Italy and pass the GMAT", says Andrea Chirolli, CEO and founder of Futura. "It is the solution for those who want to enhance their learning without having to invest significant money in private tutors". Looking ahead, Futura aims to expand in the Italian market by covering the entire range of university entrance tests.

The business model raised €1.8 million from institutional investors such as LVenture Group, United Ventures and Exor Seeds. From a commercial policy point of view, "the platform wants to replace private tutors with a personalized experience, at a fraction of the cost," underlines Chirolli, who graduated from Bocconi with the World Bachelor in Business (WBB). "By buying a monthly subscription, users can follow courses and evaluate their initial level of skills through a test". Finally, students have the opportunity to interact in real time with actual tutors and other course participants, thus creating an online community. In the coming months, Futura will implement its functions, adding, for example, the possibility of chatting with tutors, further forms of community and more simulation methods.

Did you say cotoletta?

Qodeup, a startup incubated by B4i, uses algorithms to streamline restaurant orders, explains CEO Fabio Marniga. It also translates menus for patrons

It is quick to present the menu to customers, instantly translates it into the language of the customer and provides them with a quick service at the time of payment. It progressively get acquainted and recognize all the patrons. It is called Qodeup and it is the web app of the same-name startup born in May 2019 and accelerated by B4i – Bocconi for Innovation, since February 2020. The restaurateur in the flesh remains, as well as the waiters, but the tasks of presenting the list of dishes, with knowledge of every Italian regional specialty, and ensuring digital payment can be entrusted to Qodeup (pronounced kyüdup). "It is a service via QR code, quick for the consumer and useful for streamlining the management of restaurant tables for the operator", underlines Fabio Marniga, co-founder and CEO of Qodeup. "In the near future, our platform will be enriched by new services such as the profiling of users who, whenever they return to the same establishment, will find their favorite dishes among the suggestions. We expect to start developing the new services at the end of this year", continues Marniga. "The goal is an ever-more personalized restaurant experience". In fact, the system evolves by analyzing the data relating to menu browsing and ordering, for example which dishes are consulted the most and in what order. Today there are 5 million unique monthly users. It is free for customers, says the CEO of Qodeup, while the owner pays a percentage on each digital payment through the app. On the other hand, the menu consultation system is available according to a quarterly or yearly subscription formula. At the end of last year, the start-up clinched a €2.5 million fundraiser and now "we look at a sector that is becoming increasingly aware of the quality behind our services", says Marniga. "I saw lot of improvised competition in the aftermath of Covid".

Is the city safe? Mine Crime tells you

The startup incubated by B4i, headed by Giacomo Salvanelli, uses AI to create risk indicators on crime in neighborhoods, which then sells to the market

Creating a network of commercial activities and citizens who share a participatory and circular conception of urban security: this is the aim of Mine Crime, a startup born at the end of 2020 and incubated and accelerated by B4i - Bocconi for Innovation, which transforms Big Data in risk indicators thanks to Artificial Intelligence (AI), then proposes them to the market as a new service with the intention of reiterating the social advantages AI can bring. The long-term idea is to incentivize large companies and especially SMEs to feed data on the level of security in urban neighborhoods. By providing information to local news, businesses are rewarded through Mine Crime by receiving a preferential treatment from customers who register on the system. "Italy is at the bottom of the European rankings for the use of open data according to the Open Data Barometer. It is difficult for a citizen and for a company to access certified and clear information", explains Giacomo Salvanelli, General Manager of Mine Crime. "On the corporate front, in particular, we study further solutions for SMEs but also for large companies, for example those that manage shopping malls or are present in areas such as railway stations. Thanks to their data", concludes the co-founder of the start-up which works on 16,000 data sources, "the AI learns to know any pockets of illegal activity that may exist even in neighborhoods considered peaceful or, conversely, indicates there are safe transit areas and malls even in areas considered risky". Artificial Intelligence therefore evolves, increasing its ability to analyze reality and translating its results into various knowledge tools available to civil society for money or for free. It ranges from maps of neighborhoods with data on criminal activities recorded and safety ratings for individual areas or more in-depth, organic reports. After transforming big data into risk indicators thanks to AI, the newborn company also offers among the publication of reports, the planning of crime prevention events, as well as training courses.

Olivia tells you how many customers you will have today

Thanks to algorithms, the startup incubated by B4i optimizes estimates of customer flows in restaurants so as to also avoid food waste, as CEO Tansella explains

It was born during the pandemic when the catering sector was experiencing a major downturn due to lockdown restrictions and uncertainty about the future. Now Olivia uses the power of Artificial Intelligence (AI) not only with a view to consolidating the restaurant business thanks to more accurate estimates of the expected flows of customers but also aiming to reduce food waste which remains very high in the sector. The startup that grew thanks to B4i-Bocconi for innovation, accelerator of the Bocconi University, was born to answer questions such as "How many customers will enter my restaurant today?" and even "What will they order?", in order to "optimize orders to suppliers, decide the number of staff needed in the kitchen and at the table, and organize everything accordingly", says **Francesco Tansella**, CEO of Olivia, a company developed on Zoom from an idea of a group of students of the Sant'Anna School of Advanced Studies and Scuola Normale of Pisa. Potential savings are a multiple of 10 with respect to what a restaurateur invests by purchasing Olivia, i.e., on average, more than a quarter of the EBITDA of the business, again according to Tansella who is keen to underline how "in smaller restaurants but also in large chains there is a psychological obstacle to overcome because often the entrepreneur, even in a business city like Milan, thinks he or she can do it alone", without grasping the impact of the predictive capacity of AI. Olivia works like a web app that processes the historical data of an individual restaurant every day, and combines them with calculations on third-party data such as local events on the calendar or weather forecasts, and communicates the results to the entrepreneur or manager covering a period of time up to two months. A mobile version of Olivia is being studied, but what is important to understand, Tansella comments, is that "the benefits also translate into higher customer satisfaction because the risk of running out of ingredients, and therefore dishes on the menu, becomes minimal".

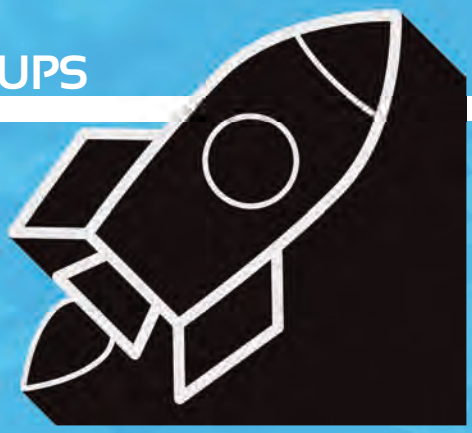

Software and sensors to manage company assets

Francesca Tosi, Head of Growth at Quick Algorithm, explains how the business of the startup incubated by B4i has evolved

Diversifying the services offered and also after-sales support policy, with the intention of strengthening the positioning of Quick Algorithm, a startup born in February 2018 and incubated by B4i - Bocconi for innovation, the University's acceleration program. Initially, Quick Algorithm only offered the Artificial Intelligence service to optimize the management of a company's physical assets, such as production machinery or even company buildings. All thanks to Scops.ai, a cloud-based proprietary software which analyzes big data from different sensors and sources, including those generated by human maintenance activities. The business advantage proposed: reducing industrial costs. For example, energy ones. Now, however, "we also supply the sensors for detection, which for the moment we do not produce ourselves", says **Francesca Tosi**, Head of Growth at Quick Algorithm and corporate board member. "We started from the needs of some companies that did not have the necessary sensors in the field or, if they did, the data was collected from different data sources or platforms that were not integrated with each other". Therefore, the business has expanded in parallel with the commercial plan that allows you to purchase, as rental or leasing, Quick Algorithm sensors. "With the increase in energy costs, the number of requests we receive has increased", adds Tosi. "With Scops.ai, companies can not only monitor and predict the numerous and varied anomalies of their plants, but also and above all discover unknown phenomena that would have been difficult to identify without the support of machine learning". An unexpectedly large and detailed environment, the boundaries of which range from incorrectly tightened clamps to the cold room door left open. Moreover, "AI based on machine learning learns from data and is able to take into account "normal" anomalies due to factors such as seasonality and the type of production. In these cases, the system does not send any alerts", concludes Tosi.

Looking at venture capitalists in the rearview mirror

Using algorithms to find patterns in historical data on startups and predict the likelihood of success for a new startup could hold entrepreneurs back, by encouraging the repetition of similarly innovative ventures at the expense of breakthrough innovation

by Clement Jonathan Mazet-Sonilhac @

The digital age has created massive amounts of data that continue to grow exponentially. The International Data Corporation estimates that the world generates more data every two days than the entire humanity did between the dawn of time to the year 2003. This rapid evolution in data availability has been accompanied by advances in statistical techniques such as machine learning and Artificial Intelligence (AI), i.e. technology designed to identify statistical patterns in large datasets, outperform humans in many forecasting (or repetitive) tasks. This big data revolution has been reshaping the financial industry: AI techniques are being increasingly deployed in finance, in areas such as asset management, algorithmic trading, corporate banking, credit underwriting or blockchain-based finance. The deployment of AI technologies in finance is expected to drive competitive advantages for financial firms - for example by reducing their costs and increasing their productivity. In turn, these competitive advantages can benefit financial consumers by providing increased quality and personalized products.

In a recent article I illustrate the latter. I examine the early adoption of new technologies by banks lending to small businesses, using the staggered deployment of high-speed internet in France from 1999 to 2007 as an exogenous shock to technology diffusion. I conclude that faster and easier data exchange between entrepreneurs seeking credit and small banks (lower search costs and improved customer–lender interaction) reduced the cost of borrowing faced by these entrepreneurs. In addition, the results show that banks are responding by changing their business practices, lending to small businesses located farther away, outside their local market, but without deteriorating the quality of their loan portfolio. This echoes recent work on fintech lending by Berg, Fuster and Puri (2022). Fintech lending is defined as the use of technology to provide lending products, with two main flavors: First, technology can be used to improve the customer-lender interaction (for example, with a fully online application process), giving rise to a better user experience, faster processing times, and lower operational costs; and second, it can be used in borrower screening or monitoring, for example, by using alternative data sources or machine learning methods. The authors find that the first channel seems to matter the most: the increase in convenience and speed have been more central to fintech lending's growth than improved screening or monitoring.

CLEMENT JONATHAN MAZET-SONILHAC
Assistant Professor of Finance, Bocconi University

However, the adoption of AI technologies by financial intermediaries has also raised concerns regarding their effects on investment decisions and, more broadly, on the allocation of capital. A specific class of financial intermediaries has recently received a lot of attention with respect to their use of AI technologies: Venture Capitalists (VCs). Venture Capitalists are private equity investors that provide capital for startups with high growth potential and play a crucial role in the financing of innovation: among public firms founded within the last fifty years, VC-funded companies account for more than 92% of R&D spending and patent value! In a recent article, Bonelli (2022) shows that dozens of VCs have adopted AI technologies for screening startups over the last decade. These VCs employ AI algorithms to detect patterns in historical data from previous startups and extrapolate them to predict a new startup's outcome. Despite using cutting-edge advances in machine learning, these forecasting algorithms are in essence backward-looking - because they are trained using past data. As such, they may not be successful at screening startups that differ radically from past companies. The paper confirms this intuition: VCs adopting AI become better at identifying good quality startups, defined as startups that survive and receive follow-on funding, but only within the pool of startups whose business is similar to that developed by past companies. At the same time, VCs that adopt AI become less likely to invest in startups that achieve major success, like an Initial Public Offering (IPO), or an innovative breakthrough, for example a highly cited patent. This finding is associated with an increase in the share of their investments being oriented towards startups developing businesses closer to those already tried-and-tested. Overall, this study highlights a potentially adverse aspect of the adoption of Artificial Intelligence on firm financing, with less funds directed to very innovative - disruptive - companies. This raises two important questions of whether this might induce entrepreneurs to produce more backward-similar ventures at the expense of breakthrough innovations, and what could be the implications for long-term economic growth. ■

Augmented creativity

The challenge is to make philosophical zombies work together with humans and thus overcome the limits of both to generate a new kind of digital marketing that manages to overcome the current obsession for optimized personalized communication via AI to reconnect individual users with the collective values of a brand

by Andreina Mandelli @

Generative AI creates content, ideas and artifacts, including music, visual art, software, products, using Large Language Models, i.e. probabilistic models that start from the linguistic patterns learned from the data inputs on which the AI was trained, which in the case of GPT-3 reaches the staggering number of 175 billion parameters.
Good marketing is made of insights, strategies, ideas and the development of new offers and constructions of relationships and market processes that create value for customers and the company.
AI can create insights from listening to consumers in social media or from the interactions with customer services to generate ideas for new products. Adidas used generative AI to create the design for a new shoe based on customer feedback, historical sales data and market trends. AI can assist in creating 3D models of products, also optimizing the final output, evaluating costs, manufacturability, consumer preferences and possible pricing options. META Advantage on Facebook and Instagram advertising, on the other hand, allows you to create and optimize campaigns, generating and testing hundreds of creative versions, up to choosing the most performing one in terms of conversion, once the campaign objectives have been indicated. Coming to the automation of customer service, we have been dealing with chatbots for some time but often with a rigidity and a quality termed by the *New York Times* "the spiral of misery". ChatGPT is certainly less rigid in its responses and for this reason it has been recently adopted by Bank of America and other financial companies.
Today, however, various criticisms are raised (among others, by Ogilvy) against the drift of digital marketing toward personalized and optimized communication via AI verging on obsession, at the expense of a more authentic connection between brands and customers based on meanings and values that are also collective and on a kind of creativity that knows how to connect in a less mechanical way to the personality and identity of the brand. These criticisms join the more specific ones pertaining to generative AI, which highlight important critical issues in these systems in terms of the non-verifiability of information, the difficulty of giving recommendations in complex contexts, and flaws in the protection of personal data and copyright.
Let us remember that, according to Margaret Boden, creativity derives from a kind of virtual reality in the human head, made up of spaces and mental maps, (often subconscious) "theories" about the world which we use to navigate reality. They are the basis of our creativity because we can modify and reinvent them. They involve three types of thinking: exploratory, combinatorial, and transformative. On its own, generative AI has excellent exploratory and combinatorial capabilities but several limitations on the ability to change these maps autonomously or even just to combine them in a way for which it has not been trained. In cooperation with humans (and in marketing we mean not only marketers but all consumers and other stakeholders), the result can overcome the limits of both. On the quality of information (without addressing issues of truth which would require a much broader discussion leading to somewhat uncertain outcomes), we recall that in the history of communication and journalism effective methods have been developed for assessing the quality of professional gatekeeping processes (community relevance, checks on sources, etc.), which can be included in generative AI systems, even in collaboration with humans. The announced disaster of the explosion of fake news and information manipulation in social media is due to human agents, i.e. the development of algorithms for optimizing attention and engagement for advertising purposes and malpractices by organized groups of humans, not to the machine intelligence itself. The ability to make recommendations/decisions in a complex context depends on the challenge between two

ANDREINA MANDELLI
Fellow of Digital Marketing at SDA Bocconi School of management

THE COURSE

Advanced Digital Marketing (in Italian) provides an overview of the levers at the disposal of today's and tomorrow's professionals, including data analysis and machine learning skills.

different visions of how AI should work. On the one hand, the exponential growth of data crunching (brute data force) that guides the learning of neural systems, to generate answers whose sources and logic are unknown ex-post; on the other, the neuro-symbolic approach promoted by the former head of the Watson project and now followed by META, where the search for probabilistically correct combinations starting from data mining and massive data crunching (by an invisible agent) becomes complementary with more traditional, rules-based, methods, and the input of already available models, which are tested and transparent to users in their causal relationships.

In summary, the bet is that we find a way to make so-called "philosophical zombies" (replicas of human intelligence that do not have most of the human capacities for intelligence) and humans, with all their inefficiencies but also great abilities, work together. As Margaret Boden said already in 1995: the challenge is to understand how humans can use AI to increase their potential to be creative, not to erase it. ■

Buyers 2.0

From Robotic Process Automation to machine learning, there are various technologies that are spreading among Chief Procurement Officers. Changing their work and always requiring new skills and competences

by Giuseppe Stabilini @

Purchasing processes are experiencing a period of great change and evolution due to heavy investments in digitization and the dynamics of transformation towards greater sustainability. The Covid19 pandemic and the war in Ukraine have also forced companies to redesign supply chains and review relationships with suppliers.
In this historical moment of great discontinuity, Chief Procurement Officers (CPOs) have taken up the various challenges by researching new solutions to support the purchasing processes. In particular, many companies have focused on adopting technologies to increase the productivity of procurement managers and the end-to-end visibility of supply chains, exploiting the value of data and the potential of Artificial Intelligence. From a research conducted in 2021 by the SDA Bocconi Procurement Lab which interviewed procurement managers of 132 medium-large companies, 37% of the sample had active projects in the AI field and 16% had already adopted the tool on a large scale.
The application of AI to purchasing processes can in fact provide an important contribution to the ability to process historical and contextual data, improving the buyer's ability to possess a more complete vision of the business scenario and to make more solid and rational decisions. Recent research has identified various applications on processes of Vendor Management, eSourcing & Tender Management, Contract Management and Spending Analysis.
In recent years, numerous companies have devoted economic resources and people to AI, starting a constructive dialogue with external subjects such as consultancy firms and software vendors and, in some cases, even with specialized start-ups, as highlighted by the research project.
The objective was first of all to make the operating machine of the purchasing processes more efficient, introducing algorithms capable of automating and governing time-consuming activities for buyers and to support them in the various phases of the purchasing process. Technologies such as basic data analysis and visualization, Robotic Process Automation (RPA), Natural Language Processing (NLP)/Text Analytics and Optical Character Recognition (OCR) have proven to be particularly popular in this area.
Secondly, as a consequent evolution, there has been a focus on the application of AI in complex areas, i.e. in support of a better reading and analysis of the competitive context. Algorithms such as Advanced & Predictive/Prescriptive Analytics, Machine Learning and Deep Learning are being used.
In this case, the effectiveness of the algorithms in analyzing and processing the data depends both on the availability of quality data (produced by the company, shared by the supply chain or provided by third parties), and on learning systems based on automatic feedback or returns from the reference buyer. To date, the output of the AI system, i.e. the final decision, is rarely left completely to the autonomy of the machine (only 4% of the projects analyzed). This aspect not only means that the buyer remains ultimate decision maker, but also highlights the need for full man-algorithm collaboration.
Large companies, early adopters to date, state very positive results from the introduction of these solutions (51% of projects), in several cases even exceeding expectations (10%). At the same time, they reveal how the critical issues for projects are linked to organizational aspects such as corporate culture, people's skills and the ability to re-engineer processes.
We are still in an early-adopter phase, but both technologies and the experience accumulated by companies and service providers are contributing to a strong acceleration of the phenomenon. Looking to the future, two investment directions can be identified.
For companies not yet engaged in AI projects, the advice is to work on corporate strategies and culture in order to ingrain greater openness to the technologies themselves, focusing on more proactive leadership towards external collaborations.
For organizations that have ongoing and already operational projects, there is a need to develop training plans for people to strengthen the skills acquired and develop new qualified skills, also through greater integration and collaboration with external parties. Furthermore, attracting young resources, more inclined to use and evaluate these technologies, can be of enormous support, especially if associated with the enhancement of experience and the ability to identify the critical points of processes brought by more senior buyers.■

GIUSEPPE STABILINI
Associate Professor of Practice of Procurement & Supply Management at SDA Bocconi School of management

THE LAB

Procurement Lab is a research center for developing skills, streamlining operational models and mapping out future scenarios in purchase and supply relationship management. Every year, SDA Bocconi's Procurement LAB promotes structured applied research on frontier areas of procurement and supply chain management.

Accounting is no long

Data entry, bank reconciliations and the processing of invoices are all routine activities that are increasingly entrusted to AI technologies, something which makes data reliability crucially relevant. And so companies are now increasingly pushed to embrace forward-thinking analytics

by Francesco Grossetti @

AI has the potential to revolutionize accounting and finance in several ways, including predicting financial market trends and automating routine accounting tasks. One of the most significant advantages of AI is its ability to process vast amounts of data quickly and accurately, which was previously done manually and was time-consuming and prone to errors. AI algorithms can identify patterns in historical data and make predictions about future market trends with a high degree of accuracy. This helps investors make more informed decisions about which stocks to buy or sell and when to do so. Another area where AI is making a significant impact is automating routine accounting tasks such as data entry, bank reconciliations, and invoice processing. These tasks can be tedious and time-consuming for accounting professionals. By using AI-powered software, these tasks can be automated, freeing up time for accountants to focus on more complex and strategic work. Overall, the future direction of AI in accounting and finance is promising, but careful consideration and planning are required to overcome some important challenges.
The use of AI in accounting centers around data and automation, with machine learning and other AI-powered technologies driving new and improved practices for data analytics in accounting. As technology continues to support much of bookkeeping, finance, and accounting, the amount of data available becomes increasingly vast and precise. AI inspires companies to embrace "forward-thinking analytics," which involves using AI to forecast weeks, months, or even years in advance, rather than relying solely on past data to make present decisions. Several benefits seem to be right around the corner. Concepts like invisible accounting, a function that quietly frees up humans from doing repetitive tasks and enable a more active involvement and engagement with strategic engagement, or continuous auditing, are becoming a reality and on the brink of physical implementation.
The use of AI in accounting has numerous benefits, but there are also concerns about its potential drawbacks. One of the most significant is the risk of bias in AI algorithms. If the data used to train the AI is biased, the results may also be biased, which could lead to inaccurate or unfair conclusions. This is a major concern that must be addressed since even small errors could have significant financial consequences. Another concern is the impact on human jobs in the accounting industry. While AI can automate many routine tasks, it cannot replace the expertise, judgement, and more generally, the common sense of human accountants and financial analysts.
To sum up, AI is rapidly gaining popularity in accounting and finance research due to its potential to provide significant benefits, including data processing, market trend prediction, task automation, auditing, risk reduction, and fraud detection. However, there are valid concerns regarding bias and the possibility of job displacement. As with any innovative technology, it is essential to understand AI's potential risks and advantages. Rejecting AI altogether is not the solution, as this could result in humans being replaced not by AI, but by other humans diligently who use this technology ethically and responsibly. ■

FRANCESCO GROSSETTI
Assistant Professor of Accounting Analytics and Data Science, Bocconi University

er what it used to be

Why we trust calculato

ors more than AI

AI creates discomfort in humans, despite our lives being immersed in algorithms: Elena Esposito, author of a book published by Bocconi University Press, explains that this depends on the fact that we think we are interacting with an actual intelligence. But since this is not the case, we should rather talk of artificial communication rather than intelligence

by Camillo Papini @

Artificial Intelligence may seem similar to a calculator but the relationship of humans with the former is not as serene as with the latter. We trust the results of computers, even if we don't know exactly how it can arrive at the result of a complex operation in a short time, while the relationship with Artificial Intelligence (AI) generates discomfort in people. Why? The reason is that machines never stop learning and the more they perform new and unexpected tasks, hitherto entrusted to human intelligence, the more users distrust them, because they do not like to find their own prerogatives embodied in machines. This is what creates discomfort. There is a fear of being overtaken by some super-intelligence and, staying in the economic sphere, of being replaced in the workplace by machines. And yet, we are immersed every day in algorithms and we deal with programs such as bots in our roles of telephone customers or bank savers. It is believed that bots are behind about 50% of internet traffic, that 40% of Wikipedia edits are the work of digital agents, without forgetting the presence on social networks of many accounts created by automated platforms and, last but not least, the boom in popular curiosity for ChatGPT, capable of producing written texts based on conversations with users.

In reality, "Algorithms work and evolve precisely because they no longer try to be intelligent. If anything, they can be seen as capable of communicating creatively and for informational purposes, but they are cannot be considered intelligent", explains **Elena Esposito**, Full Professor of Sociology of Cultural and Communicative Processes at Bielefeld and Bologna. She recently put out the book Comunicazione artificiale. *Come gli algoritmi producono intelligenza sociale* (*Artificial communication. How algorithms produce social intelligence)*, Bocconi University Press, 2022. "In fact, the purpose with which algorithms are programmed is not to understand the data provided by our online behavior. The intention is to identify correlations between data and process them so that they are informative for users", underlines Esposito, a student of Niklas Luhmann, who not by chance proposes to move from the definition of Artificial Intelligence to that of Artificial Communication, positing a new theoretical model to reiterate that the interlocutor with whom we interact is not a human being, but an algorithm. We need new rules and habits of behavior which, given the AI's multiple areas of application, must be promoted by national and supranational institutions, as well as families and individuals alike, according to their respective realms of pertinence.

ELENA ESPOSITO
Full Professor of Sociology of Cultural and Communicative Processes, Bielefeld University and Unioversità degli Studi di Bologna

➜ *How can talking about Artificial Communication reduce the discomfort felt towards AI?*

First, because I hypothesize that the analogy between the performance of algorithms and human intelligence, which generates this discomfort, is

misleading. Furthermore, because it allows for the emergence of new insights on the challenges and paradoxes that recent technologies pose. The numerous positive aspects of algorithmic intelligence remain, from the availability of more information to the higher speed with which to find it, passing through the cost-effectiveness of the process. But we can also ask ourselves how the growing intervention of AI affects, for example, our conception of the public sphere and the maintenance of social cohesion, considering the progressive customization of the information and services offered to each individual person, without he or she having even asked for them. This creates a bubble that is difficult to get out of. It becomes more difficult to realize that there may be something different from what you already know and decide whether you want to find it out or not. In other words, the individual no longer knows what others know and that common ground of shared information that makes everyone feel part of the societal whole decays. But social cohesion and markets themselves vitally depend on that common ground of shared. In addition, one may wonder what are the effects of the different versions of AI in specific fields such as education or, finally, how our perception of the relationship between reality and fiction changes. More and more often, in fact, we can intervene not only on reality but also on fiction, which is no longer the unalterable fiction of commercial movies or novel, but an area with which we can interact during the course of the story, such as for example happens in video games.

➜ Does the growing presence of AI in our lives change our faculties? For example, in knowing how to remember and what to forget?

Let's say that there is a new relationship between people and oblivion. In the past we mostly committed ourselves to remembering things and forgetting proceeded by itself, it intervened spontaneously to select the information that shouldn't last over time. Now the difficulty is reversed and lies in remembering not to remember; you need to try harder to forget as all memories and information are preserved online. We can therefore reason on the final paradox we arrive at: to forget memories it is better to multiply them, to make one climb from the first place to the eleventh in the results of a search engine, since we know that people tend to read only the first few results.

➜ Is it believable or illusory that machines can predict our future?

The future will forever be unwritten because it depends on human behavior, which is constantly changing. The future remains open, even if we can remember that AI offers a series of new tools to deal with the uncertainty of the future. However, if so far we have tried to anticipate it by relying on the calculation of probabilities, now the algorithms try to identify correlations between various possible configurations in the big data sets. And the correlations highlighted are not necessarily the most likely, even if the algorithms are partly based on probabilistic structures. The result of trying to predict the future is that algorithms produce indications about the future that are obscure to humans, since they are unable to understand how they were generated. These are predictions that end up recalling the divinatory practices of the ancient world, with their sibylline and cryptic responses. Precisely the practices from which science had to move away from. ■

THE BOOK

Big data, machine learning and bots, there are some of the terms that describe the hi-tech scenario in which we are immersed. The implications of these technologies currently unfolding in our lives are vast. Thus questions arise such as "Will we able to control something that we do not fully understand" or "Aren't the machines just getting too smart?". Elena Esposito, Full Professor of Sociology of Culture and Communication at the University of Bielefeld and the University of Bologna, answers these questions in his latest book *Comunicazione artificiale. Come gli algoritmi producono intelligenza sociale (Artificial communication. How algorithms produce social intelligence)*, Bocconi University Press, 2022, €22.


A leading
MBA
for tomorrow's
leaders
Ranking 2023
FT FINANCIAL TIMES
6th worldwide | 3rd in Europe
1-Year Full-Time MBA.
For people who just won't stop.
sdabocconi.it/mba
SDA Bocconi
SCHOOL OF MANAGEMENT

Intelligenza Artificiale: motore di innovazione

Segui Egea su

Egea

www.egeaeditore.it