seramount

WORKBOOK

Navigating the Intersection of AI and DEI

A Practical Workbook for Diversity Leaders

Table of Contents

The Intersection of AI and DEI

Artificial Intelligence (AI) and Generative AI (GAI) are rapidly transforming the business landscape, presenting both unprecedented opportunities and complex challenges. As organizations increasingly adopt these technologies, leaders are grappling with their associated risks and implications. The fast-paced development has created an urgent need for executives to understand and navigate this new terrain.

What is AI and GAI?

To effectively address these challenges, it's crucial to understand the distinctions between AI and GAI:

- **AI**, which has been around for decades, enables machines to perform cognitive functions like learning and problem-solving. It's already part of our daily lives through voice assistants and search engines.
- **GAI**, a recent and specialized form of AI, can create new, high-quality content such as audio, code, images, and text. This unique ability to generate novel content sets it apart from traditional AI.

The impact of these technologies is profound and widespread. As of 2024, McKinsey reported 65% of organizations use GAI in at least one business process, highlighting its rapid adoption and transformative potential.

As businesses increasingly adopt GAI and other AI technologies, leaders are grappling with the associated risks and challenges. The fast-paced development of AI has created a sense of urgency among executives to understand its full implications.

As AI and GAI integration accelerates, their influence extends beyond technical domains, creating new challenges for organizational leadership—particularly in areas of diversity, equity, and inclusion.

Uncertain about GAI? You're not alone.

Only

10%

of CDOs feel knowledgeable about Generative AI

1 in 2

CDOs do not feel their organization is prepared to address GAI advancements

The Role of Chief Diversity Officer

Chief Diversity Officers (CDOs) are finding themselves in a unique position, increasingly being called upon to become experts in GAI due to the profound implications its development and deployment have on workplace dynamics, culture, and inclusivity.

Despite this demand, Seramount research shows only 9% of CDOs feel prepared to address GAI.

Given these challenges, it's no surprise that the intersection of AI and Diversity, Equity, and Inclusion (DEI) has risen to the forefront of organizational priorities; Seramount research reveals it as one of the top two concerns for CDOs.

The Intersection of AI and DEI (cont.)

AI's Impact on DEI Priorities

As GAI becomes more deeply integrated into the workplace, DEI leaders must be prepared to navigate a broader range of AI-related challenges and opportunities; organizations that leverage GAI in the most thoughtful and responsible ways will emerge as the biggest winners.

That's why it is critical for DEI leaders to develop a comprehensive understanding of AI's implications: the good, the bad, and the ugly. This knowledge will enable them to engage in informed conversations with organizational leaders and guide decision-making processes regarding GAI implementation.

Why DEI Leaders Should Engage in GAI Discussions

Ensure Inclusive Development

Your input can help ensure diverse perspectives are considered in AI development and implementation.

Leverage AI for DEI Goals

When used responsibly, AI can be a powerful tool for understaffed and underbudgeted DEI departments, amplifying your impact.

Demonstrate Broader Impact

Helping mitigate AI risks showcases how DEI principles apply across the organization, enhancing your strategic value.

Prevent Unintended Consequences

Your involvement can help prevent AI systems from inadvertently perpetuating or exacerbating existing inequities.

Purpose of This Workbook

This workbook is designed to equip DEI professionals with the essential knowledge needed to contribute meaningfully to discussions surrounding AI in the workplace. Through a combination of key information and reflective exercises, we aim to empower DEI leaders to navigate the complex intersection of AI and DEI effectively.

Generative AI and Bias

SECTION 1

Understanding Generative AI and Bias

GAI bias is, to some extent, unavoidable. This is true primarily because GAI systems are trained on human-generated content, which inherently contains societal biases.

To understand this, we need to break down what "training" means in the context of GAI:

The GAI Learning Cycle

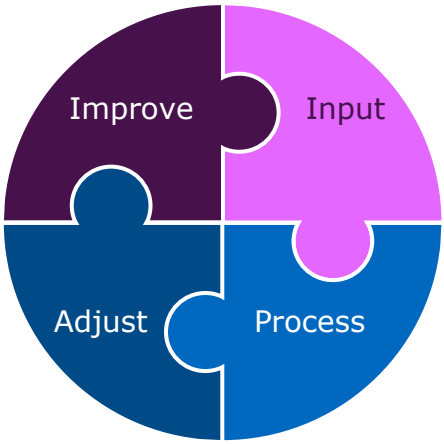

Input

GAI systems are fed vast amounts of human-generated content.

Process

The AI processes and learns from the data, including inherent societal biases.

Adjust

GAI makes adjustments based on the analyzed data and human feedback.

Improve

GAI tools are further trained through additional testing, auditing, and iterative training.

Understanding the Nature of GAI Bias

Several key points are crucial to understanding GAI bias:

- GAI isn't inherently more biased than humans; it's learning from and mimicking human-produced content.
- The technology doesn't "know" what it's producing. It makes probabilistic assumptions based on patterns recognized in its training data.
- Existing patterns of bias and discrimination in the workplace and society are reflected in the data used to train GAI.

Correcting AI's Biased Lens

In 2023, Brown University researchers found that major text-to-image AI systems consistently generated images of white male surgeons when asked to depict a surgeon. This bias was later corrected through human feedback, demonstrating the cycle of GAI learning and improvement in action.

The Role of Humans

Given these realities, human oversight and intervention are crucial in mitigating GAI bias. As one expert noted, "Humans must tell AI what they consider suitable, teach it which information is relevant, and indicate the outcomes they consider best."

Just as we've developed strategies to mitigate human biases in decision-making processes, we need to apply similar principles to GAI. The goal isn't to eliminate bias entirely (as it is likely impossible) but to recognize, understand, and mitigate it where necessary.

Reflection Exercise: Bias

Is GAI biased?

Your CEO shares with you they recently read an article about how a competitor is using AI. They ask you, "Isn't generative AI biased?"

Write 3 key points you would make to your CEO about GAI bias on your worksheet below.

What key points on GAI bias would you make?

1.

2.

3.

Mitigating Generative AI's Risks

SECTION 2

Mitigating Generative AI's Risks

While bias in Generative AI (GAI) is widely discussed, CDOs must also be aware of other, less discussed risks. Understanding these risks is crucial when contributing to organizational guidelines such as AI Acceptable Use Policies to ensure comprehensive safeguards are in place to cover a range of GAI activities.

GAI's Lesser-Known Risks

PRIVACY

GAI tools may inadvertently expose private employee data if fed sensitive information, potentially violating privacy rights and leading to legal consequences.

SECURITY

Interactions with GAI may lead to the unintended disclosure of sensitive organizational information, resulting in compliance violations and a loss of stakeholder trust.

HALLUCINATIONS

GAI can generate convincing but entirely false information that appears credible, leading to the spread of misinformation and potentially causing misguided decision-making.

This is a non-exhaustive list, underscoring why mitigation strategies are key to ensuring that emerging issues during GAI use can be swiftly identified and addressed.

Strategies to Mitigate These Risks

CDOs have a unique opportunity to shape GAI by partnering with key tech leaders at their organization.

The worst-case scenarios that have made headlines have also spurred effective solutions. While each organization's approach will vary based on GAI usage, consider these five key strategies:

1. Ensure diverse representation in development teams
2. Conduct bias audits during development and after implementation
3. Establish ethical guidelines and policies for the entire GAI spectrum of activities
4. Continuously monitor and evaluate outputs
5. Promote transparency and accountability

Urgency to Act

Remember: 70% of employees are currently using GAI for work without managerial disclosure. This means these risks likely already exist within your team and organization.

Reflection Exercise: Mitigating Risks

What does your organization need to know about adopting GAI from a DEI perspective?

Prioritize the following actions you would take to mitigate GAI bias. Consider what is in your locus of control, the greatest opportunity for impact, and your organization's current GAI maturity.

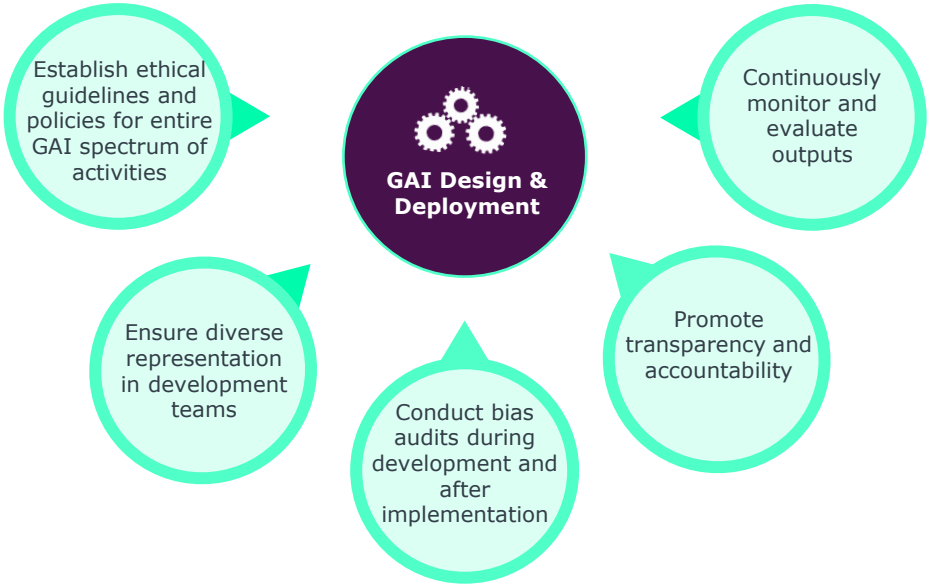

Answer these questions:

1. What is your top priority and why?

2. What stakeholders will you need to engage to accomplish this priority?

3. What barriers might you encounter in addressing this priority?

seramount

Generative AI as a DEI Catalyst

SECTION 3

Generative AI as a DEI Catalyst

While not without risk, GAI also offers powerful solutions that can drive real impact for your organization. Nowhere is this potential more evident than in staffing, where GAI tools are helping organizations overcome long-standing challenges in building diverse teams.

Overcoming Hiring Challenges with GAI

CDOs are well aware of the growing challenge of increasing the percentage of historically excluded talent (HET) in organizations.

External pressures, such as the affirmative action ruling and broader attacks on DEI initiatives, alongside with limited direct authority over hiring processes, make achieving diversity goals seem difficult.

However, GAI presents a unique opportunity to overcome these hurdles**.**

> Regardless of how hard the best recruiter can try, they may still have some unconscious human biases…This is where AI can come in. It can ignore factors like a candidate's personality, education, gender, or ethnicity. Instead, it can focus on the qualifications, experience, and skills that might make them great candidates.
>
> - **Rebekah Bastian** in *Forbes*

The GAI Advantage in Filling Open Positions

Organizations implementing GAI solutions across their staffing and hiring processes are seeing positive, measurable results. Importantly, these solutions not only benefit DEI goals but also align with broader organizational metrics of success.

Developing Job Descriptions

An Australian software company implemented Textio, a GAI program, to reword job ads to appeal to a broader demographic. After deploying this software, they say applications for open positions from women increased from 18% to 57%.

↓

Increase in HET applicants

Sourcing Applications

CareerBuilder used a GAI tool that provided real-time feedback on how well job postings performed with applicants from a target audience. Adjustments based on this feedback increased target audience applications by 175%.

↓

Increase in target applicants

Screening Candidates

Using Pymetric, an AI tool that analyzes candidate data to identify applicants most likely to succeed, Deloitte saw a 60% reduction in screening process time and a 30% increase in historically excluded talent hires.

↓

Reduction in screening time

Reflection Exercise: Recruiting

Where should DEI be involved in GAI's deployment across HR business processes?

Answer these questions:

1. What steps in your process to fill open positions do you think will benefit the greatest from GAI integration?

2. What stakeholders do you need to influence and mobilize to integrate GAI?

3. Within what steps, if any, is GAI integrated in your process to fill open positions?

4. What outstanding questions do you still need answered to push this integration forward?

seramount

15 Key Terms to Build Your Foundational Knowledge

SECTION 4

Generative AI (GAI) 101: 'Get Smart' with These Definitions[1]

15 Definitions to Build a Foundational Understanding of GAI and Its Components

1. **Artificial Intelligence (AI):** AI is a machine's ability to perform the cognitive functions we usually associate with human minds.

2. **Data:** Data refers to pieces of information that AI systems use to learn and make smart decisions.

   - **Structured Data:** Structured data is information organized in a tabular format (for example, tables, databases, or spreadsheets) that can be used to train some machine-learning models.
   - **Unstructured Data:** Unstructured data is information that lacks a consistent format or structure (for example, text, images, and audio files) and typically requires more advanced techniques to extract insights.
   - **Training Data:** Training data is the information used to teach an AI system. It is essential to use diverse and unbiased data to avoid unfair decisions being made by the AI.

3. **Algorithm:** Algorithms are sets of instructions that guide AI systems on what decisions to make. These instructions can be created either by humans or learned by the AI system itself.

4. **Rules-Based AI:** A series of "if/then" rules structured for a machine to follow. Early versions of Siri, Alexa, and Google were rules-based because they were usually predictable and repetitive.

5. **Machine Learning (ML):** ML is a way for AI systems to learn from data without explicit programming. ML uses algorithms to identify patterns and build models for decision-making on their own based on historical data.

6. **Neural Networks:** Neural networks are modeled on how human brains work. Rather than being given a set of explicit rules, these systems are "taught" by looking at the patterns in relationships between things, processing information in layers to solve problems. ChatGPT uses a neural network as it requires large data sets and massive amounts of training on the data.

7. **Deep Learning:** Deep learning is a more advanced form of neural networks, helping AI systems understand complex patterns by using multiple layers. These networks are widely used in image and speech recognition tasks.

1) Adopted from Center for Integrative Research in Computing and Learning Sciences (CIRCLS) "Glossary of Artificial Intelligence Terms for Educators," McKinsey's "Glossary of AI Terms," and MIT Technology.

Source: McKinsey (2023); CIRCLS (2023); MIT Technology Review (2023); Seramount interviews and analysis.

GAI 101: 'Get Smart Quick' on AI with These Top Definitions[1]

8 **Transformers:** Transformer models are a type of neural network used in language modeling. Transformers utilize a self-attention mechanism, which enables them to focus on relevant parts of the input and output sequences. This attention mechanism allows the model to weigh the importance of different elements within the sequence, leading to better contextual understanding and more accurate predictions. Basically, transformers can track contextual information.

For example, *MIT Technology Review* gives an example of how the word "hot dog" takes on different meanings in these sentences: *"a hot dog should be given lots of water"* versus " *a hot dog should be eaten with mustard."* Transformers can tell the difference between these two meanings of "hot dog."

- **Self-Attention Mechanisms:** These mechanisms determine the important aspects of input. The inputs are inspired by how humans can direct their attention to important features in the world, understand ambiguity, and encode information.

9 **Foundation Models (FM):** FM are deep learning models trained on large quantities of training/unstructured data that can be utilized for a wide range of tasks out of the box or adapted to specific tasks via fine-tuning. These models are the "brain" behind GAI tools such as ChatGPT, Dall-E, and Bard.

- **Natural Language Processing (NLP):** NLP is a way for computers to understand human language, which enables them to do things such as converting speech to text or correct grammatical errors.

- **Large Language Models (LLMs):** An LLM—a type of foundation model—can process massive amounts of unstructured text and learn the relationships between words or portions of words. LLMs can generate natural language text, performing tasks such as summarization or knowledge extraction. Chat-GPT is an example of an LLM.

10 **Generative AI (GAI):** GAI is a type of AI that is built using foundation models and has the ability to generate new, high-quality content efficiently, including audio, code, images, text, simulations, and videos.

11 **Chat-Based Generative Pre-trained Transformer (ChatGPT):** ChatGPT is an AI system based on a transformer model, designed for natural-language processing tasks. It can generate responses to questions because it has been trained on a great deal of information from the Internet to give helpful responses to questions. In other words, it is pre-trained on a large amount of web text to effectively process sentences.

12 **Multimodal Models:** Multimodal models generate AI content in multiple modalities (e.g., text, graphics, audio, video). These models are trained by data in various formats and learn to combine information from different sources to produce one output.

1) Adopted from Center for Integrative Research in Computing and Learning Sciences (CIRCLS) "Glossary of Artificial Intelligence Terms for Educators," McKinsey's "Glossary of AI Terms," and MIT Technology.

Source: McKinsey (2023); CIRCLS (2023); MIT Technology Review (2023); Seramount interviews and analysis.

GAI 101: 'Get Smart Quick' on AI with These Top Definitions[1]

13 **Prompt Engineering:** Prompt Engineering is an essential process that involves designing, refining, and optimizing input prompts to guide a GAI model toward producing desired and accurate outputs.

- **Prompt Chaining:** Ability of AI to use information from previous interactions and prompts to inform and influence future responses. Part of good prompt engineering is the recursive nature of funneling prompts toward the end goal.

14 **AI Hallucination:** AI hallucination is a phenomenon wherein a large language model (LLM)—often a generative AI chatbot or computer vision tool—perceives patterns or objects that are nonexistent or imperceptible to human observers, creating outputs that are nonsensical or altogether inaccurate. These can be prevented and avoided using high-quality training data, defining the exact purpose the AI model serves, limiting possible responses, rigorously testing, and ensuring informed human oversight.[2]

- **Reinforcement Learning from Human Feedback (RLHF):** RLHF is a machine learning (ML) technique that uses human feedback to optimize ML models to self-learn more efficiently. Reinforcement learning (RL) techniques train software to make decisions that maximize rewards, making their outcomes more accurate. RLHF incorporates human feedback in the rewards function, so the ML model can perform tasks more aligned with human goals, wants, and needs. RLHF is used throughout GAI applications, including in large language models (LLM).

15 **AI Bias:** AI bias, also called machine learning bias or algorithm bias, refers to the occurrence of biased results due to human biases that skew the original training data or AI algorithm—leading to distorted outputs and potentially harmful outcomes. The models upon which AI efforts are based absorb the biases of society that can be quietly embedded in the mountains of data they're trained on. Historically biased data collection that reflects societal inequity can result in harm to historically marginalized groups in use cases, including hiring, policing, credit scoring, and many others.[3]

- **Algorithm Bias:** Misinformation can result if the problem or question asked is not fully correct or specific or if the feedback to the machine learning algorithm does not help guide the search for a solution.

- **Cognitive Bias:** AI technology requires human input, and humans are fallible. Personal bias can seep in without practitioners even realizing it. This can impact either the data set or model behavior.

- **Confirmation Bias:** Closely related to cognitive bias, this happens when AI relies too much on pre-existing beliefs or trends in the data—doubling-down on existing biases and unable to identify new patterns or trends.

- **Exclusion Bias:** This type of bias occurs when important data is left out of the data being used, often because the developer has failed to see new and important factors.

- **Measurement Bias:** Measurement bias is caused by incomplete data. This is most often an oversight or lack of preparation that results in the data set not including the whole population that should be considered. For example, if a college wanted to predict the factors to successful graduation, but included only graduates, the answers would completely miss the factors that cause some to drop out.

1) Adopted from Center for Integrative Research in Computing and Learning Sciences (CIRCLS) "Glossary of Artificial Intelligence Terms for Educators," McKinsey's "Glossary of AI Terms," and MIT Technology.
2) Adopted from IBM's definition page on hallucinations.
3) Adopted from IBM's definition page and Microsoft's Inclusive AI principles.

Source: McKinsey (2023); CIRCLS (2023); IBM (2023); MIT Technology Review (2023); Microsoft; Seramount interviews and analysis.

seramount

Conclusion & More Resources

SECTION 5

Conclusion

As DEI leaders, it's natural to be concerned about the challenges GAI may present to your DEI initiatives. However, it's crucial to recognize that GAI is not just another risk—it's a powerful tool to advance your goals.

The key lies in striking a balance between caution and innovation. By understanding and addressing the potential pitfalls of GAI, you can unlock its immense potential as a driver for your efforts. This includes overcoming long-standing DEI challenges, such as building diverse teams, as highlighted in our guide.

As we move forward in this AI-driven era, your role as a DEI leader becomes even more critical in shaping how these technologies are implemented and leveraged to create positive change. The future of DEI is intertwined with technological advancement, and by mastering tools such as GAI, you can lead your organization toward a more diverse, equitable, and inclusive future.

Explore More Resources

AI and the Future of Work

AI presents both challenges and opportunities for DEI leaders. Discover the five key takeaways to consider as you and your organization work to drive equitable outcomes along your AI journey.

Artificial Intelligence in Talent Conversations

Microsoft has long led the way in advancing inclusive experiences in tech. Join this session to hear their strategies for fostering inclusive workplaces in the AI era.

Research Review: The Impact of AI on Talent

AI has revolutionized industries worldwide, yet its impact remains complex and multifaceted. This session will review Seramount's latest research on the profound effects of AI on talent.

Seramount can help you transform challenges into opportunities and position your organization at the forefront of AI-driven DEI innovation. Our resources offer actionable insights to help you navigate the AI landscape, mitigate risks, and leverage GAI to drive meaningful progress in your DEI efforts.

Contact Seramount

seramount

We empower the world's most inclusive and high-performing workplaces.

Insight-Powered Solutions Trusted by Talent and HR Leaders

TALENT SOURCING	ASSESSMENT AND STRATEGY	RESEARCH AND ADVISORY SERVICES	LEARNING AND DEVELOPMENT
Identify, Engage, and Hire Top Early Career Talent	**Uncover the Real Levers to Unlock Enduring Employee Engagement**	**Set Strong DEI Foundations and Build Adaptable Roadmaps**	**Create a More Resilient, Connected Workforce**
Branding and recruiting platform that surfaces engaged, prepared candidates more efficiently, effectively, and at scale	Actionable intelligence into the employee experience powered by innovative technology and over 40 years of workforce data and insights	Expert guidance rooted in research and experience to support and advance your strategic talent and DEI priorities in every environment	Impactful content designed and delivered to inflect productivity and positive behavior change at every level

We partner with **600+** corporations, government entities, and nonprofits **globally**.

Including **~50%** of the Fortune 100 and **~25%** of the Fortune 500.

seramount.com

202-747-1000 seramount.com