NTT DATA

Future of Generative AI: Navigating Ethical, Regulatory, and Governance Challenges

NTT DATA point of view on Unraveling Ethical Dimensions, Proactive Compliance, and Charting the Course for a Responsible Future.

NTT DATA Point of View

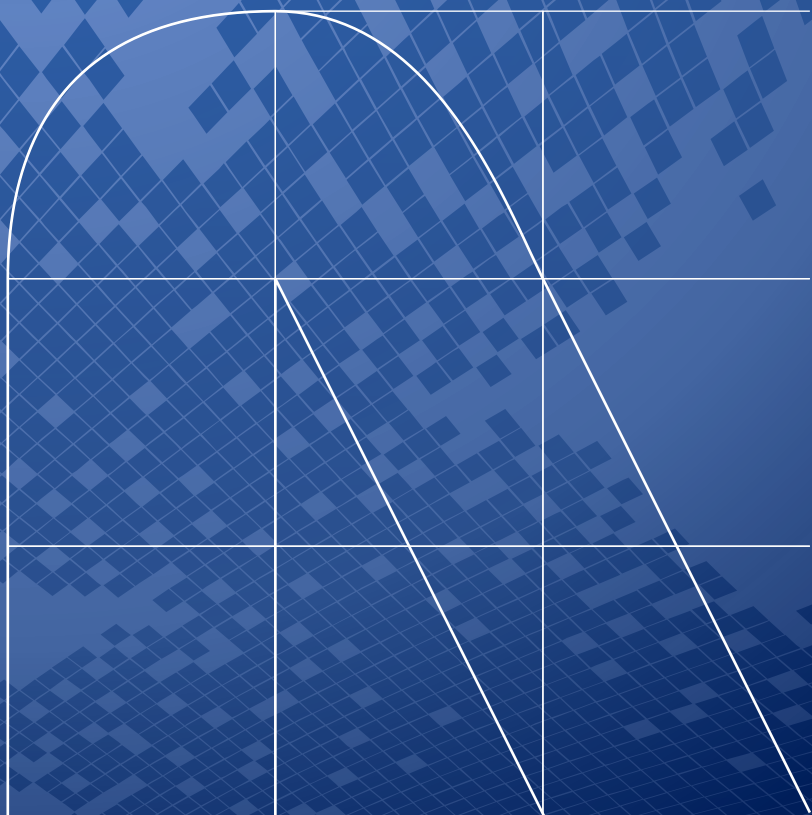

Introduction

Over the last decade, technology has experienced a swift evolution across diverse sectors, from gaming to transportation. The rapid emergence of novel technologies, such as augmented reality in games like Pokémon Go, has outpaced the development of our legal and ethical frameworks. This acceleration has given rise to notable legal challenges, exemplified by the intersection of augmented reality and the resulting legal issues—ranging from privacy concerns related to location-based data collection to disputes over virtual property rights and public safety implications.

Also, with the dawn advent of driverless cars has introduced a new dimension to the relationship between technology and law. As these autonomous vehicles become a reality, navigating the legal landscape involves addressing intricate questions surrounding liability, safety standards, and ethical considerations.

Tortious issues were brought into the regulatory discourse as policymakers sought to establish a liability framework that could account for the unique challenges posed by autonomous vehicles. Questions regarding who bears responsibility in the event of an accident involving a driverless car, whether it's the manufacturer, the software developer, or the vehicle owner, became central to the legal considerations. The principles of negligence and product liability, inherent in tort law, have been instrumental in crafting regulations that establish accountability and ensure a fair distribution of responsibility in this evolving technological landscape.

In the rapidly advancing landscape of Generative AI, upholding ethical standards, implementing effective regulations, and establishing sound governance mechanisms are pivotal. These considerations are not just safeguards but catalysts for responsible innovation, fostering trust and ethical deployment of AI technologies in our ever-evolving society.

Did Generative AI unleash array of ethical concerns for the first time in history?[3]

No, historically the relationship between technology and ethical issues has a long history, with concerns regarding suspicion of immorality arising since the introduction of computing technology in World War II. Also, discussion about the risk of generating electricity in nuclear power stations from the late 1950s after the use of nuclear (atomic) bombs during war were also questioned.

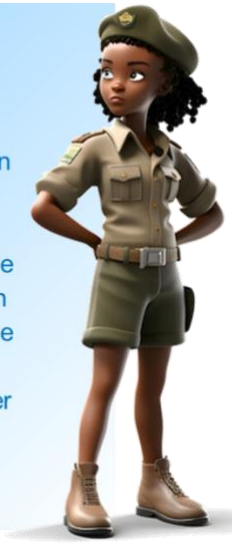

What is Generative AI?[1]

Generative artificial intelligence (AI) describes algorithms (such as ChatGPT) that can be used to create new content, including audio, code, images, text, simulations, and videos. Recent breakthroughs in the field have the potential to drastically change the way we approach content creation.

The Intersection of Law and Ethics in Generative AI

Navigating the terrain of Generative AI is like traversing on extra-terrestrial landscape where unknown and undefined challenges lurk like hidden obstacles. Generative AI is like a painting that subtly reflects the artist's biases, AI training data carries inherent biases that can manifest in the outputs.

Discriminatory outcomes, akin to shadows in this landscape, are cast when these biases take center stage. The synthesis of realistic yet fictional faces serve as a mirror reflecting privacy concerns, much like crafting imaginary characters in a story. As we embark on a detailed exploration of biases, discrimination, privacy, and more, it becomes evident that deciphering this terrain is crucial for charting a responsible course in the development and governance of Generative AI.

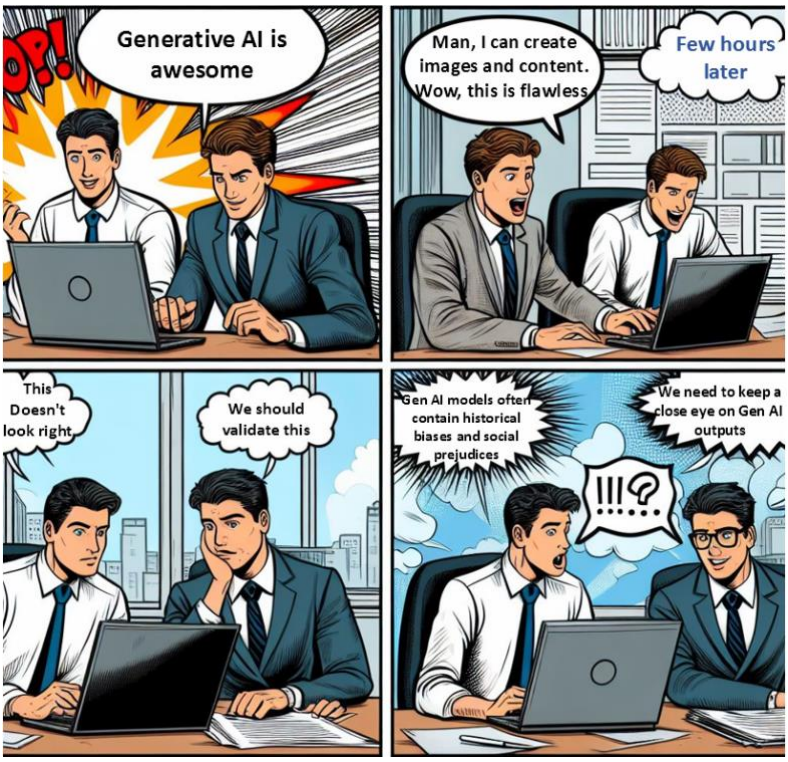

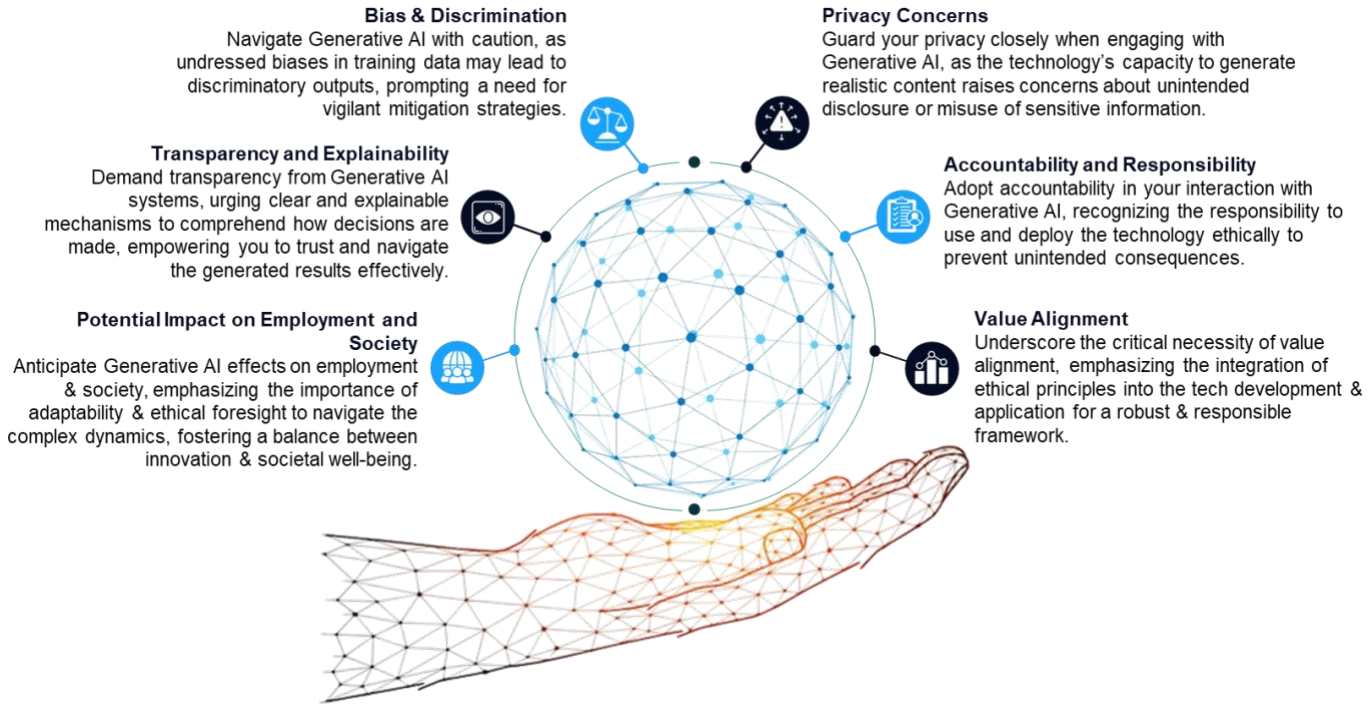

Bias and Discrimination:

Generative AI systems, if not carefully designed, may perpetuate biases present in training data, leading to discriminatory outcomes. Addressing bias is crucial to ensure fairness and equity in AI applications. For instance, if a Generative AI model is trained on biased data that disproportionately represents one demographic, it may produce outputs that perpetuate or even exacerbate existing societal biases, leading to discriminatory outcomes in areas like hiring or lending.

Privacy Concerns:

The vast capabilities of Generative AI raise significant privacy challenges, particularly in data generation and synthesis. Striking a balance between innovation and protecting individual privacy is essential to navigate this ethical terrain. We can see that in image synthesis, Generative AI might inadvertently generate realistic looking

faces of non-existent individuals, raising concerns about potential misuse for identity theft or other privacy violations.

Transparency and Explainability:

The opacity of many Generative AI models poses challenges for transparency. Ensuring that these systems are explainable is vital, enabling users to understand decisions made by AI and fostering trust in their outcomes. The issue with Black box was prevalent in AI and has cascaded in Generative AI. Hence it is important for Generative AI algorithm producing artistic content should offer an explanation for its creative choices, allowing users and stakeholders to understand the factors influencing its output, promoting transparency in the creative process.

Accountability and Responsibility:

Establishing clear lines of accountability for the actions of Generative AI is essential. Determining who is responsible in case of unintended consequences or ethical lapses ensures a framework of accountability in the development and deployment of these systems. For example, If an autonomous vehicle using Generative AI algorithms is involved in an accident, defining responsibility becomes critical – whether it lies with the vehicle owner, the manufacturer, or the developers of the AI system.

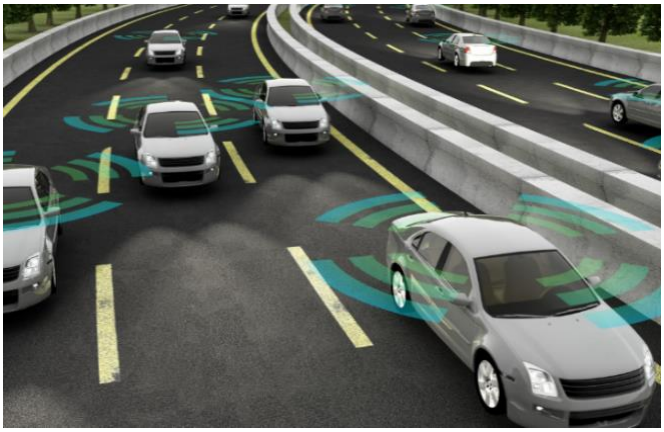

Potential Impact on Employment and Society:

Generative AI's transformative power in automating tasks may impact employment landscapes. Ethical considerations involve understanding and mitigating potential negative societal effects, ensuring a responsible integration that benefits, rather than harms, communities. Automation through Generative AI in manufacturing processes could lead to job displacement, necessitating ethical considerations for retraining and supporting affected workers to minimize negative societal impacts.

"The use of artificial intelligence[4], amid concerns from both writers and actors that unchecked AI could dramatically reshape Hollywood and undermine their roles, pitting artists against robots in a battle over human creativity."

Value Alignment:

Aligning the values embedded in Generative AI systems with societal norms and ethical principles is crucial. Ensuring that these systems contribute positively to human wellbeing and societal progress requires deliberate consideration and ethical alignment.

Exploring these ethical dimensions is vital in guiding the responsible development and deployment of Generative AI, contributing to a technology landscape that respects individual rights, fosters inclusivity, and aligns with societal values. This can be seen in the domain of healthcare, Generative AI assisting in diagnosis should align with medical ethical principles, ensuring that recommendations prioritize patient wellbeing, respect privacy, and adhere to established medical standards.

Regulating Generative AI

In the expansive world of Generative AI, the regulatory landscape serves as a critical checkpoint for responsible integration. A global survey reveals a mosaic of existing regulations governing Generative AI, highlighting variations in specificity and scope across different regions. Regulatory bodies, both at national and international levels, play pivotal roles in overseeing these technologies, but challenges persist, including the need for agile frameworks that can keep pace with the rapid evolution of Generative AI.

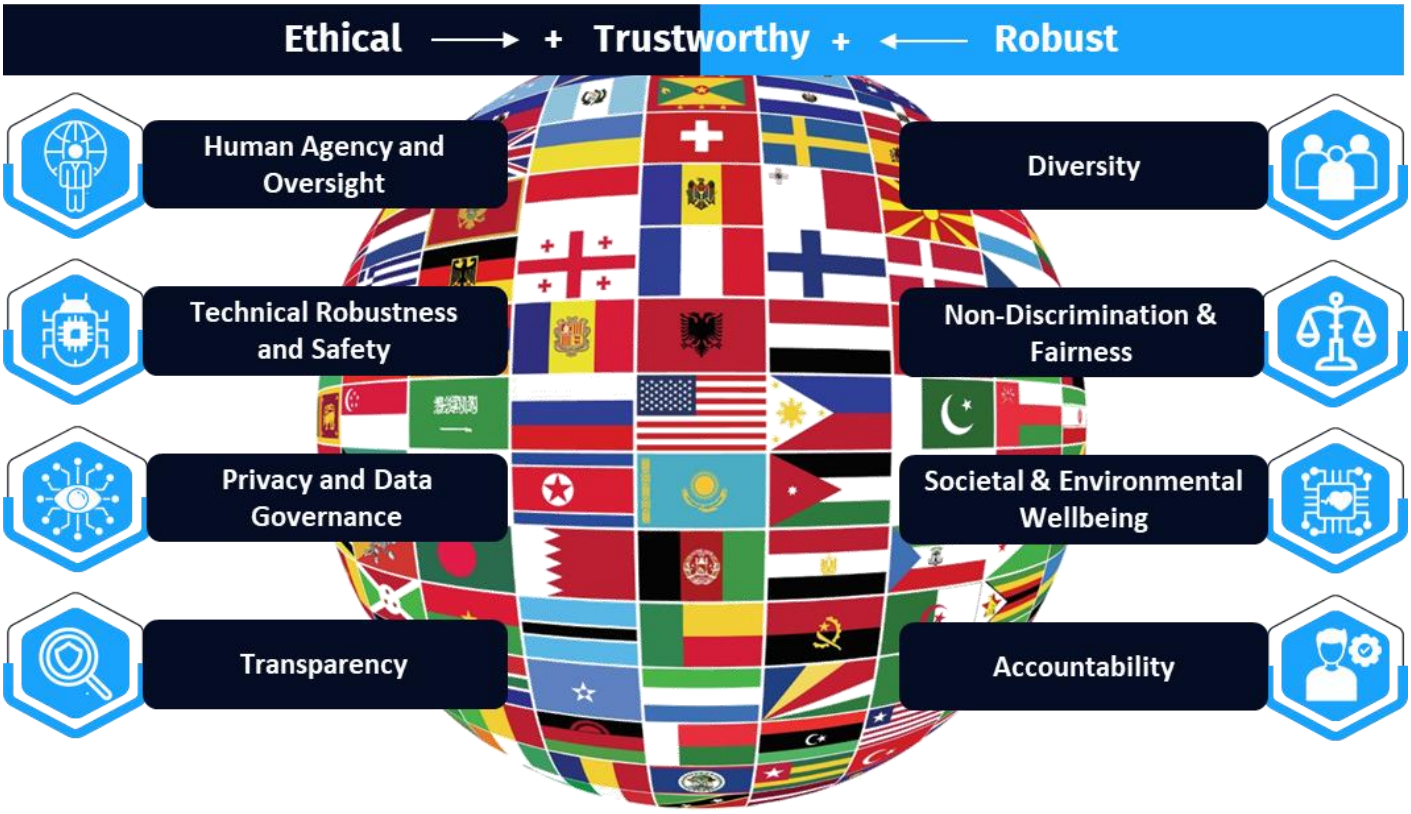

Identifying challenges and gaps within the current regulatory framework unveils potential obstacles in effectively governing Generative AI. The ambiguous legal terrain and the struggle to adapt regulations to technology's swift progression raise concerns. As Generative AI finds applications across diverse industries, from healthcare to finance, the need for industry specific regulations becomes apparent. This exploration also underscores the importance of international collaboration to harmonize standards, offering a unified approach to navigate the intricacies of Generative AI regulation on a global scale. In this dynamic landscape, understanding the current state of regulations sets the stage for the subsequent discussion on governance mechanisms crucial for the ethical development and deployment of Generative AI.

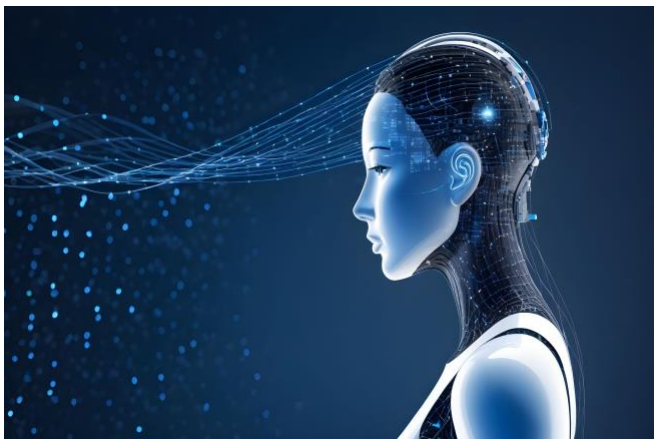

Regulations related to AI governance vary around the world.

As of November 2023, nonexhaustive

Type of policy: Nonbinding principles (eg, OECD)

- Japan
- Singapore
- United Arab Emirates
- United Kingdom
- United States
- Other OECD member countries

General AI legislation proposed or being finalized

- Brazil
- Canada
- China
- South Korea
- European Union

Example countries without general AI legislation

- Australia
- India
- New Zealand
- Saudi Arabia

McKinsey & Company[6]

Human Agency and Oversight:

Without adequate oversight, autonomous systems may make decisions leading to unintended harm. In automated content moderation, lack of human oversight can result in the wrongful removal of content, impacting freedom of expression.

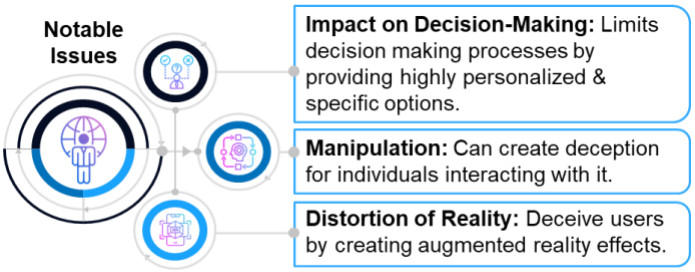

Remedies: The guidelines emphasize maintaining human control in critical decision points. This mechanism allows users to appeal AI decisions, ensuring human intervention in ambiguous cases.

"AI systems should empower human beings, allowing them to intervene and oversee decisions."

The issue of how humans[2] and AI can collaborate effectively while addressing ethical concerns was extensively disused in our paper "Working Hand in Hand with Artificial Intelligence."

Technical Robustness and Safety:

Inadequate robustness can lead to security breaches or unintended harmful outputs. In facial recognition technology, lack of robustness may result in misidentification, leading to wrongful arrests.

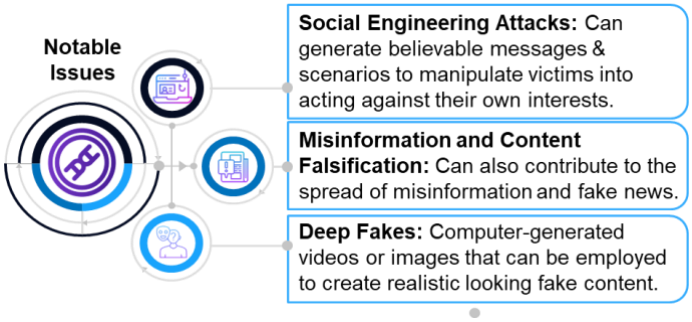

Remedies: The guidelines stress the importance of secure and resilient AI systems. Regularly updating algorithms, conducting rigorous testing, and implementing failsafe mechanisms to prevent misuse.

"AI systems should be technically robust and safe to use."

Privacy and Data Governance:

Poor data governance can lead to unauthorized access, data breaches, and privacy violations. For instance, in healthcare AI applications, insufficient data governance can compromise patient confidentiality and privacy.

Copyright and Intellectual Property: Concerns regarding the ownership of copyright for AI-generated works, as it can be difficult to determine who should be considered the creator or owner of the copyright in a machine-generated work.

Lack of Regulation: Absence of strong regulations, the responsibility falls on individual organizations and developers to ensure responsible usage of generative AI.

Remedies: Guidelines emphasize transparent and ethical data practices. Therefore, implement robust data protection measures, ensure transparent data handling, and obtain informed consent for data usages.

"AI systems should respect privacy."

Transparency:

Opacity in AI decision making may lead to distrust and unintended bias. Lack of transparency in algorithmic hiring systems can perpetuate gender or racial biases.

Black Boxes: Generative AI models lacks interpretability or transparency in the decision-making processes.

Remedies: The guidelines advocate for transparent and explainable AI systems. Implement mechanisms to provide understandable explanations for AI outputs, allowing users to comprehend and challenge decisions.

"The principle of transparency implies that decisions made by AI should be explainable."

Diversity, Non-discrimination, and Fairness:

Biases in AI systems can perpetuate discrimination and social inequalities. Bias in predictive policing algorithms can disproportionately target certain communities, exacerbating social injustice.

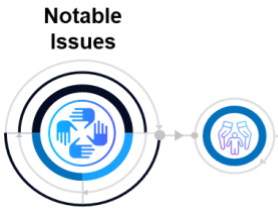

Discrimination and Bias: Can perpetuate societal discrimination and inequality. Some biases may be introduced into the training dataset due to existing societal biases.

Remedies: Guidelines stress avoiding biases and ensuring diverse representation in AI development. Regularly audit and address biases in training data, algorithms, and decision-making processes to achieve fairness.

"AI systems should be accessible, allowing everyone to benefit fairly."

Societal and Environmental Wellbeing:

Unintended negative societal impacts or environmental harm. Unchecked use of AI in high frequency trading can contribute to market volatility, negatively impacting economic stability.

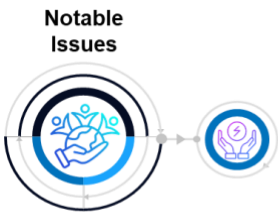

Power Consumption: AI systems require high-end computers and servers to run, consuming a significant amount of electricity to function correctly. Training and using generative AI systems can demand substantial power and computational resources.

Remedies: Guidelines highlight the need for AI systems to benefit humanity. Develop and deploy AI applications that contribute positively to societal wellbeing and environmental sustainability, aligning with ethical principles.

"AI systems should be used to enhance positive societal outcomes."

Accountability:

Establishing clear lines of accountability is crucial. In legal applications, understanding which entities are responsible for AI-generated content ensures accountability in case of misinformation or unintended legal implications.

Lack of Clear Responsibility: Results of AI's decisions and actions are sometimes unpredictable. As a result, it might be challenging to identify who, if anybody, is accountable for the outcomes of a particular AI-driven process that needs to be fixed.

"As generative AI technology advances, there is a need for precise regulation."

Challenges in Implementing Ethical Regulations in the Landscape of Generative AI

Translating the ethical regulations outlined in the 2019 European Commission Guidelines into practical implementation within the dynamic landscape of Generative AI is confronted with significant challenges. One inherent difficulty lies in the technical complexity of Generative AI algorithms, particularly within deep learning and neural networks, making it intricate to align these models perfectly with ethical principles. For instance, in scenarios like image generation, achieving transparency and explainability becomes a formidable task when the AI model generates highly detailed and complex visual outputs.

Data privacy concerns further complicate the implementation of regulations, as there is a delicate balance between the need for vast datasets to train robust Generative AI models and the ethical imperative of maintaining data privacy. In domains like medical image synthesis, where comprehensive datasets are crucial for accuracy, ensuring patient privacy becomes challenging when striving for transparency and accessibility. The rapid evolution of Generative AI poses another hurdle, outpacing the adaptability of regulatory frameworks and creating a lag that struggles to keep pace with technological advancements. For instance, the emergence of novel Generative AI techniques, such as GPT-4, may introduce new ethical considerations not covered by existing regulations, resulting in a gap between technological progress and regulatory updates.

Why is Difficult to Implement Ethical Regulations?

Technical Complexity: Generative AI's intricate algorithms, especially in tasks like image generation, present challenges in achieving transparency and explainability.

Data Privacy Balance: Balancing the need for vast datasets with maintaining privacy, crucial in domains like medical image synthesis, is a persistent challenge.

Regulatory Lag: The rapid evolution of Generative AI creates a regulatory lag, struggling to keep pace with advancements, evident with novel techniques like GPT-4.

Bias Mitigation Challenge: Identifying and addressing biases within large datasets, as seen in language generation, poses complexities in achieving unbiased AI systems.

Collaboration Hurdles: Effective collaboration between AI developers, ethicists, policymakers, and legal experts faces challenges, impacting scenarios like legal document summarization.

Global Regulatory Fragmentation: Variations in regulatory approaches globally create a fragmented landscape, adding complexities in deploying universally accessible AI applications.

Mitigating biases within Generative AI systems proves to be a complex challenge, given the intricate task of identifying and addressing subtle biases within large datasets. In language generation, biases encoded in historical text data

can inadvertently perpetuate stereotypes, making it challenging to completely eradicate bias without compromising language coherence. Interdisciplinary collaboration also presents hurdles, requiring effective alignment between AI developers, ethicists, policymakers, and legal experts. In scenarios such as developing Generative AI for legal document summarization, misalignment between technical specifications and legal nuances hinders collaboration, posing challenges in ensuring AI-generated summaries adhere to legal and ethical standards. Finally, global regulatory fragmentation introduces complexities as variations in regulatory approaches across different regions create a fragmented landscape for Generative AI governance. Deploying globally accessible Generative AI applications becomes challenging when compliance requirements vary widely, necessitating complex, region-specific adaptations. Recognizing and addressing these challenges is imperative for establishing a regulatory environment that effectively guides the responsible and ethical development of Generative AI technologies.

Navigating Generative AI Governance Solutions

The ethical considerations discussed earlier have been addressed to an extent. However, now it's imperative for us to establish an effective governance of Generative AI. As we explore the intricate ethical landscape, it is crucial to transition seamlessly into the realm of governance recommendations that address these challenges. The ethical dimensions laid out underscore the urgency for a comprehensive governance framework capable of steering the responsible development and deployment of Generative AI. Also, the critical need for governance mechanisms that not only respond to these challenges but also proactively shape the trajectory of this transformative technology. In this section we explore key recommendations for Gen AI governance, seeking to strike a balance between innovation and ethical responsibility.

Holistic Governance Approach

Propose a governance approach that transcends mere regulatory measures by embracing a holistic perspective. This involves considering not only the technical intricacies but also the broader ethical and legal dimensions of Generative AI. A comprehensive strategy ensures that the governance framework is well-rounded, addressing the multifaceted challenges posed by this transformative technology.

Dynamic Regulatory Frameworks

Advocate for regulatory frameworks that exhibit adaptability and responsiveness to the rapid evolution of Generative AI. Static regulations risk becoming outdated in the face of technological advancements, making it imperative to implement frameworks that can dynamically adjust to emerging ethical considerations. This approach strikes a balance between fostering innovation and upholding ethical standards.

Industry Collaboration

Encourage collaboration within the industry to establish shared guidelines and best practices for the development and deployment of Generative AI. By fostering a collective effort, this recommendation aims to create a cohesive and responsible approach across the tech landscape. Industry collaboration reduces the risk of fragmented governance approaches, promoting consistency and ethical practices in the development of Generative AI technologies.

Anticipated Impact

Implementing these recommendations anticipates a transformative impact on the governance of Generative AI. A holistic governance approach ensures that policies go beyond addressing immediate challenges, providing a framework capable of navigating the nuanced ethical considerations of this powerful technology. Dynamic regulatory frameworks allow for timely adaptation to emerging ethical concerns, ensuring that governance remains effective and relevant in the face of rapid technological advancements.

Collaborative efforts within the industry establish a united front, promoting ethical practices and reducing the risk of fragmented approaches. As a result, the anticipated impact is a governance ecosystem that not only effectively addresses current challenges but also proactively anticipates and mitigates future ethical concerns. This approach prioritizes transparency, accountability, and inclusivity, instilling public trust in the development and deployment of Generative AI and fostering a technologically advanced yet ethically grounded future.

Gen AI Outlook at NTT DATA

At NTT DATA, we recognize the complex landscape stakeholders and clients navigate when aligning AI systems with evolving regulatory frameworks to ensure trustworthiness and compliance. Acknowledging the diverse challenges posed by each regulatory provision, adapting AI design and development processes can be daunting. To address this, we proactively assist stakeholders and clients in anticipating their needs concerning EU AI regulations.

As an initial step, we have developed the AI Act Audit Tool—a pivotal asset designed to assess existing gaps in meeting each regulatory requirement. This tool serves as a comprehensive guide, evaluating the compliance maturity of AI models with the new European AI Act, adhering closely to the Guidelines for a Trustworthy AI. Beyond regulatory alignment, our proprietary AI Governance assessment empowers AI leaders across various industries. This assessment aids in defining a strategic roadmap, ensuring a robust and compliant AI Governance framework is in place for the successful deployment of AI technologies.

NTT DATA AI Centre of Excellence, has developed our proprietary AI Governance assessment to enable AI leaders from diverse industries, supporting the definition of their roadmap towards deploying AI Governance.

As an outcome, the assessment translates the company´s diagnosis on their AI maturity level and serves as a guide to build the action plan for becoming an AI-driven organization.

Diagnose AI Maturity of your Organization

As the first step, identifying current needs & gaps within the AI lifecycle is a valuable exercise to translate visibility on the end-to-end & build a unified AI vision.

Orchestrate AI & Business Teams

Deploying AI as a strategic lever requires cross-hierarchy & cross-functional collaboration. Particularly, AI & business teams need to work closely at defining AI opportunities in alignment with organization objectives.

Plan for Long-term, Act for Short-term

Journey towards greater AI maturity should be shaped around a common path in which capabilities-wise every AI initiate builds upon the next. By supporting the AI transformation, the action should become the backbone of the AI strategy.

At NTT DATA our commitment extends beyond mere compliance – we have been addressing these issues with our dedicated ***"AI Centre of Excellence[5]",*** where we aim to empower our clients with the tools and insights needed to foster reliability, transparency, and trustworthiness in their AI endeavors. Through our innovative audit solutions, we pave the way for organizations to navigate the intricacies of AI regulations while fostering a culture of responsible AI development and deployment.

Conclusion

While the excitement surrounding the advancements in AI is evident, it is imperative to maintain vigilance over the ethical, legal, and societal implications that accompany such progress. The need for ongoing validation and a thoughtful consideration of potential consequences is paramount. At NTT DATA, our unwavering commitment to nurturing reliability and transparency stands as a guiding principle. By prioritizing these values, we contribute to shaping a future where AI is not only innovative but also safe and trustworthy. This emphasis on responsible AI development is crucial for building a sustainable foundation that aligns with ethical standards and societal expectations, ensuring a positive impact on our collective future.

About the author

Tanvir Khan is chief digital and strategy officer focusing on technology direction, go-to-market and offering management. With more than 25 years of experience in the IT industry, he is a thought leader in digital transformation, associated core technologies and value

realization. He is also a hands-on IT practitioner with five patents and four pending patents in AI and automation. As a spokesperson for NTT DATA Services, Tanvir shares his insights to clients, media and analysts on topics ranging from Generative AI to emerging global service delivery locations.

Let's get started

See what NTT DATA can do for you.

- Deep industry expertise and market-leading technologies
- Tailored capabilities with your objectives in mind
- Partnerships to help you build and realize your vision.

Contact one of our authors or visit **nttdata.com** to learn more.

Sources

1. https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai
2. https://us.nttdata.com/en/-/media/assets/brochure/human-centric-ai-pov-final.pdf
3. https://philpapers.org/archive/MLLHOD.pdf
4. https://www.theguardian.com/culture/2023/oct/01/hollywood-writers-strike-artificial-intelligence
5. https://www.nttdata.com/global/en/news/press-release/2019/may/ntt-data-launches-artificial-intelligence-center-of-excellence
6. https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/as-gen-ai-advances-regulators-and-risk-functions-rush-to-keep-pace

All images in this report were generated using Azure Open AIs