

INTERNATIONAL MONETARY FUND

MONEY AND CAPITAL MARKETS AND
INFORMATION TECHNOLOGY DEPARTMENTS

Powering the Digital Economy

Opportunities and Risks of Artificial Intelligence in Finance

Prepared by El Bachir Boukherouaa and Ghiath Shabsigh
in collaboration with
Khaled AlAjmi, Jose Deodoro, Aquiles Farias, Ebru S. Iskender,
Alin T. Mirestean, and Rangachary Ravikumar

DP/2021/024

2021
OCTOBER



DEPARTMENTAL PAPER

INTERNATIONAL MONETARY FUND

**MONEY AND CAPITAL MARKETS AND
INFORMATION TECHNOLOGY DEPARTMENTS**

DEPARTMENTAL PAPERS

Powering the Digital Economy

**Opportunities and Risks of Artificial Intelligence in
Finance**

Prepared by El Bachir Boukherouaa and Ghiath Shabsigh

in collaboration with

Khaled AlAjmi, Jose Deodoro, Aquiles Farias, Ebru S. Iskender, Alin T. Mirestean, and
Rangachary Ravikumar

Copyright© 2021 International Monetary Fund

Powering the Digital Economy: Opportunities and Risks of Artificial Intelligence in Finance
DP/2021/024

Authors: El Bachir Boukherouaa and Ghiath Shabsigh
in collaboration with

Khaled AlAjmi, Jose Deodoro, Aquiles Farias, Ebru S. Iskender, Alin T. Mirestean, and Rangachary Ravikumar¹

Cataloging-in-Publication Data
IMF Library

Names: Boukherouaa, El Bachir. | Shabsigh, Ghiath. | AlAjmi, Khaled. | Deodoro, Jose. | Farias, Aquiles. | Iskender, Ebru S. | Mirestean, Alin. | Ravikumar, Rangachary. | International Monetary Fund, publisher.

Title: Powering the digital economy : opportunities and risks of artificial intelligence in finance / prepared by El Bachir Boukherouaa and Ghiath Shabsigh in collaboration with Khaled AlAjmi, Jose Deodoro, Aquiles Farias, Ebru S. Iskender, Alin T. Mirestean, and Rangachary Ravikumar.

Description: Washington, DC : International Monetary Fund, 2021. | 2021 September. | Departmental paper series. | Includes bibliographical references.

Identifiers: ISBN 9781589063952 (paper)

Subjects: LCSH: Artificial intelligence—Economic aspects. | Machine learning —Economic aspects. | Financial services industry—Technological innovations.

Classification: LCC HC79.I55 B68 2021

ISBN	978-1-59806-395-2 (Paper)
JEL Classification Numbers:	C40, C510, C550, E17, G21, G23, G280, O310, O330
Keywords:	Artificial Intelligence, Machine Learning, Financial Stability, Embedded Bias, Financial Regulation, Cybersecurity, Risk Management, Data Privacy
Author's E-Mail Address:	GShabsigh@imf.org; EBoukherouaa@imf.org; KAlAjmi@imf.org; JDeodoro@imf.org; AFarias@imf.org; ESonbullskender@imf.org; AMirestean@imf.org; RRavikumar@imf.org

The Departmental Paper Series presents research by IMF staff on issues of broad regional or cross-country interest. The views expressed in this paper are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Publication orders may be placed online or through the mail:

International Monetary Fund, Publication Services

P.O. Box 92780, Washington, DC 20090, USA

T. +(1) 202.623.7430

publications@imf.org

IMFbookstore.org

elibrary.IMF.org

¹ We are grateful to Aditya Narain and other IMF colleagues for valuable comments, and to Javier Chang for production support.

Executive Summary

This paper discusses the impact of the rapid adoption of artificial intelligence (AI) and machine learning (ML) in the financial sector. It highlights the benefits these technologies bring in terms of financial deepening and efficiency, while raising concerns about its potential in widening the digital divide between advanced and developing economies. The paper advances the discussion on the impact of this technology by distilling and categorizing the unique risks that it could pose to the integrity and stability of the financial system, policy challenges, and potential regulatory approaches. The evolving nature of this technology and its application in finance means that the full extent of its strengths and weaknesses is yet to be fully understood. Given the risk of unexpected pitfalls, countries will need to strengthen prudential oversight.

AI and ML are technologies with the potential for enormous societal and economic impact, bringing new opportunities and benefits. Recent technological advances in computing and data storage power, big data, and the digital economy are facilitating rapid AI/ML deployment in a wide range of sectors, including finance. The COVID-19 crisis has accelerated the adoption of these systems due to the increased use of digital channels.

AI/ML systems are changing the financial sector landscape. Competitive pressures are fueling rapid adoption of AI/ML in the financial sector by facilitating gains in efficiency and cost savings, reshaping client interfaces, enhancing forecasting accuracy, and improving risk management and compliance. AI/ML systems also offer the potential to strengthen prudential oversight and to equip central banks with new tools to pursue their monetary and macroprudential mandates.

These advances, however, are creating new concerns arising from risks inherent in the technology and its application in the financial sector. Concerns include a number of issues, such as embedded bias in AI/ML systems, the opaqueness of their outcomes, and their robustness (particularly with respect to cyber threats and privacy). Furthermore, the technology is bringing new sources and transmission channels of systemic risks, including greater homogeneity in risk assessments and credit decisions and rising interconnectedness that could quickly amplify shocks.

AI/ML in finance should be broadly welcome, together with preparations to capture their benefits and mitigate potential risks to the financial system's integrity and safety. Preparations include strengthening the capacity and monitoring frameworks of oversight authorities, engaging stakeholders to identify possible risks and remedial regulatory actions, updating relevant legal and regulatory, and expanding consumer education. It is important that these actions are taken in the context of national AI strategies and involve all relevant public and private bodies.

Cooperation and knowledge sharing at the regional and international level is becoming increasingly important. This would allow for the coordination of actions to support the safe deployment of AI/ML systems and the sharing of experiences and knowledge. Cooperation will be particularly important to ensure that less-developed economies share the benefits.

Contents

Executive Summary	1
Acronyms and Abbreviations	4
1. Introduction	5
2. Artificial intelligence in the Financial Sector	7
A. Forecasting	7
B. Investment and Banking Services	7
C. Risk and Compliance Management	9
D. Prudential Supervision	9
E. Central Banking	12
3. Risks and Policy Considerations	14
A. Embedded Bias	14
B. Unboxing the “Black Box”: Explainability and Complexity	15
C. Cybersecurity	16
D. Data Privacy	17
E. Robustness	17
F. Impact on Financial Stability	18
4. Conclusion	20
Annexes	
Annex 1. How Machine Learning Algorithms Work	21
Annex 2. Artificial Intelligence in Finance—Risk Profile	24
Annex 3. National Artificial Intelligence Strategies	25
References	28
BOXES	
Box 1. Artificial Intelligence and Machine Learning Capabilities	6
Box 2. Artificial Intelligence in Investment Management—Sample Use Cases	8
Box 3. Artificial Intelligence in Credit Underwriting	8
Box 4. Artificial Intelligence in Regulatory Compliance—Sample Use Cases	10
Box 5. Artificial Intelligence in Supervision—Sample Applications	11
Box 6. Artificial Intelligence in Central Banking—Sample Applications	13
Box 7. Explaining the “Black Box”	16
FIGURES	
Figure 1. Top Five Technologies Employed in Regulatory Technology Offerings	9
Figure 2. Technologies Used in Supervisory Technology Tools	10
Annex Figure 1.1. Machine Learning Paradigms	22

Annex Figure 1.2. Example of an Input Attack.....	23
Annex Figure 3.1. National Artificial Intelligence Strategy Landscape.....	25
Annex Figure 3.2. Key Features of National Artificial Intelligence Strategies	26

Acronyms and Abbreviations

AI	Artificial Intelligence
AML/CFT	Anti-Money Laundering/Combating the Financing of Terrorism
Fintech	Financial Technology
ML	Machine Learning
NLO	Natural Language Processing
OECD	Organisation for Economic Co-operation and Development
Regtech	Regulatory Technology
Suptech	Supervisory Technology

1. Introduction

This paper explores the use of artificial intelligence (AI) and machine learning (ML) in the financial sector and the resultant policy implications.¹ It provides a nontechnical background on the evolution and capabilities of AI/ML systems, their deployment and use cases in the financial sector, and the new challenges they present to financial sector policymakers.

AI/ML systems have made major advances over the past decade. Although the development of a machine with the capacity to understand or learn any intellectual task that a human being performs is not within immediate grasp, today's AI systems can perform quite well tasks that are well defined and normally require human intelligence. The learning process, a critical component of most AI systems, takes the form of ML, which relies on mathematics, statistics, and decision theory. Advances in ML and especially in deep learning algorithms are responsible for most of the recent achievements, such as self-driving cars, digital assistants, and facial recognition.²

The financial sector, led by financial technology (fintech) companies, has been rapidly increasing the use of AI/ML systems (Box 1). Recent adoption by the financial sector of technological advances, such as big data and cloud computing, coupled with the expansion of the digital economy, made the effective deployment of AI/ML systems possible. A recent survey of financial institutions (WEF 2020) shows that 77 percent of all respondents anticipate that AI will be of high or very high overall importance to their businesses within two years. McKinsey (2020a) estimates the potential value of AI in the banking sector to reach \$1 trillion.

AI/ML capabilities are transforming the financial sector.³ AI/ML systems are reshaping client experiences, including communication with financial service providers (for example, chat bots), investing (for example, robo-advisor), borrowing (for example, automated mortgage underwriting), and identity verification (for example, image recognition). They are also transforming the operations of financial institutions, providing significant cost savings by automating processes, using predictive analytics for better product offerings, and providing more effective risk and fraud management processes and regulatory compliance. Finally, AI/ML systems provide central banks and prudential oversight authorities with new tools to improve systemic risk surveillance and strengthen prudential oversight.

The COVID-19 pandemic has further increased the appetite for AI/ML adoption in the financial sector. BoE (2020) and McKinsey (2020b) find that a considerable number of financial institutions expect AI/ML to play a bigger role after the pandemic. Key growth areas include customer relationship and risk management. Banks are exploring ways to leverage their experience of using AI/ML to handle the high volume of loan applications during the pandemic to improve their underwriting process and fraud detection. Similarly, supervisors relying on off-site intensive supervision activities during the pandemic could further explore AI/ML-supported tools and processes in the post-pandemic era.

The rapid progress in AI/ML development could deepen the digital divide between advanced and developing economies. AI/ML deployment, and the resulting benefits, have been concentrated largely in advanced economies and a few emerging markets. These technologies could also bring significant benefits to developing economies, including enhanced access to credit by reducing the cost of credit risk assessments, particularly in countries that do not have an established credit registry (Sy and others 2019). However, these economies are falling behind, lacking

¹ Following the Oxford Dictionary, AI is defined as the theory and development of systems able to perform intellectual tasks that usually require human intelligence. ML is the learning component of an AI system, and is defined as the process that uses experience, algorithms, and some performance criterion to get better at performing a specified task. Given that AI and ML heavily overlap and that most statements in this paper hold true for both concepts, the terms are often used as a pair (AI/ML).

²See Annex 1 for more details.

³This includes revenue gains and cost savings.

the necessary investment, access to research, and human capital.⁴ Bridging this gap will require developing a digital-friendly policy framework anchored around four broad policy pillars: investing in infrastructure; investing in policies for a supportive business environment; investing in skills; and investing in risk management frameworks (IMF 2020).

Cooperation among countries and between the private and public sectors could help mitigate the risk of a widening digital divide. So far, global initiatives—including the development of principles to mitigate ethical risks associated with AI (UNESCO 2021; OECD 2019), calls for cooperation on investing in digital infrastructure (see, for example, Google and International Finance Corporation (2020)), and the provision of access to research in low-income countries (see, for example, AI4Good.org)—have been limited. Multilateral organizations could play an important role in transferring knowledge, raising investments, building capacity, and facilitating a peer-learning approach to guide digital policy efforts in developing economies. Similarly, the membership in several intergovernmental working groups on AI (such as the Global Partnership on Artificial Intelligence and the OECD Network of Experts on AI, among others) could be expanded to include less-developed economies.

AI/ML adoption in the financial sector is bringing new unique risks and challenges that need to be addressed to ensure financial stability. AI/ML-based decisions made by financial institutions may not be easily explainable and could potentially be biased. AI/ML adoption brings in new unique cyber risks and privacy concerns. Financial stability issues could also arise with respect to the robustness of the AI/ML algorithms in the face of structural shifts and increased interconnectedness through widespread reliance on few AI/ML service providers. Chapter 2 explores the adoption of AI/ML in the financial sector and possible associated risks, Chapter 3 discusses related policy concerns, and Chapter 4 provides some conclusions.

Box 1. Artificial Intelligence and Machine Learning Capabilities

- *Forecasting.* Machine learning algorithms are used for forecasting and benefit from using large data sets. They usually perform better than traditional statistical or econometric models.¹ In the financial sector, this is used in such areas as credit risk scoring, economic and financial variables forecasting, risk management, and so on.
- *Natural language processing.* Artificial intelligence systems can communicate by understanding and generating human language. Boosted by deep learning and statistical models, natural language processing has been used in the financial sector in such applications as chat bots, contract reviewing, and report generation.
- *Image recognition.* Facial and signature recognition is being used by some financial institutions and financial technology companies to assist with carrying out certain anti-money laundering/combating the financing of terrorism (AML/CFT) requirements (for example, the identification and verification of customers for customer due diligence process), and for strengthening systems security.
- *Anomaly detection.* Classification algorithms can be applied to detect rare items, outliers, or anomalous data. In the financial sector, insider trading, credit card and insurance fraud detection, and AML/CFT are some of the applications that leverage this capability (Chandola, Banerjee, and Kumar 2009).

⁴See Alonso and others (2020) for a broader discussion about possible implications of AI on developing economies. In particular, the paper finds that the new technology risks widening the gap between rich and poor countries by shifting more investment to advanced economies where automation is already established, with negative consequences for jobs in developing economies.

2. Artificial Intelligence in the Financial Sector

The capability of acquiring large sets of data from the environment and processing it with artificial intelligence (AI) and machine learning (ML) is changing the financial sector landscape. AI/ML facilitates enhanced capacity to predict economic, financial, and risk events; reshape financial markets; improve risk management and compliance; strengthen prudential oversight; and equip central banks with new tools to pursue their monetary and macroprudential mandates.

A. Forecasting

AI/ML systems are used in the financial sector to forecast macro-economic and financial variables, meet customer demands, provide payment capacity, and monitor business conditions. AI/ML models offer flexibility compared to traditional statistical and econometric models, can help explore otherwise hard-to-detect relationships between variables, and amplify the toolkits used by institutions. Evidence suggests that ML methods often outperform linear regression-based methods in forecast accuracy and robustness (Bolhuis and Rayner 2020).

While the use of AI/ML in forecasting offers benefits, it also poses challenges. Use of nontraditional data (for example, social media data, browsing history, and location data) in AI/ML could be beneficial in finding new relationships between variables. Similarly, by using AI natural language processing (NLP), unstructured data (for example, the information in email texts) can be brought into the forecasting process. However, the use of nontraditional data in financial forecasting raises several concerns, including the governing legal and regulatory framework; ethical and privacy implications; and data quality in terms of cleanliness, accuracy, relevancy, and potential biases.

B. Investment and Banking Services

In the financial sector, advances in AI/ML in recent years have had their greatest impact on the investment management industry. The industry has used technology for decades in trading, client services, and back-office operations, mostly to manage large streams of trading data and information and to execute high-frequency trading. However, AI/ML and related technologies are reshaping the industry by introducing new market participants (for example, product customization), improved client interfaces (for example, chatbots), better analytics and decision-making methods, and cost-reduction through automated processes (Box 2).

Compared to the investment management industry, the penetration of AI/ML in banking has been slower. The banking industry has traditionally been at the forefront of technological advancements (for example, through the introduction of ATMs, electronic card payments, and online banking). However, confidentiality and the proprietary nature of banking data have slowed AI/ML adoption. Nonetheless, AI/ML penetration in the banking industry has accelerated in recent years, in part on account of rising competition from financial technology (fintech) companies (including fintech lenders), but also fueled by AI/ML's capacity to improve client relations (for example, through chatbots and AI/ML-powered mobile banking), product placement (for example, through behavioral and personalized insights analytics), back-office support, risk management, credit underwriting (Box 3), and, importantly, cost savings.⁵

⁵The aggregate potential cost savings for banks from AI/ML systems is estimated at \$447 billion by 2023 (Digalaki 2021).

Box 2. Artificial Intelligence in Investment Management—Sample Use Cases¹

- Increased market liquidity provision through a wider use of high-frequency algorithmic trading and more efficient market price formation.
- Expanded wealth advisory services by providing personal and targeted investment advice to mass-market customers in a cost-effective manner, including for low-income populations.
- Enhanced efficiency with artificial intelligence and machine learning (AI/ML) taking on a growing portion of investment management responsibilities.
- More customized investment portfolios based on AI/ML targeted customer experiences.
- Development of new return profiles through the use of AI/ML instead of established strategies.

¹ See WEF (2018) for a more detailed discussion.

Box 3. Artificial Intelligence in Credit Underwriting

- Artificial intelligence/machine learning (AI/ML) predictive models can help process credit scoring, enhancing lenders' ability to calculate default and prepayment risks. Research finds that ML reduces banks' losses on delinquent customers by up to 25 percent (Khandani, Adlar, and Lo 2010). There is also evidence that, given their greater accuracy in predicting defaults, automated financial underwriting systems benefit underserved applicants, which results in higher borrower approval rates (Gates, Perry, and Zorn 2002), as does the facilitation of low-cost automated evaluation of small borrowers (Bazarbash 2019).
- AI/ML-assisted underwriting processes enable the harnessing of social, business, location, and internet data, in addition to traditional data used in credit decisions. AI/ML reduces turnaround time and increases the efficiency of lending decisions. Even if a client does not have a credit history, AI/ML can generate a credit score by analyzing the client's digital footprint (social media activity, bills payment history, and search engine activity). AI/ML also has the potential to be used in commercial lending decisions for risk quantification of commercial borrowers.¹ However, financial institutions and supervisors should be cautious in using and assessing AI/ML in credit underwriting and build robust validation and monitoring processes.

¹ See Bazarbash (2019) for a discussion of the potential strengths and weaknesses of AI/ML-based credit assessment.

AI/ML introduces new challenges and potential risks. The use of AI/ML in investment and banking depends on the availability of large volumes of good-quality, timely data. With the storage and use of large quantities of sensitive data, data privacy and cybersecurity are of paramount importance. Difficulties in explaining the rationale of AI/ML-based financial decisions is increasingly an important issue as AI/ML algorithms may uncover unknown correlations in data sets that stakeholders may struggle to understand because the underlying causality is unknown. In addition, these models may perform poorly in the event of major and sudden movements in input data resulting in the breakdown of established correlations (for example, in response to a crisis), potentially providing inaccurate decisions, with adverse outcomes for financial institutions or their clients.


C. Risk and Compliance Management

AI/ML advances in recent years are changing the scope and role of technology in regulatory compliance. Regulatory technology (regtech)⁶ has assumed greater importance in response to the regulatory tightening and rising compliance costs following the 2008 global financial crisis. For the most part, technology has been used to digitize compliance and reporting processes (Arner, Barberis, and Buckley 2017). However, advances in AI/ML over the past few years are reshaping risk and compliance management by leveraging broad sets of data, often in real time, and automating compliance decisions. This has improved compliance quality and reduced costs.

Maturing AI/ML technology has the potential to propel further adoption of regtech in the financial sector. According to a recent global survey, AI/ML is the top technology under consideration among regtech firms (Schizas and others 2019; Figure 1). Increased adoption of AI/ML in regtech has significantly expanded its use cases, cutting across banking, securities, insurance, and other financial services, and covering a wide variety of activities. These include identity verification, anti-money laundering/combatting the financing of terrorism, fraud detection, risk management, stress testing, microprudential and macroprudential reporting, as well as compliance with COVID-19 relief requirements (Box 4).

Regulators have generally been supportive of the adoption of regtech by regulated financial entities. Various regulators (for example, in Hong Kong Special Administrative Region) have developed strategies to promote the adoption of regtech that include boosting awareness, promoting innovation, and enhancing regulatory engagement within the regtech ecosystem. Even where there are no explicit strategies, many authorities have supported regtech adoption. For example, AI/ML systems, in conjunction with facial and voice recognition and NLP capabilities, could play an important role in helping digital banks secure licenses (for example, in Hong Kong Special Administrative Region and Malaysia) and for digital onboarding of customers.

Figure 1. Top Five Technologies Employed in Regulatory Technology Offerings
(Percent)



Source: Schizas and others (2019).

D. Prudential Supervision

Although decisions will ultimately depend on the judgment of supervisors, there is a place for AI/ML, primarily in data collection and data analytics. Many Financial Stability Board member country authorities are currently using ML and NLP tools in data analysis, processing, validation, and plausibility (FSB 2020). With AI, supervisors can draw deeper insights from any type of data and make more informed, data-driven decisions. AI can identify patterns that humans fail to spot and thereby enhance the quality of supervision. It can also make supervision more agile by flagging anomalies to supervisors in real time (ECB 2019). Additionally, supervisory technology (suptech)⁷ applications that leverage AI can provide predictive analyses, with the potential to improve the quality of supervision. Despite all its

⁶Regtech refers to the use of technologies by regulated financial entities to meet their regulatory requirements.

⁷Suptech refers to the use of technologies by supervisory agencies to support supervision.

Box 4. Artificial Intelligence in Regulatory Compliance—Sample Use Cases


- *Anti-money laundering/combating the financing of terrorism (AML/CFT) compliance.* Artificial intelligence/machine learning (AI/ML)-powered technologies, including analysis of unstructured data and consumer behavior, are used to reduce false positives, which constitute the vast majority of AML/CFT alerts, allowing financial institutions to devote more resources to cases that are likely to be suspect.
- *Mapping and updating regulatory obligations.* AI/ML-based applications are helping financial institutions identify and update relevant regulations, reducing costs and improving regulatory compliance in the process.
- *Conduct risk management.* AI/ML and natural language processing are used to monitor sales calls by financial institutions' employees to ensure compliance with regulatory requirements to accurately disclose the offered financial products' features and risks.
- *Stress-testing.* Many global banks are using AI/ML-enabled data analytics to improve the analysis of complex balance sheets and stress testing models to meet stress testing regulatory requirements.
- *COVID-19 relief.* AI/ML and robotic process automation are used to process COVID-19-mandated credit relief by processing large volumes of applications and loan documents to determine if the loans are eligible for payment modification terms.

potential, AI/ML is not a silver bullet, and the effectiveness of supervision will always very much depend on human judgement and an organization's risk culture.

Currently, supotech use cases are mainly concentrated in misconduct analysis, reporting, and data management (Box 5). Virtual assistance, microprudential, macroprudential, and market surveillance make up a smaller share of cases (di Castri and others 2019). AI in market conduct supervision is used in the collection and analytics of structured and unstructured data,⁸ work-flow automation, and conduct risk profiling and early warnings. Off-site surveillance, on-site inspection, and complaints handling, are areas of interest for market conduct authorities (FinCoNet 2020).

Many supervisory authorities are actively exploring the use of AI/ML systems in their risk-based supervision process (Figure 2). AI/ML in microprudential supervision could improve the assessment of risks, including credit and liquidity risks, and governance and risk culture in financial institutions, allowing supervisors to focus on risk analysis and on forward-looking assessments. Risks and trends could be assessed at the sectoral level, improving the quality of macroprudential supervision. AI/ML systems are also applied for market surveillance purposes to detect collusive behavior and price manipulation in the

Figure 2. Technologies Used in Supervisory Technology Tools



Source: FSB (2020).

Note: DLT = distributed ledger technology; "Self-sovereign ID" refers to technologies that give individuals and businesses control of their digital identities.

⁸Unstructured data analytics are used in a number of areas, such as identifying potential violations of advertising guidelines, ascertaining key issues of concern and interest for consumers, and predicting potential harm to consumers in real time using information from social networks.

securities market—potential misconducts that can be especially hard to detect using traditional methods (IIF 2017). The greater use of remote working arrangements as a result of the COVID-19 pandemic is also driving authorities to use technology to improve the effectiveness of remote supervision. Drawing on these experiences, supervisors, post pandemic, might lean more on AI/ML-supported off-site supervision engagements.

Box 5. Artificial Intelligence in Supervision—Sample Applications

Banca d'Italia	- Exploring how loan default forecasting can benefit from the use of machine learning (ML) algorithms using data from different sources.
Banco de España	- Natural language processing (NLP) helps process institutions' environmental, social, and governance disclosures to improve the understanding of the domestic green economy. - Supervised ML assists with misconduct detection.
Bank of Russia	- Operating a system that allows for the analysis of retail loan portfolios using algorithms for processing large data arrays.
Bank of Thailand	- Using an artificial intelligence/ML system to analyze board meeting minutes of financial institutions, which is used by supervisors to assess the regulatory compliance of the board and give recommendations as part of the ongoing supervision.
De Nederlandsche Bank	- ML uses transactional data to detect networks of related entities to assess the exposure of financial institutions to networks of suspicious transactions.
European Central Bank	- Machine reading of the "fit and proper" questionnaire helps flag problems. - ML assists in early identification of distress in "less significant institutions." - NLP and ML are used to search for information in supervisory review decisions to facilitate the identification of emerging trends and clusters of risks.
Monetary Authority of Singapore	- Working on a project where credit risk assessments by supervisors are conducted using algorithms instead of sampling.
Oesterreichische Nationalbank	- ML and deep learning algorithms are used to predict the probability that a data set contains errors that need to be rectified by the reporting entity.

The use of AI/ML by supervisory authorities comes with challenges and risks that need to be carefully considered. The effectiveness of AI/ML-driven supervision will depend on data standardization, quality, and completeness, which could be challenging for authorities and regulated institutions alike, particularly when leveraging nontraditional sources of data, such as social media (FSB 2020). Resource and skills gaps could pose a challenge to the effective and safe adoption of AI/ML by supervisory authorities. The personnel pool could be expanded to include AI/ML specialists and data scientists. Finally, deploying AI/ML systems presents supervisors with risks, including those associated with, privacy, cybersecurity, outcome explainability, and embedded bias.

E. Central Banking

AI/ML applications could help central banks implement their mandates, but they have largely been slow to embrace the technology.⁹ AI/ML systems could help central banks improve their understanding of economic and financial developments, support better-tuned monetary and macroprudential policies, and improve central banks' operations. They could also strengthen systemic risk monitoring and potentially help predict the buildup of systemic risks and speed up crisis response.¹⁰ Notwithstanding these potential benefits, the use of AI/ML processes for policymaking should remain subject to good judgment. Furthermore, AI/ML could provide increased efficiency and better internal control opportunities for central banks, including monitoring of internal operations and resource allocation across functions.¹¹ The technology to develop these applications largely exists; central banks' slow adoption thereof can be attributed to cultural, political, and legal factors, and lack of adequate capacity (Danielsson, Macrae, and Uthemann 2020). Nevertheless, some central banks are exploring the possibilities of AI/ML (Box 6). Experiments and research by central banks have focused on improving the near-term forecasting capacity and the monitoring of market sentiment to inform policy decisions. Some central banks are also developing applications to improve internal processes and back-office functions (for example, cash management).

AI/ML use in central banking does not seem to raise significant concerns for the time being, but this could change with broader deployment. Expanded use of large nontraditional and unstructured data sets to produce analytics in support of monetary and macroprudential policymaking could expose a central bank to data biases that are usually avoided in sampling methods. ML algorithms are still susceptible to sudden structural shifts in the data from unforeseen events, and the resulting errors could undermine the effectiveness of a central bank's crisis monitoring and response capacity. Central banks could also face challenges in acquiring quality and representative data as well as data privacy and security issues.¹² These concerns would be amplified if the central bank does not have the resources and skills required to safely operate AI/ML or mitigate the risks associated with it.

⁹The focus of this section is on central banks' monetary and macroprudential functions. AI/ML could also be relevant for other, potential central bank functions, in such areas as consumer protection, financial integrity, financial inclusion, or even climate change.

¹⁰Conceivably, AI/ML could strengthen the central bank's rapid crisis response to a fast-moving financial crisis.

¹¹See Khan and Malaika (2021) for more details.

¹²For a more detailed discussion, see Doerr, Gambacorta, and Serena (2021).

Box 6. Artificial Intelligence in Central Banking—Sample Applications

Strengthening nowcasting

- Sveriges Riksbank has developed real-time indicators to support its policy analysis. These include investigating whether fruit and vegetable prices, scraped daily from the Internet and machine learning (ML)-processed into an index, could improve the accuracy of short-term inflation forecasts.
- Reserve Bank of New Zealand is experimenting with using ML to process large, real-time data sets of about 550 macroeconomic indicators to improve its nowcasts of GDP growth, with results so far outperforming comparable statistical benchmarks.

Assessing market sentiment

- Banca d'Italia has developed a real-time tracking system of consumer inflation expectations, using ML and textual analysis of millions of daily Italian Twitter feeds. The Twitter-based indicators seem to be highly correlated with standard statistical measures, provide superior forecasting for survey-based monthly inflation expectations than all other available sources, and accurately anticipate consumers' expectations.
- Bank Indonesia is testing ML-based techniques for identifying the expectation of stakeholders, scraped from published news articles, of Bank Indonesia's policy rate, in support of the monthly Board of Governor's meeting.

Monitoring uncertainty

- Banco de Mexico staff has built a sentiment-based risk index using an artificial intelligence (AI)/ML system to analyze Twitter messages in response to positive or negative shocks to the Mexican financial sector. The research finds that index shocks correlate positively with an increase in financial market risk, stock market, volatility, sovereign risk, and foreign exchange rate volatility.
- Banco Central de Chile staff has developed a daily-frequency index of economic uncertainty for Chile using AI/ML to analyze Twitter feeds. It tracks the level of general disagreement—a proxy for economic uncertainty—on economic developments and policies. The index shows significant spikes that coincide with several episodes of substantial economic uncertainty of both local and international origin.
- De Nederlandsche Bank is exploring whether AI is useful for detecting liquidity problems at banks in anticipation of potential deposit runs (Triepels, Daniels, and Heijmans 2018).

Improving internal processes

- Banco de España has developed an AI tool for sorting banknotes between fit and unfit for circulation.

3. Risks and Policy Considerations

The rapid deployment of AI/ML systems in finance will have significant impact that will require robust policy responses to ensure the integrity and safety of the financial system. Concerns are rising regarding a number of issues, such as embedded bias in AI/ML systems, the ability to explain the rationale of their decisions, their robustness (particularly with respect to cyber threats and privacy), and their potential impact on financial stability. This chapter discusses those concerns. Annex 2 provides a consolidated profile of risks arising from AI/ML use in the financial sector.

A. Embedded Bias

The growing use of AI/ML in the financial sector, which is highly regulated and where public trust is an essential factor, has stirred discussions on the risk of embedded bias. Friedman and Nissenbaum (1996) defines embedded bias as computer systems that systematically and unfairly discriminate against certain individuals or groups of individuals in favor of others.¹³ Customer categorization processes used in AI/ML can lead to bias in the financial sector through differentiation in pricing or service quality. Bias in AI/ML decisions is often the result of biased training data that comes from existing biased processes and data sets, which will teach AI/ML models to be biased too (Wang 2016). Data biases, or inaccurate and insufficient information, could result in greater financial exclusion and feed distrust for the technology, especially among the most vulnerable (Sahay and others 2020). Data collection could result in bias in two ways:

- The data used to train the system might be incomplete or unrepresentative. For example, predictive algorithms (such as for a loan approval) favor groups that are better represented in the training data, given that there will be less uncertainty associated with those predictions (Goodman and Flaxman 2016).
- The data may underpin prevailing prejudices (Hao 2019). For example, Amazon discovered that its internal recruiting tool was [dismissing female candidates](#) because it was trained on historical hiring decisions, which favored men over women.¹⁴

Human bias can give rise to bias in the algorithm during the design and training of AI/ML systems. For instance, when a researcher decides on the features to include or exclude in the ML model, the choice can be influenced by various psychological, social, emotional, and cultural factors, which in turn influence the researcher's feedback on the outputs when training the AI/ML system.¹⁵

Notwithstanding the potential for bias in AI/ML systems, they may help reduce prevailing biases. Mayson (2019) and Silberg and Manyika (2019) argue that AI/ML systems can reduce human bias in decision making because ML algorithms can eliminate irrational biases that stem from the subjective interpretation of data.¹⁶ Miller (2018) notes that AI systems, while susceptible to embedded bias, can still improve the decision-making process by mitigating human bias. Finally, Silberg and Manyika (2019) suggests that, even though many AI/ML systems could be perceived as black boxes, their prediction and decision-making processes could be scrutinized more than those of individuals and, thus, existing biases could be identified and mitigated.

¹³This bias should not be confused with statistical bias, which is defined as the difference between the expected value of the estimator and its true value (Lehmann 1951). For further discussion on bias-variance trade-off, see Annex 1.

¹⁴The algorithm was taught to recognize word patterns; AI software penalized any resume that contained the word "women."

¹⁵For an in-depth discussion of how psychological, social, emotional, and cultural factors already play a role in the financial sector, see Khan (2018).

¹⁶For an in-depth discussion on the relation between human bias and algorithms, see Plous (2002); Corbett-Davies and others (2017); and Kleinberg and others (2018a, 2018b, and 2019).

Given that bias may arise as an unintended consequence of AI/ML systems, regulators could view this as a potential source of operational and reputational risks. Financial institutions that deploy AI/ML systems in a significant manner, particularly with regard to the provision of credit and financial services and in risk management, should develop and implement bias mitigation and detection plans as part of their operational risk management framework. Such plans could include adequate assurances about the algorithm’s robustness against systemically generating biased decisions, disclosure of data sources, awareness of potential bias-generating factors in the data, monitoring and evaluation tools, and, more generally, how the institution intends to meet the standing anti-discrimination rules in the context of AI/ML deployment.

Policy responses to AI/ML embedded bias issues could be advanced by developing and deploying a broader framework for the governance and ethical use of AI/ML applications. In recent years, several initiatives have been launched to help develop such a framework. For example, in April 2019 the European Union released its “Ethics Guidelines for Trustworthy AI,” which outlines key requirements for a trustworthy AI system.¹⁷ Similarly, in May 2019 the Organisation for Economic Co-operation and Development adopted a set of principles to promote ethical and innovative AI development. These principles contributed to the June 2019 Group of Twenty declaration on “human-centered” AI.

B. Unboxing the “Black Box”: Explainability and Complexity

Explainability of AI/ML systems outcomes is an important issue, particularly when used in the financial sector. ML models are often referred to as black boxes because they are not directly explainable by the user (Guidotti and others 2019). This characteristic could make detection of the appropriateness of ML decisions difficult¹⁸ and could expose organizations to vulnerabilities—such as biased data, unsuitable modeling techniques, or incorrect decision making—and potentially undermine the trust in their robustness (Silberg and Manyika 2019).

Explainability is a complex and multifaceted issue. There are several reasons why ML models are frequently considered to be black boxes: (1) they are complicated and cannot be easily interpreted, (2) their input signals might not be known, and (3) they are an ensemble of models rather than a single independent model. Furthermore, stronger explainability might enable outsiders to manipulate the algorithm (Molnar 2021) and create risks in the financial system. Generally, there is a trade-off between model flexibility—which refers to its capacity to approximate different functions and is directly related to the number of parameters of the model—and its explainability. ML models are more flexible and accurate, but are less explainable than, for example, linear models that produce easily interpretable results but are less accurate. Box 7 provides a brief overview of ML explanation methods.

Ongoing research and regulatory initiatives may help address AI/ML explainability issues. The literature points to various levels of explainability related to different stages of the modeling process, or different objectives, such as explaining individual predictions or the overall model behavior (Guidotti and others 2019). To advance the policy discussion, regulators could explore the possibility that financial sector AI/ML models may require different levels of explainability, depending on the impact of the model’s outcome or the governing regulation. Broadly, regulatory guidance that creates a proper framework and strategy for managing explainability risks at different levels is needed for the financial sector. Currently, ML explainability is included in only a few regulatory frameworks as part of the

¹⁷Similarly, the Monetary Authority of Singapore published “Principles to Promote Fairness, Ethics, Accountability, and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore’s Financial Sector” in November 2018; De Nederlandsche Bank published “General Principles for Use of Artificial Intelligence in Finance” in July 2019; and the Hong Kong Monetary Authority published “High-Level Principles on Artificial Intelligence” in November 2019.

¹⁸In this context, “explainability” (sometimes referred to as “interpretability”) in a model means showing how input variables contribute to both the model’s aggregate results and explain individual outcomes (HKMA 2020).

accountability and transparency requirements (for example, in the European Union, Hong Kong Special Administrative Region, Netherlands, and Singapore).

Box 7. Explaining the “Black Box”

- *Explainability by design.* Some machine learning (ML) models are inherently interpretable by design, such as linear models and decision trees, but this may constrain their accuracy.
- *Model inspection.* During model development, various techniques could be used for testing the sensitivity of a model to changes in variables. This approach can help avoid unintended bias and other modeling pitfalls, however, the techniques have limited applicability to deployed models.
- *Individual prediction explanation.* After the model is deployed, some methods can provide local approximations that help explain individual predictions.¹ For instance, why credit was or was not granted to an individual or why a buy or sell order was posted, thereby making it easier to audit results.
- *Global model explanation.* This method consists of using an interpretable model, such as linear regression or decision tree, to explain the behavior of the more complex ML models used for prediction. The interpretable model would not be used for prediction, but to provide a general explanation of the behavior of the main model. While useful to comprehend the overall behavior of the predictive model, there is no assurance that the surrogate model can fully mimic the more complex model while keeping a reasonable level of explainability.
- *Models to explain ML models.* Research is increasingly focused on developing techniques that use ML models or their outcomes as training input to achieve explainability. This is accomplished by repeatedly disturbing the ML inputs or variables and comparing them to the ML outcome to ascertain the decision-making process.²

¹For example, Shapley values (Shapley 1953) and Lime values (Ribeiro, Singh, Guestrin 2016).

²See examples of explainability in Cloudera Fast Forward (2020).

C. Cybersecurity

AI/ML adoption increases the scope for cyber threats and brings in new unique cyber risks. In addition to traditional cyber threats from human or software failures, AI/ML systems are vulnerable to novel threats. Such threats focus on manipulating data at some stage of the AI/ML lifecycle to exploit inherent limitations of AI/ML algorithms (Comiter 2019). Such manipulation allows attackers to evade detection and prompts AI/ML to make the wrong decisions or to extract information. Due to their complexity and the potential impact for financial sector institutions, ML models require constant oversight to certify that such attacks are detected consistently and managed promptly. In addition to typical cybersecurity concerns, specific cyber threats to AI/ML can broadly be grouped as:

- Data poisoning attacks intend to influence an ML algorithm during the training stage by adding special samples to its training data set. These attacks cause the AI to incorrectly learn to classify or recognize information. Data poisoning also may be used to create Trojan models, which hide malicious actions that wait for special inputs to be activated (Liu, Dolan-Gavitt, and Garg 2018). Data poisoning attacks require privileged access to model and training information, but once executed successfully, and as long as the malicious behavior does not interfere in regular diagnostics tests, infected models may be undetectable.

- Input attacks¹⁹ allow attackers to introduce perturbations to data inputs and mislead AI systems during operations. For example, attackers could alter images with elements unperceivable to human vision, but which provoke AI/ML image recognition systems to mislabel the images.
- Model extraction or model inversion attacks attempt to recover training input and data, or the model itself. Membership inference (a variant) simply aims to check whether a specific data instance was in the training set. Such attacks may be performed as black-box attacks, whereby attackers have merely read-only access to the model, possibly through business application programming interfaces. Current privacy protection laws (for example, the European Union’s General Data Protection Regulation) are not designed to deal with the methods of model inversion attacks, raising privacy and copyright concerns.

AI/ML cybersecurity is increasingly becoming a concern for financial sector regulators. AI/ML cyber threats could undermine the integrity of, and trust in, the financial sector. Corrupted systems could undermine the financial sector’s capacity to accurately assess, price, and manage risks, which could lead to the buildup of unobserved systemic risks. Attackers could also acquire training data sets that contain sensitive financial and personal information.

The regulatory perimeter of cybersecurity requirements in the financial sector could be expanded to cover AI/ML-specific cyber threats. Providers and users of AI/ML applications in the financial sector should be required to put in place, as part of their broader cybersecurity framework, mitigating strategies. These could include detection and reporting systems, robust protection of training data feeds, and strategies to secure model and data privacy.

D. Data Privacy

AI/ML introduces new unique privacy issues. Privacy concerns regarding big data are well known and predate the emergence of AI/ML into the mainstream.²⁰ Tools have been developed to help maintain data anonymity and data subjects’ privacy. Legal data policy frameworks are being put in place across the globe to address these concerns. The robustness of AI/ML models, however, in preventing data leakage from the training data set raises new privacy concerns. For example, AI/ML has the capacity to unmask anonymized data through inferences (that is, deducing identities from behavioral patterns). Similarly, AI/ML may remember information about individuals in the training set after the data is used, or AI/ML’s outcome may leak sensitive data directly or by inference. Tools are being developed to address these issues and strengthen AI/ML models’ robustness in safeguarding sensitive data, but more work is needed, along with an appropriate update to the legal and regulatory framework that requires AI/ML systems and related data sources to adhere to enhanced privacy standards, along with relevant anti-money laundering/combating the financing of terrorism requirements.

E. Robustness

Robust AI/ML algorithms will help build public trust in an AI-driven financial system and safeguard financial stability and the integrity of the financial system. This is particularly important, given the relatively high concentration of AI/ML service providers and the lack of adequate oversight capacity in many jurisdictions to effectively engage financial institutions using AI/ML systems. Several challenges would need to be addressed by the industry and regulators to ensure the robustness of ML models with respect to cyber threats and safeguarding privacy, as discussed above, as well as with respect to their performance. The latter covers issues related to minimizing ML models’ false signals

¹⁹Input attacks are sometimes referred in specialized literature as “adversarial examples” (a form of adversarial attack). Furthermore, “input attack” should not be confused with “input validation attack,” a popular vulnerability often used to hack websites.

²⁰See Haksar and others (2021) on broader data policy issues.

during periods of structural shifts and having in place an appropriate governance over the AI/ML development process.

AI/ML systems in the financial sector have performed well in a data environment that is relatively stable and produces reliable signals, but that could quickly change in periods of rapid structural shift. In a relatively stable environment, AI/ML models are reasonably able to incorporate evolving data trends without significant loss in prediction accuracy. However, they face a more challenging task with structural changes in their data environment, when a previously reliable signal becomes unreliable or when behavioral correlations shift significantly. The recently observed misalignment of AI/ML-generated risk assessments during the COVID-19 pandemic is a good illustration. The crisis negatively affected the performance of ML models because they were not originally trained for such an event. Harker (2020) and BoE (2020) highlight that ML algorithms trained using pre-COVID-19 data may experience performance deterioration, for example, by pointing to aggregate credit scores in the United States, which, despite the record job loss and rising defaults, have improved during the acute part of the crisis, reflecting temporary relief measures that were not captured by the algorithms.

New governance frameworks for the development of AI/ML systems are needed to strengthen prudential oversight and to avoid unintended consequences. Given that the process of creating AI/ML systems is very similar to software development, proper quality control and agility of the process is required. The process could encompass all phases of development, testing, and deployment, and focus on the relevant risks and controls. Checks for embedded bias, data poisoning, other security risks, and performance should comply with best practices of software development, including separation of duties and constant monitoring.

F. Impact on Financial Stability

The widespread deployment of AI/ML systems in the financial sector will be transformational, and their impact on financial stability is yet to be fully assessed. As highlighted above, on one hand, with carefully designed and tested algorithms satisfying a high level of controls to limit risks and performance issues, AI/ML systems may bring increased efficiencies; better assessment, management, and pricing of risks; improved regulatory compliance; and new tools for prudential surveillance and enforcement—all of which will contribute positively to financial stability. On the other hand, AI/ML systems bring new and unique risks arising from the opacity of their decisions, susceptibility to manipulation, robustness issues, and privacy concerns. These could undermine the public's trust in the integrity and safety of an AI/ML-driven financial system. Furthermore, AI/ML could potentially bring about new sources and transmission channels of systemic risks. More specifically:

- AI/ML service providers could become systemically important participants in financial market infrastructure due to the high specialization of AI/ML systems and network effects, which could increase the financial system's vulnerability to single points of failure.
- The concentration of third-party AI/ML algorithm providers could drive greater homogeneity in risk assessments and credit decisions in the financial sector, which, coupled with rising interconnectedness, could create the conditions for a buildup of systemic risks. The likely concentration of data and growing use of alternative data in AI/ML could result in, respectively, the risk of uniformity (herding) and out-of-sample risk that eventually could lead to systemic risk.
- The widespread use of AI/ML could potentially increase procyclicality of financial conditions. For instance, credit underwriting and risk management processes are inherently procyclical, reacting to financial conditions in ways that reinforce and amplify those changes (that is, falling default rates lead to more lending, which can further lower default rates). AI/ML may automate and accelerate the procyclicality and potentially obscure it (due to explainability issues).

- Moreover, in the case of a tail risk event, an inaccurate risk assessment and reaction by ML algorithms could quickly amplify and spread the shock throughout the financial system and complicate or even undermine the effectiveness of the policy response.
- Challenges in interpretation, sustainability of analytical power, and prediction of models raise the concern that economic policies or market strategies based on these models will be challenging to interpret or predict by the relevant counterparties, creating additional asymmetric information in the market, with an uncertain impact on financial stability.
- Finally, regulatory gaps could adversely impact financial stability if technological advances outpace existing regulations. Such advances are often led by providers who may fall outside existing regulatory perimeters.²¹

The rapid evolution of AI/ML has led to a range of regulatory responses. While some jurisdictions have taken a more holistic approach to addressing the issues involved (for example, the Monetary Authority of Singapore²² and De Nederlandsche Bank²³), others have concluded that existing regulations and expectations on good governance are sufficient to address the emerging issues. Whether it is by way of new regulations or existing ones, regulators have generally focused on AI/ML governance frameworks, risk management, internal controls, and better controls over the model and data.

Addressing these challenges requires broad regulatory and collaborative efforts. An adequate policy response requires developing clear minimum standards and guidelines for the sector, coupled with stronger focus on securing the necessary technical skills. Given the inherent interconnectivity of issues related to the deployment of AI/ML systems in the financial sector, collaboration among financial institutions, central banks, financial supervisors, and other stakeholders is important to avoid duplication of work and to help counter potential risks. Many leading jurisdictions in the AI/ML sector have relied on well-articulated national AI strategies for promoting AI/ML development while ensuring that regulatory gaps do not materialize. Annex 3 provides an overview of country approaches to developing AI national strategies.

²¹For a more detailed discussion of new transmission channels of systemic risks, see Gensler and Bailey (2020).

²²See MAS (2018).

²³See DNB (2019).

4. Conclusion

The deployment of artificial intelligence (AI) and machine learning (ML) systems in the financial sector will continue to accelerate. This trend is driven by rapid increases in computational powers, data storage capacity, and big data, as well as by significant progress in modeling and use-case adaptations. The COVID-19 pandemic is accelerating the shift toward a more contactless environment and increasingly digital financial services, which will further strengthen the appeal of AI/ML systems to providers of financial services.

Use of AI/ML will bring important benefits but will also raise significant financial policy challenges. AI/ML systems offer financial institutions the potential for significant cost savings and efficiency gains, new markets, and better risk management; bring customers new experiences, products, and lower costs; and offer powerful tools for regulatory compliance and prudential oversight. However, these systems also bring about ethical questions and new unique risks to the financial system's integrity and safety, of which the full extent is yet to be assessed. The task facing financial sector policymakers is further complicated by the fact that these innovations are still evolving and morphing as new technologies come into play. These developments call for improvements in oversight monitoring frameworks and active engagement with stakeholders to identify possible risks and remedial regulatory actions.

In step with the Bali Fintech Agenda's call on national authorities to embrace the fintech revolution, regulators should broadly welcome the advancements of AI/ML in finance and undertake the preparations to capture its potential benefits and mitigate its risks. This includes the timely strengthening of institutional capacity, recruiting relevant expertise, building up knowledge, improving external communication with stakeholders, and expanding consumer education. Deployment of AI/ML systems in the financial sector has proven to be most effective when there are national AI strategies in place that involve all relevant public and private bodies.

Cooperation and knowledge sharing at the regional and international level is becoming increasingly important. This would allow for the coordination of actions to support the safe deployment of AI/ML systems and the sharing of experiences and knowledge. Cooperation will be particularly important to ensure that less-developed economies have access to knowledge related to techniques and methods, use cases, and regulatory and supervisory approaches.

Finally, the evolving nature of the AI/ML technology and its applications in finance mean that neither the users, the technology providers and developers, nor the regulators understand, currently, the full extent of the strengths and weaknesses of the technology. Hence, there may be many unexpected pitfalls that are yet to materialize, and countries will need to strengthen their monitoring and prudential oversight.

Annex 1. How Machine Learning Algorithms Work

While machine learning (ML) is a common fixture of today’s artificial intelligence (AI) systems, this has not always been the case. Early AI relied heavily on expert systems consisting of a set of precise rules created by human experts. These were rules that a computer could follow, step by step, often giving the impression that it was responding to changing situations. However, given that these rules were often expressed in an “if-then” format, they imposed heavy limitations, basically rendering such systems useless when facing situations outside of the prescribed environment. Advances in ML offer AI systems the ability to learn from experience, thus making them more adaptable to new states of the world.

The key element that constitutes the basis of ML is that all learning can be reduced to learning the representation of a mapping between some input and output. For example, if a classification problem is considered in which objects are split into predetermined classes, the objective of ML becomes learning a mapping between the features of the objects and the predetermined classes. An important fact to understand is that the mapping does not necessarily have to have a functional form, that is, a mathematical expression. There is an important distinction between fitting and learning processes. If the functional forms for the mapping were known, with the only unknowns being its parameters, then the problem becomes a fitting task, which means that all that is left to do is to estimate the parameters. When the functional form is not known, or does not exist, fitting is not helpful, and a learning process needs to be employed.

A ML model usually consists of the selected relevant features from the input data along with a learning algorithm and performance criteria. Regardless of the learning paradigm, the learning process requires the selection of a set of features from the available data on which a learning algorithm, or even a set of learning algorithms, is applied. The algorithms are evaluated based on some performance criteria that vary with the problem for which a solution needs to be learned. The combination of features, learning algorithms, and performance criteria can be thought of as the ML model that provides the learned mapping from the inputs to the outputs.

The learning paradigms depend on the data available for the learning process:

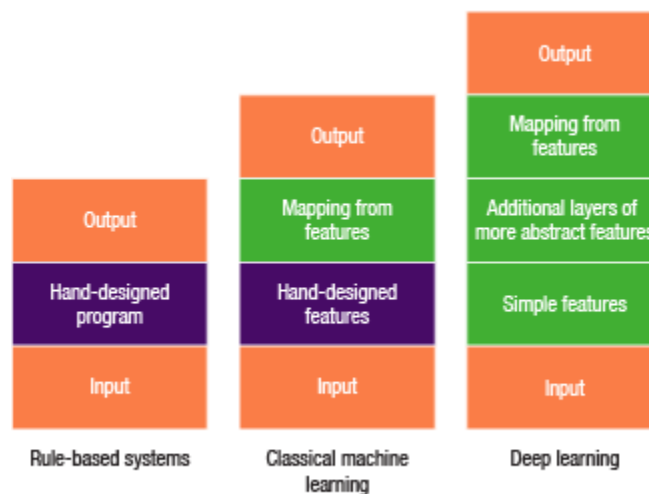
- *Supervised learning* requires that the training data contain both the inputs (or features) and the correct associated outputs (labels, or targets). In this case, the objective of ML is to find a mapping between the known inputs and outputs through the use of an algorithm. The idea is that the mapping (or the model)¹ learned this way can be used on unseen inputs and accurately predict the outputs. Overfitting is one of the main challenges in ML and it occurs when the model performs very well on the training data set, but poorly on new, unseen data (test data). The better a model performs on the training data, the less biased it is. However, its performance can change dramatically on a set of unseen data. In this case, in order to achieve robustness in performance, the researcher may opt to choose a model that does not perform as well on the training set but maintains its performance on the test set. This situation is known as the bias-variance trade-off.
- *Unsupervised learning* occurs when the available data does not contain the outputs. Given that only the features are available (and without an associated label), the algorithm employed will not only have to learn the mapping, but also generate the labels. For example, in anomaly detection problems, no prior information exists about the data points that represent an anomaly. In this situation, using a clustering algorithm is a common approach to learn a viable mapping. Under the assumption that all normal activity occurs in clusters, the data points that are identified as belonging to clusters can be labeled as normal activity, while the rest of the data points will be labeled as anomalies.

¹In practice, it is common to refer to the mapping as a model; the terms are used interchangeably from here on.

- *Reinforcement learning* is used in situations when the available data is not fixed, and it changes through a feedback loop from the environment to the AI system. In addition, the labels or outputs associated with inputs are not necessarily correct or desired. The most common example of reinforcement learning is learning to play a game, such as chess.

Deep learning is an approach to AI/ML that has been responsible for most of the recent developments of AI/ML systems, including image and speech recognition (Hinton, LeCun, and Bengio 2015). The main difference between deep learning and the classical ML approach resides in the way they pre-process the raw input data before feeding it into the algorithm (Annex Figure 1.1). Classical ML modeling requires the modeler to transform the original raw data into a set of variables (features) that can be used by the AI/ML system to train the model. For instance, if the task is to recognize handwritten numbers, the modeler would have to transform the image pixels into a set of variables, such as curvature, size, and density, and then use this transformed data to train a ML model. Deep learning, conversely, embeds this process into the AI/ML system, which would be trained using the raw data. One of its tasks would be to extract from the data the important features required to maximize its performance criteria. For instance, considering the aforementioned task, the modeler would feed the images directly into the system and it would select the features set that works best (Goodfellow, Bengio, and Courville 2016).

Annex Figure 1.1. Machine Learning Paradigms



Sources: IMF staff; adapted from Goodfellow, Bengio, and Courville (2016).

Another feature of deep learning models is that they are comprised of several learning layers, where each layer is responsible for representing concepts that are expressed in terms of previous layers that represent simpler concepts (Hinton, LeCun, and Bengio 2015). For example, the first layer of the system can recognize limbs, the second one may recognize human limbs, and the third and final one can leverage these previous layers and recognize humans in a picture. The number of layers is related to the depth of the deep learning system. In general, the deeper the model, the more flexible (and more complex) it is. On one hand, it enables great performance in complex tasks, such as face recognition and translation but, on the other hand, it is usually referred to as a black box system and is at the lower end of explainable models (Guidotti and others 2019). The depth and complexity of such models has implications for the security of the system, given that it could be more prone to adversarial attacks, which consist of building examples with small perturbations that can easily trick the AI system (Goodfellow, Shlens, and Szegedy 2014). A well-known example is in Annex Figure 1.2, where a human imperceptible perturbation is added to a panda picture that tricks the AI to classify it as a gibbon with high level of confidence.

Natural language processing (NLP) is a set of computational techniques that allow machines to learn, understand, and produce human language contents. It dates to the 1940s, when scientists began to work with populating computers with vocabularies and human language rules. From the 1990s onward scientists started to favor the currently used statistical or corpus-based methods largely because of the difficulties of interpreting human language in a structured manner (Hirshberg and Manning 2015). Nevertheless, NLP suffers from a major limitation: only a limited set of the languages currently spoken generate online material in the necessary magnitude to train algorithms. Even though many of the remaining languages are spoken and written by a large share of the world's


population, they lack proper online resources for NLP algorithms, yielding inaccurate models for such tasks as translations or conversational agents on such languages, therefore contributing to a widening of the digital divide. Researchers aim to rectify this situation by developing multilingual NLP using unsupervised models and novel techniques to transfer resources between various languages (Ponti and others 2019).

NLP is currently used in several applications, such as:

- To facilitate human-to-human communication via automated language translation, a capability that was boosted by the exceptional quantities of parallel text brought by the Internet in the 1990s. In a context-aware blind evaluation in 2020, human judges thought a deep learning system outperformed news translation in accuracy, with comparable fluency.
- To provide spoken dialogue systems and conversational agents, leading to such devices as Apple Siri or Amazon Alexa personal assistants. Although dialogue management and contextualization are generally deficient at this moment, conversational technology is widespread and performs well under suitable conditions, such as when conversation topics are known in advance.
- To extract knowledge by reading and understanding free text. Such capability may be used to create structured information databases from textual records. For instance, in 2015 researchers were able to predict drug-gene relationships by using algorithms to read and analyze over 23 million articles containing the results of biomedical research abstracts (Percha and Altman 2015).
- To analyze social media, by retrieving information from sources such as Twitter, Facebook, YouTube, among other. Even though data from these sources is often unreliable, it enables the extraction of useful analytics from language used by users in forums and discussions, such as demographics. Furthermore, it is often used to analyze speaker states: the opinions, beliefs, emotions, and other personal views expressed in language. This is usually done through sentiment analysis, whereby one may determine the positive or negative emotional attitude of subjects toward a predetermined subject.

Annex Figure 1.2. Example of an Input Attack

In an input (or adversarial example) attack, researchers add an imperceptible perturbation (middle column, pictures magnified by 10×) to a previously recognized image (left column), causing the machine learning model to misclassify all the resulting images as ostriches.



Source: Adapted from "Adversarial examples generated for AlexNet" by Szegedy and others (2014), licensed under CC BY 3.0.

Annex 2. Artificial Intelligence in Finance—Risk Profile

Sources of Risk	Risk Category	Risk Materialization
Data	Quality and availability	Inaccurate or suboptimal outcome
	Privacy	Data leaks undermining public trust
	Embedded bias	Discriminatory practices
Algorithm	Performance robustness	Inability to recognize structural shifts in the data
	Explainability	Difficulties in explaining or detecting the appropriateness of machine learning (ML) decisions
	Privacy	ML operations creating inadvertent data leaks undermining public trust
	Embedded bias	Discriminatory practices
Cybersecurity	Development governance	Inadequate deployment of ML systems in production environments
	Traditional cyber threats	Human or software failures
	Data poisoning	Corrupting artificial intelligence (AI)/ML training
	Input manipulation	Misleading AI/ML systems during operations
Technology Management	Model extraction/inversion	Recover training data, or the model itself
	Operational capacity	Inadequate deployment, operation, and risk management of AI/ML systems
Systemic	Outsourcing AI/ML services	Exposure to third-party risks
	Network effects	Higher concentration of AI/ML service providers who could become systemically important
	Risk assessment	Buildup of systemic risk due to greater homogeneity in risk assessments and credit decisions
	Interconnectedness	Inaccurate ML risk assessment and reaction could quickly amplify and spread in the financial system
	Regulatory gaps	Development outpaces regulations and providers who may fall outside regulatory perimeter

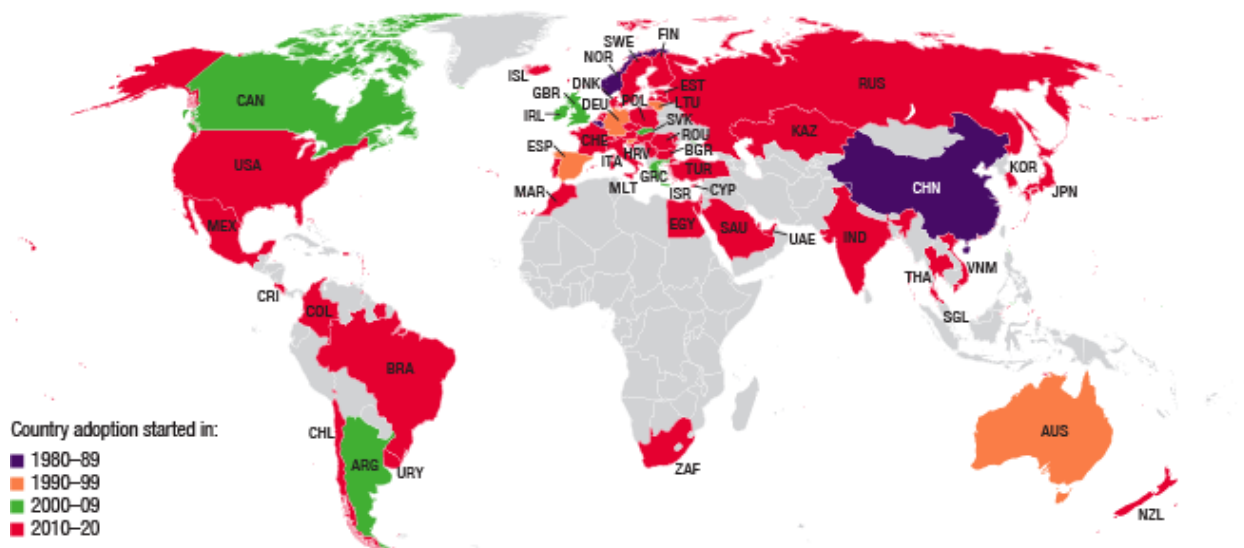
Annex 3. National Artificial Intelligence Strategies

This annex provides an overview of national artificial intelligence (AI) strategies. The overview is intended to highlight efforts and approaches to developing AI strategies across the globe as well as priorities and key lessons. This overview is not intended to provide details on specific national AI strategies or views on their qualities.

Considering the promise and complexity of AI, a rising number of countries are developing national AI strategies. About 60 countries around the world have developed national AI strategies (most since 2010 (OECD 2021)). Some national authorities have pursued overarching multiyear strategies of AI implementation that are intended to attract private investment, foster innovation, and develop a skilled workforce for the future. Success indicators often include mapping AI national strategies to the Sustainable Development Goals adopted by United Nations member states in 2015 (Stanford University 2019).

Developing national strategies on AI is not new but has accelerated significantly in recent years (Annex Figure 3.1). Recent years have witnessed a rapid increase in the number of countries developing national AI strategies, driven by technological advances that are facilitating fast deployments of AI systems across a wide range of sectors (for example, security, financial, health).

Annex Figure 3.1. National Artificial Intelligence Strategy Landscape



Source: IMF staff, based on information in OECD (2021).
 Note: Data labels use International Organization for Standardization (ISO) country codes.

Approaches to national AI strategies vary across countries, reflecting differences in national priorities, as well as available skills and resources. The coverage of a national strategy on AI is defined by the number of policy initiatives a country develops and the spread of these policies across multiple themes (for example, finance, education, health, climate change). AI strategies are sometimes embedded in broader science and technology initiatives. Furthermore, AI strategies are linked to strategies and regulations related to, among others, data access, sharing and privacy frameworks, intellectual property management, national security, and ethical use of AI. Regulations, such as the US Health Insurance Portability and Accountability Act and General Data Protection Regulation, will need to be continuously reviewed to address new concerns arising from AI systems deployment.

OECD Artificial Intelligence Policy Observatory data show that 60 countries and the European Union have developed AI strategies that encompass, collectively, more than 600 policy initiatives covering a broad range of issues. The widest coverage is attributed to the United States, where more than 47 initiatives have been developed on AI usage and coverage. Regionally, the European Union has more coverage than any other region, with 51 initiatives.¹ Common elements of these strategies touch on leadership and vision, focus and specialization, partnership and collaboration, including with academia, AI research and development, developing human capital, governance, and risk management (Annex Figure 3.2).

A review of current AI national and regional strategies reveals several common lessons in designing and implementing AI development strategies:

- Designing and implementing a successful AI strategy is linked to clear objectives, often the United Nations Sustainable Development Goals. This linkage enables robust assessment and monitoring, which in turn translates to valuable outcomes on national, regional, and global levels.
- There is a need for enhanced open access to research on AI. Publishing and sharing AI research, including data and code, would aid practitioners and policymakers.
- Stronger emphasis should be placed on developing human capital. Those skilled in AI are heavily recruited by major companies, resulting in a brain drain. The slow growth in skilled graduates with formal AI education must be addressed.
- Focus on ethics in AI. While technical standards are designed to address concerns about using technology and managing associated risks, they do not address, by design, the ethical aspects of developing AI technologies, which need to be considered when designing AI strategies.

Collaboration at the international level is picking up. The European Commission and its 28 state members, in addition to Norway, signed the “Declaration of Cooperation on AI” in 2018 (Stix 2020). It aims to define and implement a comprehensive and integrated approach and to review and modernize individual state

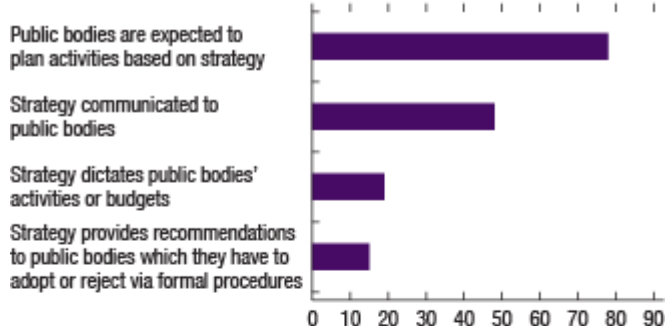
Annex Figure 3.2. Key Features of National Artificial Intelligence Strategies
(Number of policy initiatives embedded in 60 national AI strategies)

AI strategies’ priorities focus on developing an enabling environment, including, in particular, the governance of AI development, digital infrastructure, upgrading skills, and promoting innovative businesses and sectors.

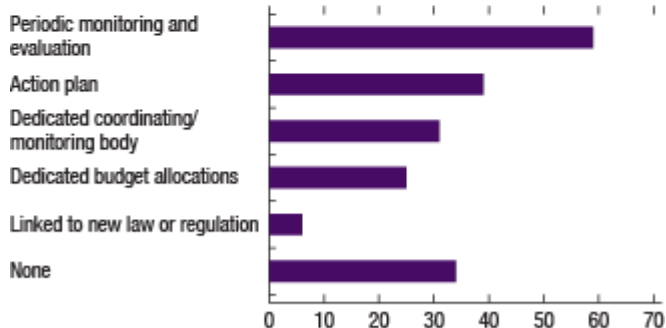
An important aim is to build up local skills and promote innovative businesses and sectors.



Most strategies involve actions and coordination across multiple public bodies.



Progress monitoring is embedded in most strategies and is often performed by a dedicated body.



Source: IMF staff, based on information provided by OECD (2021).
 Note: AI = artificial intelligence.

¹ Of the 60 countries and regions that have developed AI strategies, 11 (Argentina, Brazil, Colombia, Estonia, Germany, India, Israel, Japan, Lithuania, Turkey, and the United States) have included initiatives that aim to address the COVID-19 pandemic (OECD 2021).

policies “to ensure that opportunities arising from AI are seized and the emerging challenges addressed.” The Chief Executive Board within the United Nations has launched an initiative with a concrete roadmap to build capacity in harnessing the potential of AI and raise awareness regarding its risks (UN 2019). The Group of Twenty (G20) Digital Economy Task Force has taken the lead in advancing the G20 AI principles. The main objective of those principles is to foster public trust and confidence in AI technologies and realize their potential through a human-centered approach (G20 2019).

References

- Alonso, C., A. Berg, S. Kothari, C. Papageorgiou, and S. Rehman. 2020. "Will the AI Revolution Cause a Great Divergence?" IMF Working Paper 20/184, International Monetary Fund, Washington, DC.
- Arner, D., J. Barberis, and R. Buckley. 2017. "FinTech, RegTech, and the Reconceptualization of Financial Regulation." *Northwestern Journal of International Law & Business* 37 (3). <https://scholarlycommons.law.northwestern.edu/njilb/vol37/iss3/2/>.
- Bank of England (BoE). 2020. "The Impact of COVID on Machine Learning and Data Science in UK Banking." *Quarterly Bulletin 2020 Q4*. <https://www.bankofengland.co.uk/quarterly-bulletin/2020/2020-q4/the-impact-of-covid-on-machine-learning-and-data-science-in-uk-banking>.
- Bazarbash, M. 2019. "FinTech in Financial Inclusion: Machine Learning Applications in Assessing Credit Risk." IMF Working Paper 19/109, International Monetary Fund, Washington, DC.
- Bolhuis, M. A., and B. Rayner. 2020. "Deus ex Machina? A Framework for Macro Forecasting with Machine Learning." IMF Working Paper 20/45, International Monetary Fund, Washington, DC.
- Chandola, V., A. Banerjee, and V. Kumar. 2009. "Anomaly Detection: A Survey." *ACM Computing Surveys* 41 (3). <https://doi.org/10.1145/1541880.1541882>.
- Cloudera Fast Forward. 2020. "Interpretability." Cloudera Fast Forward Labs Research, Santa Clara, CA. <https://ff06-2020.fastforwardlabs.com/ff06-2020-interpretability.pdf>.
- Corbett-Davies, S., E. Pierson, A. Feller, S. Goel, and A. Huq. 2017. "Algorithmic Decision Making and the Cost of Fairness." In *Proceedings of the 23rd International Conference on Knowledge Discovery and Data Mining*, 797–806.
- Comiter, M. 2019. "Attacking Artificial Intelligence: AI's Security Vulnerability and What Policymakers Can Do About It." Belfer Center for Science and International Affairs, Harvard Kennedy School. <https://www.belfercenter.org/publication/AttackingAI>.
- Danielsson, J., R. Macrae, and A. Uthemann. 2020. "Artificial Intelligence As a Central Banker." VOX CEPR Policy Portal, March 6. <https://voxeu.org/article/artificial-intelligence-central-banker>.
- De Nederlandsche Bank (DNB). 2019. "General Principles for Use of Artificial Intelligence in Finance." DNB, Amsterdam.
- di Castri, S., S. Hohl, A. Kulenkampff, and J. Prenio. 2019. "The Suptech Generations." FSI Insights on Policy Implementation, Bank for International Settlements, Basel. <https://www.bis.org/fsi/publ/insights19.pdf>.
- Digalaki, E. 2021. "The Impact of Artificial Intelligence in the Banking Sector & How AI is Being Used in 2021." *Insider*, January 13. <https://www.businessinsider.com/ai-in-banking-report>.
- Doerr, S., L. Gambacorta, and J. M. Serena. 2021. "Big Data and Machine Learning in Central Banking." BIS Working Paper 930, Bank for International Settlements, Basel.
- European Central Bank (ECB). 2019. "Bringing Artificial Intelligence to Banking Supervision." https://www.bankingsupervision.europa.eu/press/publications/newsletter/2019/html/ssm.nl191113_4.en.html.
- Friedman, B., and H. Nissenbaum. 1996. "Bias in Computer Systems." *ACM Transactions on Information Systems* 14 (3): 330–47.
- Financial Stability Board (FSB). 2020. "The Use of Supervisory and Regulatory Technology by Authorities and Regulated Institutions: Market Developments and Financial Stability Implications." Financial Stability Board, Basel, Switzerland. <https://www.fsb.org/wp-content/uploads/P091020.pdf>.
- FinCoNet. 2020. "SupTech Tools for Market Conduct Supervisors." Paris.

- Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther. 2020. “Predictably Unequal? The Effects of Machine Learning on Credit Markets.” Mimeo. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3072038.
- Gates, S. W., V. G. Perry, and P. M. Zorn. 2002. “Automated Underwriting in Mortgage Lending: Good News for the Underserved?”, *Housing Policy Debate*, 13 (2): 369–91.
- Gambacorta, L., Y. Huang, H. Qiu, and J. Wang. 2019. “How Do Machine Learning and Nontraditional Data Affect Credit Scoring? New Evidence from a Chinese Fintech Firm.” BIS Working Paper 834, Bank for International Settlements, Basel.
- Gensler, G., and L. Bailey. 2020. “Deep Learning and Financial Stability.” MIT Working Paper, Massachusetts Institute of Technology, Boston, MA. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3723132.
- Goodfellow, I., Y. Bengio, and A. Courville. 2016. *Deep Learning*. Cambridge, MA: MIT Press.
- Goodfellow, I. J., J. Shlens, and C. Szegedy. 2014. “Explaining and Harnessing Adversarial Examples.” Cornell University. <https://arxiv.org/abs/1412.6572>.
- Goodman, B., and S. Flaxman. 2016. “European Union Regulations on Algorithmic Decision Making and a ‘Right to Explanation’.” <https://arxiv.org/pdf/1606.08813.pdf>.
- Google and International Finance Corporation. 2020. “e-Conomy Africa 2020.” <https://www.ifc.org/wps/wcm/connect/e358c23f-afe3-49c5-a509-034257688580/e-Conomy-Africa-2020.pdf?MOD=AJPERES&CVID=nmuGYF2>.
- Group of Twenty (G20). 2019. “G20 Ministerial Statement on Trade and Digital Economy.” <https://www.mofa.go.jp/files/000486596.pdf>.
- Guidotti, R., A. Monreale, S. Ruggieri, F. Turini, D. Pedreschi, and F. Giannotti. 2019. “A Survey of Methods for Explaining Black Box Models.” *ACM Computing Surveys* 51 (5): 1–42.
- Haksar, V., Y. Carrière-Swallow, E. Islam, A. Giddings, K. Kao, E. Kopp, and G. Quiros. 2021. “Towards A Global Approach to Data in the Digital Age.” IMF Staff Discussion Note, International Monetary Fund, Washington, DC.
- Hao, K. 2019. “This is How AI Bias Really Happens—And Why It’s So Hard to Fix.” *MIT Technology Review*, February 4, Massachusetts Institute of Technology, Cambridge, MA.
- Harker, P. 2020. “The Economics of Artificial Intelligence and Machine Learning.” Speech at the Official Monetary and Financial Institutions Forum, September 29, Philadelphia, PA.
- Hinton, G., Y. LeCun, and Y. Bengio. 2015. “Deep Learning.” *Nature* 521 (7553): 436–44.
- Hirshberg, J., and C. D. Manning. 2015. “Advances in Natural Language Processing.” *Science* 349 (6245): 261–66.
- Hong Kong Monetary Authority (HKMA). 2020. “Reshaping Banking with Artificial Intelligence.” https://www.hkma.gov.hk/media/eng/doc/key-functions/financial-infrastructure/Whitepaper_on_AI.pdf.
- Institute of International Finance (IIF). 2017. “Deploying Regtech against Financial Crime.” Report of the Regtech Working Group, IIF, Washington, DC. https://www.iif.com/portals/0/Files/private/32370132_aml_final_id.pdf.
- International Monetary Fund (IMF). 2020. *Regional Economic Outlook for Sub-Saharan Africa*, April, Washington, DC.
- [Khan, A. 2018. “A Behavioral Approach to Financial Supervision, Regulation, and Central Banking.” IMF Working Paper 18/178, International Monetary Fund, Washington, DC.](https://www.imf.org/en/Publications/Working-Papers/2018/178)
- Khan, A., and M. Malaika. 2021. “Central Bank Risk Management, Fintech, and Cybersecurity.” IMF Working Paper 2021/105, International Monetary Fund, Washington DC.
- Khandani, A., K. Adlar, and A. Lo. 2010. “Consumer Credit-Risk Models via Machine Learning Algorithms.” *Journal of Banking & Finance* 34 (11): 2767–787.

- Kleinberg, J., H. Lakkaraju, J. Leskovec, J., Ludwig, J., and S. Mullainathan. 2018a. "Human Decisions and Machine Predictions." *The Quarterly Journal of Economics* 133 (1): 237–93.
- Kleinberg, J., J. Ludwig, S. Mullainathan, A. Rambachan. 2018b. "Algorithmic Fairness." *AEA Papers and Proceedings* 108: 22–27.
- Kleinberg, J., J. Ludwig, S. Mullainathan, C. Sunstein. 2019. "Discrimination in the Age of Algorithms." *Journal of Legal Analysis* 10: 113–74.
- Lehmann, E. L. 1951. "A General Concept of Unbiasedness." *The Annals of Mathematical Statistics* 22 (4): 587–92.
- Lessmann, S., B. Baesens, H. V. Seow, and L. C. Thomas. 2015. "Benchmarking State-of-the-Art Classification Algorithms for Credit Scoring: An Update of Research." *European Journal of Operational Research* 247 (1): 124–36.
- Liu, K., B. Dolan-Gavitt, and S. Garg. 2018. "Fine-Pruning: Defending against Backdooring Attacks on Deep Neural Networks." In *Research in Attacks, Intrusions, and Defenses*, edited by M. Bailey, T. Holz, M. Stamatogiannakis, and S. Ioannidis, 273–94. Cham, Switzerland: Springer.
- Mayson, S. G. 2019. "Bias In, Bias Out." *The Yale Law Journal* 128 (8): 2218–300.
- McKinsey. 2020a. "AI-Bank of the Future: Can Banks Meet the AI Challenge?" <https://www.mckinsey.com/industries/financial-services/our-insights/ai-bank-of-the-future-can-banks-meet-the-ai-challenge#>.
- McKinsey. 2020b. "The State of AI in 2020: Survey." <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/global-survey-the-state-of-ai-in-2020>.
- Miller, A. P. 2018. "Want Less-Biased Decisions? Use Algorithms." *Harvard Business Review*, July 26. <https://hbr.org/2018/07/want-less-biased-decisions-use-algorithms>.
- Molnar, C. 2021. *Interpretable Machine Learning: A Guide for Making Black Box Models Explainable*. <https://christophm.github.io/interpretable-ml-book/>.
- Monetary Authority of Singapore (MAS). 2018. "Principles to Promote Fairness, Ethics, Accountability, and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector."
- Murphy, K. P. 2012. *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: MIT Press.
- Organisation for Economic Co-operation and Development (OECD). 2019. "Recommendation of the Council on Artificial Intelligence." OECD Legal Instruments, OECD, Paris. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.
- Organisation for Economic Co-operation and Development (OECD). 2021. "STIP Compass Database." OECD, Paris. <https://stip.oecd.org/stip.htm>.
- Percha, B., and R. B. Altman. 2015. "Learning the Structure of Biomedical Relationships from Unstructured Text." *PLoS Computational Biology* 11 (7): e1004216.
- Petropoulos, A., V. Siakoulis, E. Stavroulakis, and A. Klamargias. 2019. "A Robust Machine Learning Approach for Credit Risk Analysis of Large Loan Level Datasets Using Deep Learning and Extreme Gradient Boosting." *IFC Bulletin* 49: 1486–506, Bank for International Settlements, Basel.
- Plous, S. 2002. "The Psychology of Prejudice, Stereotyping, and Discrimination: An Overview." In *Understanding Prejudice and Discrimination*, edited by S. Plous. New York: McGraw-Hill.
- Ponti, E.M., H. O'Horan, Y. Berzak, I. Vulić, R. Reichart, T. Poibeau, E. Shutova, and A. Korhonen. 2019. "Modeling Language Variation and Universals: A Survey on Typological Linguistics for Natural Language Processing." *Computational Linguistics* 45 (3): 559–601.

- Ribeiro, M. T., S. Singh, and C. Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier." Cornell University. <https://arxiv.org/abs/1602.04938>.
- Sahay, R., U. E. von Allmen, A. Lahreche, P. Khera, S. Ogawa, M. Bazarbash, and K. Beaton. 2020. "The Promise of Fintech: Financial Inclusion in the Post-COVID-19 Era." Departmental Paper 20/09, International Monetary Fund, Washington, DC.
- Schizas, E., G. McKain, B. Zhang, A. Ganbold, P. Kumar, H. Hussain, K. J. Garvey, et al. 2019. "The Global Regtech Industry: Benchmark Report." Cambridge Centre for Alternative Finance. University of Cambridge, UK. <https://www.jbs.cam.ac.uk/wp-content/uploads/2020/08/2019-12-ccaf-global-regtech-benchmarking-report.pdf>.
- Shapley, Lloyd S. 1953. "A Value for N-Person Games." *Contributions to the Theory of Games* 2 (28): 307–17.
- Silberg, J., and J. Manyika. 2019. "Notes from the AI Frontier: Tackling Bias in AI (and in Humans)." McKinsey Global Institute.
- Stanford University. 2019. "The AI Index 2019 Annual Report." AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA. https://hai.stanford.edu/sites/default/files/ai_index_2019_report.pdf.
- Stix, C. 2020. "A Survey of the European Union's Artificial Intelligence Ecosystem." JRC Science Hub Communities, European Commission, Brussels. https://ec.europa.eu/jrc/communities/sites/default/files/ff3afe_1513c6bf2d81400eac182642105d4d6f.pdf.
- Sy, A., R. Maino, A. Massara, H. Perez-Saiz, and P. Sharma. 2019. "Fintech in Sub-Saharan African Countries: A Game Changer?" Departmental Paper 19/04, International Monetary Fund, Washington, DC.
- Szegedy, C., W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. 2014. "Intriguing Properties of Neural Networks." Cornell University. <https://arxiv.org/abs/1312.6199>.
- Triepels, R., H. Daniels, and R. Heijmans. 2018. "Detection and Explanation of Anomalous Payment Behavior in Real Time Gross Settlement Systems." In *Enterprise Information Systems*, edited by S. Hammoudi, M. Śmiałek, O. Camp, and J. Filipe, 145–61. Cham: Springer Verlag.
- United Nations (UN). 2019. A United Nations System-Wide Strategic Approach and Road Map for Supporting Capacity Development on Artificial Intelligence, CEB/2019/1/Add.3 (17 June 2019). [rom undocs.org/en/CEB/2019/1/Add.3](https://undocs.org/en/CEB/2019/1/Add.3).
- United Nations Educational, Scientific and Cultural Organization (UNESCO). 2021. Intergovernmental Meeting of Experts (Category II) related to a Draft Recommendation on the Ethics of Artificial Intelligence. <https://unesdoc.unesco.org/ark:/48223/pf0000373434>.
- Wang, J. C., and C. B. Perkins. 2019. "How Magic a Bullet Is Machine Learning for Credit Analysis? An Exploration with FinTech Lending Data." Federal Reserve of Boston Working Paper 19–16.
- Wang, T. 2016 "The Human Insights Missing from Big Data." https://www.ted.com/talks/tricia_wang_the_human_insights_missing_from_big_data.
- World Economic Forum (WEF). 2018. "The New Physics of Financial Services: Understanding How Artificial Intelligence is Transforming the Financial Ecosystem." http://www3.weforum.org/docs/WEF_New_Physics_of_Financial_Services.pdf.
- World Economic Forum (WEF). 2020. "Transforming Paradigms: A Global AI in Financial Services Survey." World Economic Forum and Cambridge Centre for Alternative Finance. http://www3.weforum.org/docs/WEF_AI_in_Financial_Services_Survey.pdf.



PUBLICATIONS

Powering the Digital Economy

Departmental Paper No. [DP/2021Y024

