

2/3

AI Governance Alliance
Briefing Paper Series 2024

Unlocking Value from Generative AI: Guidance for Responsible Transformation

IN COLLABORATION
WITH IBM CONSULTING

Contents

Executive summary	3
Introduction	4
1 New opportunities with generative AI	5
2 Assessing use cases for adoption	7
2.1 Evaluation gate: business impact	8
2.2 Evaluation gate: operational readiness	8
2.3 Evaluation gate: investment strategy	9
3 Responsible transformation	10
3.1 The case for responsible transformation	10
3.2 Addressing accountability: defined governance for immediate and downstream outcomes	11
3.3 Addressing trust: enabling transparency through communication	11
3.4 Addressing challenges to scale: diverse and agile operations structures	12
3.5 Addressing human impact: value-based change management	12
Conclusion	12
Contributors	13
Endnotes	17

Disclaimer

This document is published by the World Economic Forum as a contribution to a project, insight area or interaction. The findings, interpretations and conclusions expressed herein are a result of a collaborative process facilitated and endorsed by the World Economic Forum but whose results do not necessarily represent the views of the World Economic Forum, nor the entirety of its Members, Partners or other stakeholders.

© 2024 World Economic Forum. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, including photocopying and recording, or by any information storage and retrieval system.

Executive summary

Organizations should emphasize responsible transformation with generative AI to build a sustainable future.

Generative AI entered the popular domain with the launch of OpenAI's ChatGPT in November 2022, igniting global fascination surrounding its capabilities and potential for transformative impact. As generative AI's technical maturity accelerates, its adoption by organizations seeking to capitalize on its potential is maturing at pace while also swiftly disrupting business and society and forcing leaders to rethink their strategies in real time. This paper addresses the impact of generative AI on industry and introduces best practices for responsible transformation.

Leaders have realized new generative AI opportunities for their organizations, from streamlining enterprise processes to supporting artists in reimagining furniture design or even aiding nations in addressing global climate challenges. From the public to the private sector, organizations are witnessing generative AI's ability to enhance enterprise productivity, create net new products or services, and redefine industries and societies. In adopting generative AI, leaders report a shift towards a use-case-based approach, focusing on evaluating and prioritizing use cases and structures that enable the successful deployment of generative AI technologies and compound value generation.

Organizations should evaluate potential use cases across the following domains: business impact, organisational readiness and investment strategy.

- Strategic alignment with the organization's goals, revenue and cost implications, and impact on resources are key factors when leaders prioritize use cases based on their potential for **business impact**.
- The requisite technical talent and infrastructure, the ability to track data and model lineage, and the governance structure to manage risk are

considerations when leaders evaluate use cases against their **operational readiness**.

- Balancing upfront development cost with reusability potential, projected time to value and an increasingly complex regulatory environment are criteria when leaders select use cases in alignment with an organization's **investment strategy**.

Following use case selection, organizations weigh benefits against downstream impacts such as impact to the workforce, sustainability or inherent technology risk such as hallucinations. A multistakeholder approach helps leaders to mitigate risk and scale responsibly.

- Multistakeholder governance with distributed ownership is central to **addressing accountability**.
- Communications teams that shape a cohesive narrative are essential to **addressing trust** through transparency.
- Operational structures that roadmap and cascade use cases to extract, realize, replicate and amplify value across the entire organization are key to **addressing challenges to scale**.
- Value-based change management is critical to **addressing human impact** and ensuring the workforce remains engaged and upskilled.

The findings in this briefing paper provide leaders with insights on how to realise the benefits of generative AI while mitigating its downstream impacts. Future publications will build on these recommendations for responsible transformation as generative AI becomes increasingly able to mimic human skills and reasoning, and technology advances in pursuit of artificial general intelligence.

Introduction

Generative AI raises new questions about responsible transformation for industry executives, government leaders and academia.

Generative artificial intelligence (AI) has captured global imagination with its human-like capabilities and has shown the potential to elevate creativity, amplify productivity, reshape industries and enhance the human experience. As a result, cross-sector executives, government leaders and academia are considering the potential impact of this technology as they weigh answers to critical questions:

- Where are the growing opportunities and novel application areas to drive sustainable economic growth?

- What are the new challenges and downstream impacts?
- What are the best practices for scaling responsibly and bringing about exponential transformation?

Finally, as the curiosity to replicate or even exceed human intelligence grows in the future, what does this mean for organizations seeking to capitalize on the opportunities offered by this technology?



1

New opportunities with generative AI

Generative AI creates new opportunities but requires a distinctive approach to value generation focused on use cases and experimentation.

Generative AI is expected to unlock opportunities that will significantly impact the global economy. Organizations are already using generative AI to enhance existing products, services, operations and provide hyper-personalized customer experiences. While most use cases focus on boosting human capabilities, some have the potential to radically accelerate benefits to humanity. For example, novel synthetic protein structures generated to help fix DNA errors can significantly accelerate the creation

of new cancer therapies.¹ Generative AI is also used to orchestrate deep synthesis of numerous data catalogues to enable work to protect the oceans.² These bolder bets have the potential to reshape not just entire industries but economies and societies at large. In general, use cases can be considered under different categories that include enhancing enterprise productivity, creating new products or services and, eventually, redefining industries and societies.

TABLE 1 Snapshot of sample generative AI case studies in the market

Category	Company	Challenge	Action	Impact
Enhancing enterprise productivity	Brex: automating corporate card expenses ³	Support corporate card customers to categorize transactions and add notes to meet company policies and Internal Revenue Service (IRS) compliance.	Brex, with OpenAI and Scale, used generative AI to create the Brex Assistant to streamline expense reporting, automatically classify expenses and create IRS-compliant notes.	Brex Assistant fully handles 51% of card swipes, saving time and improving expense accuracy and compliance. It generated over 1.4 million receipts and 1 million receipt memos.
Enhancing enterprise productivity	IKEA: reimagining furniture design ⁴	Seek creative solutions to aid furniture designers in crafting new designs inspired by their iconic past.	IKEA and SPACE10 used generative AI to explore furniture design concepts, training a model on 1970s and 1980s catalogues for students to create future-focused designs inspired by the past.	Furniture designers collaborate with AI, expanding design possibilities and speeding up cycles.
Enhancing enterprise productivity and net-new product or service	Google: streamlining software prototyping ⁵	Reduce software development cycles internally and simplify access to generative AI models.	Google created Google AI Studio, a generative AI tool to simplify software prototyping and democratize access to their foundation models, which were first used internally.	Increased proactive UX and product prototyping, provided an efficient UI for easy model prompting and was later launched as a new product in 179 countries and territories.
Net-new product or service	Synthia and PepsiCo: reinventing the football fan experience ⁶	Connect brand and performance marketing efforts into one seamless experience.	Fans could generate and share personalized videos using Lionel Messi's AI avatar in eight languages, bypassing traditional production limits.	Seven million videos were regenerated, attracting over 38 million website visits in 24 hours.

TABLE 1 | Snapshot of sample generative AI case studies in the market (continued)

Category	Company	Challenge	Action	Impact
Redefining industries and societies	Insilico Medicine: accelerating drug discovery ^{7,8}	Discover and develop new treatments for serious diseases more quickly and cheaply compared to traditional processes.	Generative AI was used during the preclinical drug discovery process to identify a novel drug candidate for idiopathic pulmonary fibrosis.	A preclinical drug candidate was discovered in less than 18 months and at one-tenth of the cost of a conventional programme. The drug candidate has now entered phase two trials.
Redefining industries and societies	NASA and IBM: unique global planning for climate phenomena and sustainability ⁹	Build a unique foundation model to generate insights from over 250 terabytes (TBs) of mission satellite imagery.	NASA and IBM created the first open-source geospatial foundation model, available via Hugging Face, using NASA data to enhance and democratize global environmental research and planning.	The model is estimated to increase geospatial analysis speed by four times with 50% less labelled data; used to solve global climate challenges, including reforestation in Kenya and other development efforts in the Global South.

“ Organizations are shifting towards smaller, use-case based approaches that emphasize ideation and experimentation.

The speed of adoption and implementation of generative AI is unparalleled to any other technological advancement. The technology is no longer dependent on the manual labelling of significant amounts of data – often the most time-consuming and costly part of traditional AI workflows.

Across the board, leaders report a new approach to generative AI opportunities that extends beyond rapid proofs of concept (POCs) based on large models. Instead, organizations are shifting towards smaller, use-case based approaches that emphasize ideation and experimentation. They are involving the workforce in the use case discovery and ideation process. Smaller use cases with low complexity are often applied first, allowing

leaders to find value while minimizing downstream implications. In either case, leaders start with diverse POCs, which are scaled across the enterprise once value is proven.

In many instances, generative AI experiments may yield unexpected learnings about where value, and often also cost and challenges, truly lie. Organizations may realize the compound benefits of generative AI when implementing it in tandem with technologies such as causal AI models¹⁰ to increase explainability, advances in quantum technologies to accelerate the generative AI life cycle, or 5G to increase reach. These compounding benefits will help organizations to prioritize use cases for adoption.

2


Assessing use cases for adoption

Generative AI use cases may be assessed by business impact, organizational readiness and investment strategy prior to adoption.

As organizations consider generative AI, they must assess all factors involved to move a use case from concept to implementation. Leaders need to ensure that each use case benefits the organization, its customers, its workforce and/or society. While evaluation criteria can differ between organizations,

the following gates comprise the most common approaches adopted by industry leaders to evaluate the viability and value-generation potential of use cases. The order is not sequential and can differ depending on each organization and use case.

FIGURE 1 Funnelling use cases through evaluation gates



2.1 Evaluation gate: business impact

Leaders evaluate the use case's value alignment with the organization's strategic objectives and its stakeholder responsibility. After alignment on the outcomes and generative AI as the best technology to address a specific use case, the impact of each use case on an organization can be categorized as follows:

1. **Scaling human capability** by enhancing productivity and existing human skills (e.g. near instant new content generation for rapid idea iteration; creation of multiple versions of an advertising campaign).
2. **Raising the floor** by increasing accessibility to technologies and capabilities previously requiring specific resources, skills and expertise (e.g. giving everyone the ability to code).
3. **Raising the ceiling** by solving problems thus far unsolvable by humans (e.g. generating new

molecular structures, which could aid the creation of novel and more effective therapeutic agents.¹¹

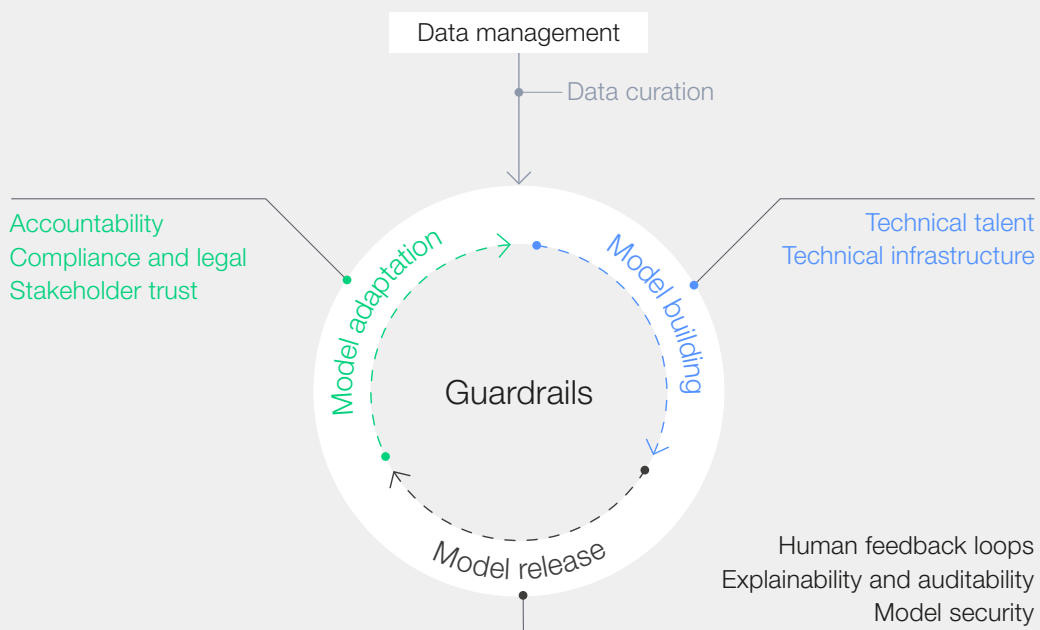
Generative AI opportunities have created strong competitive pressures and inaction can come with significant opportunity costs.¹² In industries such as marketing or consumer goods, understanding the criticality of time to market and improved experience for users, helps leaders prioritise use cases and resource allocation. Reputation is another important consideration – will the use case enhance the organization's brand as a pioneer of innovation? Enabling the workforce to access generative AI tools can be an important factor for talent attraction and retention. When generative AI performs administrative tasks that previously required significant time and effort, the workforce can repurpose their time from rote activities to those that allow them to explore their creativity and hone their unique skillset.

2.2 Evaluation gate: operational readiness

Responsible adoption of generative AI requires operational readiness for technological dependencies and outcomes. Before organizations expose generative AI to their data, data curation is essential to ensure it is accurate, secure, representative and relevant. In developing or implementing generative AI technologies, organizations must consider if they have the right technical talent and infrastructure, such as appropriate models and necessary

computing power. In deploying generative AI technologies, organizations should ensure human feedback loops are in place to mitigate risks by ensuring user feedback is elicited, standardized and incorporated into the continuous fine-tuning of the model. Additionally, organizations require the ability to track model lineage and data sources that inform model outputs, as well as vet models and systems for cybersecurity robustness.

FIGURE 2 Operational readiness considerations (non-exhaustive) across the model life cycle



Organizations will be held responsible for the outcomes of their AI technology and must, therefore, ensure compliance with the global complexity of regulation and policies as cited in *Generative AI Governance: Shaping the Collective Global Future*.¹³ This will require new skills and roles for accountability, compliance and legal responsibilities as a multistakeholder approach. Generative AI's

evolutionary nature and its inherent potential for downstream implications create a greater need to continually evaluate even if the necessary guardrails are in place. Finally, organizations need a plan to enhance stakeholder trust with a technology that can elicit great scepticism to ensure their workforce, customers and other critical parties responsibly adopt generative AI.

2.3 Evaluation gate: investment strategy

While investment considerations are important to any organizational decision-making, they are particularly significant for generative AI opportunities. Use cases often require a higher upfront investment, the regulatory environment is becoming increasingly complex and the technology is evolving at a rapid pace.

When prioritizing use cases, leaders must consider if each merits the use of models adopted from open-source communities, acquired from other third parties or developed in-house. Model selection must account for alignment with the use case, speed to market, requisite resource investments, including capital and talent, licensing and acceptable use policies, risk exposure and competitive differentiation offered by each option.

Leaders evaluate the reusability potential of a use case across the organization, as it can offset development costs and curtail sustainability

footprints. Additionally, they evaluate whether the use case can operate viably within the current regulatory environment and whether the organization can monitor compliance to minimize legal risk. This can require significant investment of capital and human resources, such as developers, lawyers, senior leadership and ethics boards.

Talent availability is central to an organization's investment strategy as well. Total investment may include upskilling, re-skilling or hiring additional employees with appropriate generative AI skills, such as content creation, model development or model tuning.

Following the evaluation of use cases by business impact, organizational readiness and investment strategy, the next step is to implement and scale selected use cases. How can they maximize opportunities while mitigating risks to ensure a responsible and successful transformation?



Responsible transformation

A multistakeholder approach creates value while balancing challenges of trust, accountability, scale and the workforce.

3.1 The case for responsible transformation

As *The Presidio Recommendations on Responsible Generative AI* detail, responsible transformation requires specific considerations for generative AI's unique capabilities, along with multistakeholder collaboration and proper steering during the transformation journey. Global generative AI regulations and standards (NIST et al.) are changing, and so the current need for self-governance is shared by organizations and leaders. There is also a need to ensure that the technology is accessible to all. Organizations are committed to aligning with global environmental and sustainability goals, pledging to adopt AI in a responsible and accessible manner.

The lack of responsibility in an organization's transformation can have many negative consequences, which are multi-fold and compounded for a technology as revolutionary as generative AI. From perpetuating biases, introducing security vulnerabilities and spreading misinformation – causing severe reputational damage – irresponsible generative AI applications and practices not only threaten the organization itself but can also negatively impact society at speed and scale.

Generative AI comes with several downstream implications associated with more traditional forms of AI, together with amplified and new ones. The following are most often noted for their potential impact, with a further list to be explored in future work.

1. Workforce and talent impact

While AI is commonly used to automate tasks, the scale at which generative AI can accomplish this amplifies its impact on the workforce. The potential risk of job displacement presents significant challenges for society that can exacerbate inequality. Research indicates that generative AI's automation capabilities provide the greatest exposure for clerical jobs, which have traditionally been held by women. In some cases, particularly in developing countries, these types of jobs may cease to exist, removing an avenue that has historically served as an entry for women into the labour market.¹⁴ Additionally, generative AI's novel capability to create,

generate and simulate human-like interactions may now overlap with tasks in creative industries, and its ability to rapidly learn domain expertise may influence the roles of knowledge workers.

Skills and workloads are changing, and organizational structures need to evolve at pace.¹⁵ Generative AI is profoundly changing the way employees view their jobs and the value work brings. Nevertheless, the technology presents a unique opportunity for organizations to re-evaluate their working practices and skills: to inspire, incentivize, motivate, upskill and reskill workers, while evaluating the agility of their own organizational structures.

2. Hallucination impact

Generative AI introduces the risk of hallucinations, which can propagate misinformation, leading to confusion, mistrust and even potential harm. Equally, hallucinations are a corollary of generative AI's capability to create net-new content, which is central to its power to accelerate creativity. Organizations need to understand whether the benefit of content creation outweighs the risk of hallucination for each use case.

Hallucinations are particularly concerning when generative AI outputs appear authoritative but are factually inaccurate, especially when used to influence decision-making that may impact global communities in areas such as health, politics and science. Organizations that rely on digital content production or customer engagement face challenges as brand reputation and customer trust could be damaged. Guardrails from *Presidio AI Framework: Towards Safe Generative AI Models* need to be considered and embedded in the process.¹⁶

3. Sustainability impact

Training and fine-tuning generative AI models demand very high energy consumption.¹⁷ Growing global efforts to offset or mitigate their sustainability footprint are ongoing, such as advancements in model, runtime and hardware

optimization, as well as improved education on model choices. Algorithmic approaches like federated computing can further minimize the energy consumption of data collection and processing. Organizations also consider their choices in data needs as a growing move towards smaller, more targeted, and more energy-efficient models underlines.

In addition to ensuring generative AI models are more sustainable, the technology itself can be used to improve sustainability, for example, through use cases focussed on energy modelling and supply chain optimization.¹⁸

As the risks associated with generative AI amplify and expand, traditional organizational structures need to pivot with agility. Leaders need to ensure cross-functional connectivity from the board level down and across all impacted functions. The following are four interconnected and interdependent functions that support this organizational effort to balance the opportunities and benefits of generative AI with its downstream impacts as organizations implement and scale generative AI applications.

3.2 Addressing accountability: defined governance for immediate and downstream outcomes

“ An AI ethics council modelled on value-based principles is indispensable for any organization.

Multistakeholder governance with distributed ownership is central to responsible transformation in the age of generative AI. This approach is characteristic of industry leaders, with legal, governance, IT, cybersecurity, human resources (HR), as well as environmental and sustainability representatives requiring a seat at the table to ensure responsible transformation across the organization. The positive and negative externalities of generative AI expand the conventional responsibilities in governance towards a more holistic, human-centred and values-driven approach.

An AI ethics council modelled on value-based principles¹⁹ is indispensable for any organization; larger organizations appoint members from their stakeholder and shareholder groups, while smaller organizations may need to rely on a limited committee or an external ethics council. Councils must collaborate with stakeholders on aspects such

as workplace policies, even if they do **not** deploy generative AI, as the workforce is likely already using it at work on personal devices. The council should expand to incorporate a diverse set of members from across the entire organization to ensure the responsible adoption of not just individual use cases but also emerging and intersecting strategies on open technologies, artificial general intelligence (AGI), 5G and quantum technology.

The evolving nature of generative AI requires rigorous self-regulation and internal AI governance leads may serve as the sentinels of the organization. Generative AI supports human-led analysis in regulatory, environmental and sustainability efforts. It assists in algorithm monitoring and policy formulation, but crucially, it requires human oversight to ensure responsible and effective application, addressing potential risks and maintaining quality outcomes.

3.3 Addressing trust: enabling transparency through communication

Generative AI evokes mixed reactions from stakeholders, placing a high demand on communications teams. These teams shape a cohesive narrative to showcase how their organization optimizes transparency, explainability, coherence and trustworthiness on a use case basis. They play a role in educating stakeholders and shareholders on the capabilities and fallibilities of the technology while managing expectations. They can inspire and instruct end-users about the benefits on the horizon, thus building trust and increasing adoption.

External communications need to assuage stakeholders that seek innovation, but not at the cost of ethical behaviour, trust and actions that prove that the organization is committed to the greater good of humanity. Internal accountability and advocacy are needed from top leadership to obtain buy-in from the workforce and establish a culture that benefits from generative AI. Examples of effective trust programmes include taking a prominent ethics stance in policy or the executive community, buddy programmes for all employees seeking (generative) AI immersion and novel career pathways that can lead to increased trust and ownership from the workforce.

3.4 Addressing challenges to scale: diverse and agile operations structures

Initial adoption of generative AI across organizations has focused on targeted, often isolated, use cases. However, as leaders plan their strategic roadmaps, many are challenged with how to scale these use cases across their organizations to realize the compound benefits of generative AI.

Operations teams are the primary implementers of use cases. Data analysts, research and development teams, resource managers, HR executives and business leaders ensure use

cases are roadmapped and cascaded across the organization for maximum benefit. In their initial development, use cases require a diverse operational structure to ensure a multistakeholder approach to extracting, realizing, replicating and amplifying value. However, as use cases become integrated and scale, an interlocking and agile operational structure is needed to understand how compound value can be unlocked, and corollary impacts to other parts of the workforce or other lines of business can be anticipated.

3.5 Addressing human impact: value-based change management

Technologies that develop as rapidly as generative AI require adoption by a workforce that evolves at pace. The implications of generative AI on the workforce are central to business and need to be managed well. The chief human resources officer, the chief information officer, and the chief financial officer teams should come together to support the workforce as needed when implementing and scaling generative AI use cases.

Leaders plan and implement talent transformation while ensuring staff have access to the necessary technological tools and training. This starts with communicating the vision for generative AI pilots that clearly states desired benefits for customers and employees alike, together with emerging professional development pathways for staff. Competencies, capabilities and skills are rapidly evolving as generative AI use cases are implemented across the organization.

Change management responsibilities across the organization are significant. HR professionals engage with the implementation of use cases from the beginning so they can proactively assess the impact on staff and put workforce transformation plans in place. Including employees in idea generation for use cases and encouraging them to own their career paths can increase engagement. Hackathons and company-wide training days are effective in upskilling the workforce while also encouraging experimentation and innovation.

The immense potential of generative AI for benefit as well as for harm requires that all four of these primary functions are dynamic, interlocked and in equilibrium. The effectiveness of this interlock correlates directly with the extent to which an organization scales generative AI applications responsibly.

“Technologies that develop as rapidly as generative AI require adoption by a workforce that evolves at pace.”

Conclusion

New technologies driving productivity have always been positioned as repurposing workers to higher-value work, which has traditionally required human oversight and creativity. However, with generative AI becoming increasingly advanced in its ability to mimic human skills and capabilities, it opens more questions about its impact on the organizations choosing to adopt it. Technological advances towards human reasoning in the pursuit of artificial general intelligence demand ongoing discourse on the responsibility of organizations to their workforce, customers and wider society.

Future work through the World Economic Forum's AI Governance Alliance will build on this foundation and address essential considerations, such as internal metrics for responsibility, understanding organizational barriers to responsible transformation, as well as broader issues such as intellectual property, regulatory alignment and workforce considerations. Generative AI is reimagining the status quo for every organization. Providing a roadmap for organizations that guides them to innovate responsibly is key to adopting and scaling this powerful technology.

Contributors

This paper is a combined effort based on numerous interviews, discussions, workshops and research. The opinions expressed herein do not necessarily reflect the views of the individuals or organizations

involved in the project or listed below. Sincere thanks are extended to those who contributed their insights via interviews and workshops, as well as those not captured below.

World Economic Forum

Hubert Halopé

Lead, Artificial Intelligence and Machine Learning

Devendra Jain

Lead, Artificial Intelligence, Quantum Technologies

Daegan Kingery

Early Careers Programme, AI Governance Alliance

Connie Kuang

Lead, Generative AI & Metaverse Value Creation

Benjamin Larsen

Lead, Artificial Intelligence and Machine Learning

Cathy Li

Head of AI, Data and Metaverse; Deputy Head, Centre for the Fourth Industrial Revolution; Member of the Executive Committee

AI Governance Alliance Project Fellows

Ann-Sophie Blank

Managing Consultant, IBM

Alison Dewhirst

Senior Managing Consultant, IBM

Heather Domin

Executive Fellow, Director of Responsible AI Initiatives, IBM

Sophia Greulich

Senior Consultant, IBM

Michelle Hannah Jung

Senior Managing Consultant, IBM

Jennifer Kirkwood

Executive Fellow, Partner, IBM

Avi Mehra

Associate Partner, IBM

Sandra Misiaszek

Associate Partner, IBM

Acknowledgements

Sincere appreciation is extended to the following working group members, who spent numerous hours providing critical input and feedback to the drafts. Their diverse insights are fundamental to the success of this work.

Martin Adams

Co-Founder, METAPHYSIC

Basma AlBuhairan

Managing Director, Centre for the Fourth Industrial Revolution, Saudi Arabia

Uthman Ali

Senior Product Analyst, AI Ethics SME, BP

Mohamed Alsharid

Chief Digital Officer, Dubai Electricity and Water Authority (DEWA)

Stefan Badža

Director, Team for Special Projects, Office of the Prime Minister of Serbia

Ricardo Baptista Leite

Chief Executive Officer, Health AI, The Global Agency for Responsible AI in Health

Elisabeth Bechtold

Head, AI Governance, Zurich Insurance Group

Sébastien Bey

Senior Vice-President and Global Head of IT at Siemens Smart Infrastructure, Siemens

Lu Bo

Vice-President; Head, Corporate Strategy, Lenovo Group

Ting Cai
Group Senior Managing Executive Officer;
Chief Data Officer, Rakuten Group

Cansu Canca
Director, Responsible AI Practice, Institute for
Experiential AI, Northeastern University

Nadia Carlsten
Vice-President, Product, SandboxAQ

Will Cavendish
Global Digital Services Leader, Arup Group

Rohit Chauhan
Executive Vice President, AI & Security Solutions,
Mastercard International

Adrian Cox
Managing Director, Thematic Strategist,
Deutsche Bank Research, Deutsche Bank

Bhavesh Dayalji
Chief Executive Officer, Kensho Technologies

Evren Dereci
Chief Executive Officer, KocDigital

Dan Diasio
Global Artificial Intelligence Consulting Leader, EY

P. Murali Doraiswamy
Professor of Psychiatry and Medicine,
Duke University School of Medicine

Elena Fersman
Vice-President and Head of Global AI
Accelerator, Ericsson

Ryan Fitzpatrick
Senior Vice-President, Strategy, Vindex

Lucas Glass
Vice-President, Analytics Center of Excellence, IQVIA

Mark Gorenberg
Chair, Massachusetts Institute of Technology (MIT)

Mark Greaves
Executive Director, AI2050, Schmidt Futures

Olaf Groth
Professional Faculty, Haas School of Business,
University of California, Berkeley

Sandeep Grover
Trust & Safety Leadership, TikTok

Sangeeta Gupta
Senior Vice-President, National Association of
Software and Services Companies (NASSCOM)

Bill Higgins
Vice-President, watsonx Platform Engineering
and Open Innovation, IBM

Matissa Hollister
Assistant Professor of Organizational Behaviour,
McGill University

Michael G. Jacobides
Professor of Strategy; Sir Donald Gordon
Professor of Entrepreneurship and Innovation,
London Business School

Fariz Jafarov
Executive Director, Centre for the Fourth Industrial
Revolution, Azerbaijan

Reena Jana
Head, Content & Partnership Enablement,
Responsible Innovation, Google

Jeff Jarvis
Professor, Graduate School of Journalism,
City University of New York

Emilia Javorsky
Director, Futures Program, Future of Life Institute

Siddhartha Jha
AI and Digital Innovation Lead, Botnar Foundation

Shailesh Jindal
Vice-President of Corporate Strategy,
Palo Alto Networks

Athina Kanioura
Executive Vice-President, Chief Strategy
and Transformation Officer, PepsiCo

Vijay Karunamurthy
Head and Vice-President, Engineering, Scale AI

Sean Kask
Chief AI Strategy Officer, SAP

Faisal Kazim
Head, Centre for the Fourth Industrial Revolution,
United Arab Emirates

Rom Kosla
Chief Information Officer, Hewlett Packard Enterprise

Nikhil Krishnan
Chief Technology Officer, Products, C3 AI

Sebastien Lehnerr
Chief Information Officer, SLB

Giovanni Leoni
Head, Business Development and Strategy,
Credo AI

Art Levy
Chief Strategy Officer, Brex

Leland Lockhart
Director, Artificial Intelligence & Machine Learning,
Vista Equity Partners

Harrison Lung

Group Chief Strategy Officer, e&

Manny Maceda

Chief Executive Officer, Bain & Company

Jim Mainard

Chief Technology Officer and Executive Vice-President, Deep Technology, XPRIZE Foundation

Naveen Kumar Malik

Chief of Staff, Office of the Chief Technology Officer, HCL Technologies

Thomas W. Malone

Professor of Management and Director, Center for Collective Intelligence, MIT Sloan School of Management

Darren Martin

Chief Digital Officer, AtkinsRéalis

Francesco Marzoni

Chief Data & Analytics Officer, Ingka Group (IKEA)

Darko Matovski

Chief Executive Officer, causalens

Andrew McMullan

Chief Data and Analytics Office, Commonwealth Bank of Australia

Nicolas Mialhe

Founder and President, The Future Society (TFS)

Steven Mills

Partner and Chief Artificial Intelligence Ethics Officer, Boston Consulting Group

Angela Mondou

President and Chief Executive Officer, TECHNATION

Yao Morin

Chief Technology Officer, JLL

Mashaël Muftah

International and Regional Organizations Adviser, Ministry of Information and Communication Technology (ICT) of Qatar

Abhishek Pandey

Global Head of Services Business Development, GEP

Charna Parkey

Real-Time AI Product and Strategy Leader, DataStax

Cyril Perducat

Senior Vice-President and Chief Technology Officer, Rockwell Automation

Andreas Prösch

Vice-President and Head, Aker AI Unit, Aker ASA

Philippe Rambach

Chief AI Officer, Schneider Electric

Mary Rozenman

Chief Financial Officer and Chief Business Officer, Insitro

Crystal Rugege

Managing Director, Centre for the Fourth Industrial Revolution, Rwanda

Prasad Sankaran

Executive Vice-President, Software and Platform Engineering, Cognizant Technology Solutions US

Isa Scheunpflug

Head, Automation Office, UBS

Mikkel Skovborg

Senior Vice-President, Innovation, Novo Nordisk Foundation

Genevieve Smith

Founding Co-Director, Responsible & Equitable AI Initiative, Berkeley Artificial Intelligence Research Lab (UC Berkeley)

Eric Snowden

Vice-President, Design, Digital Media, Adobe

Jim Stratton

Chief Technology Officer, Workday

Murali Subbarao

Vice-President, Generative AI Solutions, ServiceNow

Norihiro Suzuki

Chairman of the Board, Hitachi Research Institute, Hitachi

Behnam Tabrizi

Co-Director and Teaching Faculty of Executive Program, Stanford University

Amogh Umbarkar

Vice-President, SAP Product Engineering, SAP

Ingrid Verschuren

Executive Vice-President, Data and AI; General Manager, Europe, Middle East and Africa, Dow Jones

Daniel Verten

Strategy Partner, Synthesia

Judy Wade

Managing Director, CPP Investments

Anna Marie Wagner

Senior Vice-President, Head of AI, Ginkgo Bioworks

Min Wang

Chief Technology Officer, Splunk

Amy Webb

Chief Executive Officer, Future Today Institute

Chaoze Wu

Head of R&D Department, Managing Director, China Securities

Joe Xavier
Chief Technology Officer, Grammarly

Alice Xiang
Global Head, AI Ethics, Sony

Zhang Ya-Qin
Chair Professor and Dean, Tsinghua University

Zhang Ying
Professor of Marketing and Behavioral Science, Guanghua School of Management, Peking University

Zhang Yuxin
Chief Technology Officer, Huawei Cloud, Huawei Technologies

Yijie Zeng
Chief Technology Officer, Beijing Langboat Technology

World Economic Forum

John Bradley
Lead, Metaverse Initiative

Karyn Gorman
Communications Lead, Metaverse Initiative

Jenny Joung
Specialist, Artificial Intelligence and Machine Learning

Hannah Rosenfeld
Specialist, Artificial Intelligence and Machine Learning

Supheakmungkol Sarin
Head, Data and Artificial Intelligence Ecosystems

Stephanie Teeuwen
Specialist, Data and AI

Karla Yee Amezaga
Lead, Data Policy and AI

Hesham Zafar
Lead, Digital Trust

IBM

Phaedra Boinodiris
Associate Partner

Frank Madden
Privacy and Regulatory Risk Adviser

Jesús Mantas
Global Managing Director

Christina Montgomery
Chief Privacy & Trust Officer

Catherine Quinlan
Vice-President, AI Ethics

Sencan Sengul
Distinguished Engineer

Jamie VanDodick
Director AI Ethics and Governance

Production

Laurence Denmark
Creative Director, Studio Miko

Sophie Ebbage
Designer, Studio Miko

Martha Howlett
Editor, Studio Miko

Endnotes

1. Gordon, Rachel, "Generative AI imagines new protein structures", *MIT News*, 12 July 2023, <https://news.mit.edu/2023/generative-ai-imagines-new-protein-structures-0712#:~:text=>.
2. "Navigating the Ocean of Data: Harnessing the Power of Knowledge Graphs in Data Catalogs", *HUB Ocean*, n.d., <https://www.hubocean.earth/blog/ocean-knowledge-graph>.
3. "Brex Gives Every Employee an Expense Assistant with AI", *Brex*, September 2023, <https://www.brex.com/journal/press/brex-gives-every-employee-an-expense-assistant-with-ai>.
4. "Exploring the future with vintage designs in AI", *IKEA Newsroom*, 20 April 2023 <https://www.ikea.com/global/en/stories/design/to-nyttillverkad-and-beyond-ikea-space10-and-designers-of-tomorrow-explore-future-with-ai-230420/>.
5. "Make with MakerSuite – Part 1: An Introduction", *Google for Developers*, 26 September 2023, <https://developers.googleblog.com/2023/09/make-with-makersuite-part1-introduction.html>. [Note: Google renamed the product to Google AI Studio on December 6, 2023].
6. "See our personalized Messi video campaign", *Synthesia Newsroom*, 30 October 2023, <https://www.synthesia.io/post/messi>.
7. "New Milestone in AI Drug Discovery: First Generative AI Drug Begins Phase II Trials with Patients", *Insilico Newsroom*, 1 July 2023, https://insilico.com/blog/first_phase2.
8. "Insilico: linking target discovery and generative chemistry AI platforms for a drug discovery breakthrough", *Nature Research Media*, n.d., <https://www.nature.com/articles/d43747-021-00039-5>.
9. "IBM and NASA are building an AI foundation model for weather and climate", *IBM Newsroom*, 30 November 2023, <https://research.ibm.com/blog/weather-climate-foundation-model>.
10. Definition for causal AI taken from: Forney, Andrew, "Casual Inference in AI Education: A Primer", *Journal of Causal Inference*, 2022, https://ftp.cs.ucla.edu/pub/stat_ser/r509.pdf.
11. Centre for Trustworthy Technology, *A New Frontier for Drug Discovery and Development: Artificial Intelligence and Quantum Technology*, n.d., <https://c4tt.org/1155-2/>.
12. "Building a Value-Driving AI Strategy for Your Business", *Gartner*, n.d., <https://www.gartner.com/en/information-technology/topics/ai-strategy-for-business>.
13. World Economic Forum, *Generative AI Governance: Shaping the Collective Global Future*, 2024.
14. International Labour Organization (ILO), *Generative AI and jobs: A global analysis of potential effects on job quantity and quality*, 2023, https://www.ilo.org/wcmsp5/groups/public/---dgreports/---inst/documents/publication/wcms_890761.pdf.
15. World Economic Forum, *The Future of Jobs Report 2023*, 2023, <https://www.weforum.org/publications/the-future-of-jobs-report-2023/>.
16. World Economic Forum, *Presidio AI Framework: Towards Safe Generative AI Models*, 2024.
17. Strubell, Emma, Ananya Ganesh and Andrew McCallum, *Energy and Policy Considerations for Deep Learning in NLP*, Cornell University Department for Computer Science, Computation and Language, 5 June 2019, <https://arxiv.org/abs/1906.02243>.
18. "Generative AI: The Next Frontier in Energy & Utilities and Oil & Gas Innovation", *BirlaSoft Newsroom*, 26 October 2023, <https://www.birlasoft.com/articles/generative-ai-frontier-energy-utilities-oilgas-innovation>.
19. OECD AI Principles overview adopted in May 2019: "OECD AI Principles overview", *Organisation for Economic Co-operation and Development*, n.d., <https://oecd.ai/en/ai-principles>.



COMMITTED TO
IMPROVING THE STATE
OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org