Generative Artificial Intelligence and Machine Learning for Identity and Access Management

Anastasios Liveretos, Milena Lazarova

Abstract—**In recent years, the advent of generative artificial intelligence (GenAI) has significantly impacted identity and access management (IAM). Generative AI democratizes the use of Large Language Models (LLMs), enabling various IAM use cases, from adaptive access to identity threat detection and response. However, the rise of generative AI also introduces new challenges, such as creating deepfakes and sophisticated phishing attacks. This paper discusses the balance between leveraging AI for improved IAM and addressing the associated risks. It emphasizes the importance of selecting the right vendors, maintaining language models, and ensuring privacy and security. The future of IAM lies in intelligent automation, where AI-driven tools can self-configure and adapt to evolving threats. This paper provides insights into the current trends, challenges, and opportunities in integrating generative AI with IAM, offering a roadmap for organizations to navigate this rapidly changing landscape.**

***Keywords*—Generative Artificial Intelligence (GenAI), Identity and Access Management (IAM), Large Language Models (LLM), Machine Learning (ML)**

I. Introduction

There has been much hype about generative AI in the last year. Machine learning approaches to AI have been a crucial part of identity and access management for many years. Recent advancements in generative AI have democratized it and changed the game for best practices by enabling many people to leverage Large Language Models (LLMs) for various identity and access management use cases [1]. Artificial intelligence has tremendous potential, but caution is still advised.

Generative AI has the potential to create deepfakes that are very effective in circumventing biometrics. It can also turbocharge the creation of phishing emails, making them more detailed and cheaper.

The good news is that generative AI can also help prevent such fraudulent approaches. Thus, we see various use cases, some of which are still maturing.

Vendors vary widely in what they offer. However, we are at a stage where one should investigate and prototype generative AI for specific identity and access management use cases. Gartner predicts that by 2025, 35% of organizations will leverage LLMs for various identity and access management use cases to help identity fabric functions, help people be more efficient, and improve user experience [2].

II. Analysis

There are many things to bear in mind when considering AI. Getting the right vendors, as well as optimizing and maintaining the language models, can make it hard to know what to do next. The sunlight of improved security and operational usability can be offset by what is lurking in the shadows. It is not just the deepfakes created by attackers that can cause problems; misusing artificial intelligence can also cause many problems. Leveraging a third-party LLM, one has to ensure that they are using it safely, that they are not leaking their sensitive information, and that their privacy is not violated.

So, where we are now is that many use cases are driving the greater adoption of GenAI. A lot of this comes down to dealing with the rapid pace of change. There is so much change in identity and access management that it is impossible to keep up with it without automation.

In fact, more than mere automation is required. It would be best to have the automated tools to teach themselves and change. In other words, we need artificial intelligence.

There is a demand for continuous insight and more frequent updates on protecting against rapidly evolving threats and actors [3]. There is also a demand for a better user experience to keep up with all these changes. Even if one had an unlimited budget and could hire everyone needed, they could only manage a large identity deployment today with machine learning.

Years ago, when the industry started talking about machine learning analytics, one of the first use cases was adaptive access; to this point, adaptive access is just table stakes. Nowadays, many other use cases are associated with identity governance and administration, which vary in their use of artificial intelligence.

III. Use Cases

One of the reasons machine learning-based analytics helps with identity and access management is that it tunes for particularly narrow use cases. So, when thinking about rolling out more analytics for identity and access management, one should consider rolling it out for a specific use case at a time, testing, and ensuring that the way that piece was tuned will be appropriate for their organization.

A. Access Requests

Some use cases are based on simple analytics, such as access requests, which can use advanced clustering techniques to help group entitlements into roles.

The access request is one of the more interesting ones. It is a combination of automation helping humans. It is an augmented use of artificial intelligence rather than autonomous use. Access requests, when leveraging artificial

Received: 04.11.2024
Published: 03.12.2024
https://doi.org/10.47978/TUS.2024.74.03.019

Anastasios Liveretos is with the Technical University of Sofia, Sofia, 1000, 8 Kl. Ohridski Blvd, Bulgaria (Anastasis.liveretos@cchellenic.com)
Milena Lazarova is with the Technical University of Sofia, Sofia, 1000, 8 Kl. Ohridski Blvd, Bulgaria (milaz@tu-sofia.bg)

intelligence, can leverage the richness of the user's request [4].

If the request is low enough risk, access can be granted automatically. If it is a higher risk, it can go through and have management approval as usual.

This automated risk-based granting of access saves time for both the user who gets their access quickly and the management.

B. Behavioral Biometrics

Today, there are a much wider number of identity and access management use cases that are readily used in production. In access management, we have passive behavioral biometrics. In jurisdictions where they are allowed, passive behavioral biometrics look at how the user interacts with a mouse or a keyboard in their device [5]. It can help detect whether the user is a robot, another human, a human behaving like the user typically does, or an account takeover fraud. Many machine learning-generated analytics are helping to ensure that the users are who they say they are.

C. Cloud Infrastructure Entitlement Management

Just as we want a regular end user protected by machine learning, we certainly want a sensitive administrator protected by that. And there is a relatively newer use case category: cloud infrastructure entitlement management. An end user or an organization with many different software developments in the cloud, particularly in the multi-cloud, with many other computing and storage resources, needs machine learning to check and ensure adequate access configurations.

Inadequate access configurations are still one of the significant sources of cloud breaches. Organizations often need help to configure adequate access to some of these resources [6]. Some vendors can consider this cloud infrastructure entitlement management capability part of privileged access management, but others part of identity governance and administration. The point here is that nobody wants to think about features belonging to separate tools. One should think about capabilities and use cases suitable for their organization.

D. Threat Detection and Response

Another relatively new area is identity threat detection and response, which also uses these same analytic techniques for a narrow use case of protecting the identity systems themselves. Machine learning analytics are great for detecting anomalies. Thus, they are also very helpful in hygiene use cases. This works because many identity data are used to train the machine-learning models to behave as they should. Identity sources come from many identity tools, whether the identities themselves or policies and roles, entitlements, or log files. Keeping these data clean is very important, and one of the significant threats is injecting damaging data into machine-learning models to confuse them [7].

Identity data also helps specific identity functions operate better. There is also a big trend to integrate identity systems more closely with security systems and vice versa.

A distributed intelligence model is evolving in identity access management and security spaces. The central model is in security systems with an extended detection and response or security information event management approach. Still, more minor, focused use cases that quickly examine a small amount of data to be more responsive are also met.

Identity threat detection and response examples are smaller case examples. This may be referred to as distributed intelligence, but machine learning data scientists would call it federated machine learning.

IV. Machine Learning, Artificial Intelligence, and Generative AI Trends

Several trends exist in applying machine learning, artificial intelligence, and generative AI models, but there is an apparent move towards a more federation use. This is part of making identity systems much more responsive and adaptive.

Three different trend lines are influencing the evolution of identity and access management analytics.

A. Security Information Event Management

The first of these trends is coming from security. Security information event management systems have been around for a very long time. They collect the data logs in one place to be analyzed, which is intendedly incredibly labor intensive. User entity behavior analytics was invented to enable those systems to analyze log files with machine learning intelligence to do the job much more quickly [8].

B. Advanced Analytics and Fraud Detection

The second thread of evolution comes from the banking industry. Banks have long been doing advanced analytics and fraud prevention and have been evolving to meet the needs of more specific identity use cases, such as new account fraud and account takeover fraud [9].

C. Generative Artificial Intelligence

And then we have the generative AI thread. Generative AI, which is relatively newer, tends to require more human intervention at this stage and has much potential to be layered on with some of the other techniques. One of the most beautiful and essential things about generative AI is that its speech is made in natural language. It enables people to communicate with computers more efficiently, reduces training, and enables the identity systems to give alerts and responses to them more naturally in sentences, paragraphs, and numbered task lists. One of the significant advantages that will further drive the evolution of artificial intelligence in the identity space is the ability to use LLMs to change models themselves. That is a game changer. We are still in the very early stages of seeing what that will do in identity and other use-case situations. In the future, more identity use case models will be composite models, where composite might be multiple machine learning ones or generative AI with machine learning.

V. Vendor Landscape and Enterprise Considerations

Many vendors are starting to leverage analytics and

machine learning, particularly in AI. Some vendors have announced that they are using and making available generative AI features for identity and access management use cases.

Things such as agents or copilots of different vendors have different names for this. Still, companies such as Microsoft, Radiant Logic, and Axiomatics have already had this available. So, there are starting to be specific packaged generative AI capabilities to help manage identity and access. Organizations may be encouraged to experiment with them as they are available and see their suitability.

There is a subtext to this: getting significant language and machine learning models to work correctly for specific identity and access management use cases is a lot of work. Unless they are the largest, most organizations should buy their artificial intelligence prepackaged as part of an identity and access management tool. Of course, applying such powerful technologies to a class of use cases has many advantages, strengths, and weaknesses [10].

Some of the advantages of using machine learning-driven artificial intelligence are the same as those for large language models, and there are places where they are different. Machine learning improves accuracy, makes decisions more rapidly, and mitigates risk.

Both machine learning and generative AI can help with the automation process, but generative AI benefits the user experience, communication, and interaction. Both involve complexity and can be contaminated by flawed data or incorrectly designed models. Therefore, the responses from artificial intelligence must be checked to make sense.

This is a big issue, particularly in the area of generative AI. Generative AI does not have a good model of reality. What it does is understand language and communication. It is a large language model. Stanford calls them foundation models. At this stage, most generative AI use cases want human supervision. So, it is beneficial to generate a draft; before doing anything with it, somebody who knows what is going on reviews it [11].

Machine learning is a vital part of the identity fabric architecture. Analytics will be central to modern identity systems in the future, not just in automation but also in auto-discovery.

As we advance, the Identity system should evolve so that it helps self-configure itself, and auto-discovery is a big part of that. Still, human augmentation has been needed for a very long time.

Continuous risk evaluation is critical for all access management, incredibly remote access management. Third-party risk management is one key area organizations may benefit from paying more attention to.

Detecting threats is about separating the signal of danger from the noise of the background environment. The analogy for identity and access management infrastructure is ensuring that not many unnecessary items are lying around. That means paying a lot more attention to identity hygiene. Excess entitlements must not be lying around, no orphan accounts must exist, and policies must be tidied up more carefully so that they do not leave an environment where a threat can prevail [12].

Another critical part of keeping identity systems in an organization's shape is ensuring that the detection mechanism is finely tuned enough to give sound alerts and doesn't give many false warnings. This is handled in security operation centers. The latest surveys show that 67% of organizations have either already deployed ITDR or are in the process of deploying it, a significant shift in the last two years. Over time, alerts from applications indicating that they may have been misused will start evolving. There is a huge opportunity to automate the response to alerts for security and identity [13].

Until 2023, it was rare for this detection and response process to be highly automated. By 2025, a typical organization will be highly automated. There has been tremendous activity and innovation in this space to leverage machine learning, generative AI, and LLMs, working quickly to prioritize, process, and act on alerts from security systems.

In the future, many threats will need to be remediated on the identity side in an identity-first security world. So, from now on, the sunlight of a narrow straight path involves the appropriate use of artificial intelligence with appropriate guardrails supervised by humans and other machines to help better detect threats, prioritize them, and respond to them to help evolve faster than the predators.

VI. Future Outlook of Artificial Intelligence in Identity Management

There is an underlying theme here about the need for evolution and the increasing speed of evolving machine learning and artificial intelligence for identity. One of the big hopes for generative AI is that it'll reduce training needs, enable people to communicate more simply and then naturally with identity and security systems, and allow people to receive information back from them naturally.

Of course, the recommendation is to adopt an identity-first security approach. Work on evolving identity infrastructure to experiment with and incorporate these tools as they become available. Consider architecting identity infrastructure to use a more federated, distributed intelligence approach that provides more agility.

Most importantly, any evolution requires iteration; set expectations with the organization to incorporate these advancements into part of an ongoing identity and access management program that will safely and effectively incorporate machine learning and generative AI in the future.

We are at the point where IAM is becoming too complex and fast-changing. There are too many environments, entitlements, and entities. Identifying all these entities and managing their access is challenging for humans. Over time, some automation in IAM tools has been made, but most of this automation still depends on humans.

To meet the current complexity, we must focus on intelligent automation. This means that systems and users must perform most of the tasks with the help of intelligent intermediary agents.

We already have some machine learning capabilities in our IAM tools, most of them purpose-built. However, generative AI will open the door to a new design aspect that can potentially revolutionize IAM. Generative AI will enable business and technical users to interact with IAM systems using natural language, and that is a significant change once we start being able to talk to our IAM infrastructure.

The question is how we build intelligent automation as

part of our IAM. This is so fundamental that generative AI will become the bread and butter of IAM. The central concept is that identity and access intelligence combine generative AI and machine learning capabilities.

Identity and access intelligence capabilities help build intelligent assistant services. These services can help further automate identity and access management capabilities, providing new experiences and journeys for all users.

Generative AI is a capability that enables systems to learn from many representations of artifacts of certain types to generate similar types of artifacts. This can be for text, images, audio, video, or other artifacts.

This is done by foundational models trained on many unlabeled data. We can customize and fine-tune these foundational models for specific use cases like identity and access management. These large language models, or LLMs, are a type of foundational model trained on large amounts of text. They can produce human-like text.

When applying GenAI to any use case or application, we need to know what it is and how it works, and it must be much more profound. We need to understand how to productize generative AI. Should we build it? Should we buy it? We must also understand how to train generative AI on enterprise data and context. Finally, we must understand what assurance control is needed to ensure these things work as promised. We need to understand what interfaces and integrations are available and examine other aspects of these solutions. For example, how we deploy generative AI. Will it be embedded in different applications, and we use it as another feature? Or do we need to work with a standalone tool parallel to other tools, capabilities, and anything in between? Also, we need to understand the attack vector and how it changes with the introduction of generative AI [14].

Once we have generative AI in our environments, we will encounter new types of attacks, such as query attacks, model manipulation, prompt injection, or data poisoning for models.

There are so many attacks out there. The MITRE Atlas has excellent references to all the different types of attacks that can happen in machine learning or AI systems. Each of these issues has subtleties that need attention [15].

Generative AI can enable more effective new attacks on IAM systems. Unfortunately, there will be a lot of fake identities and fake data that can wreak havoc in the IAM system. For example, an administrator may be sitting somewhere and getting a ticket that says, go and wipe out that server; we don't need it, and the executive vice president signs it. The administrators may do that, not knowing that the vice president's signature is mimicked by a rogue agent enabled by generative AI, learning from the behavior of interactions of that executive vice president to impersonate them. So, the challenge is that this type of attack happens by damaging the trust foundation on which we build IAM controls.

When we cannot trust our trust foundation, that can cause some problems.

One is that we must rely on something other than our identity assurance model. Without that, our identity-first security model goes. In this case, there are two paths one can follow. One is full of promises to ensure our IAM controls are consistent, contextual, and continuous, require less manual effort, have better coverage, and reduce error. The other path is full of challenges.

We must ensure our model is trustworthy, address all potential ethical issues to ensure neutrality and deal with potential reality alteration.

We also need to deal with privacy threats and defend against new security and adversarial attacks at a scale. If malware can deploy one agent, it can multiply that by many more.

The scale of the attacks may also increase. Moreover, on the cost front, the generative AI features will cost more because they require more computing power than the classic features. From a suitability perspective, generative AI-enabled features are more suitable for probabilistic administrative tasks and configuration management. Finally, we need lots of data to train and fine-tune our models for different capabilities.

We need to rethink our assurance model and develop a set of safety controls for our generative AI-enabled capabilities to ensure that everything works as intended safely and responsibly. So, the first call to action is that this is an excellent time to start planning how we want to adopt generative AI as part of our identity and access management solutions.

There is so much investment and hype around it, and it has the potential to evolve quickly and change the threat landscape for identity and access management and how we implement and run our IAM controls.

Some use cases are more general, and some can be technical. In the general use cases, simple questions and answers can be related to system search, support, and different queries.

The assumption is that IAM controls leverage generative AI. The models are trained on enterprise data and enterprise context. Thus, we need to fine-tune models with some identity and access management-related data and train them on the context, for example, organization structure, operating model, all the rules and regulations, and policies and procedures. GenAI reads all the text of all rules, potentially comparing and contrasting that with policies and procedures and identifying areas where gaps exist. These models also need to be aware of business processes and applications that automate those business processes.[16]

These are not just point-in-time data. They can be trained on historical data on how these business processes have operated over the years. Thus, it can learn patterns of changes and behaviors.

This is interesting in one way because the system learns about how things are done, but it can also be dangerous.

Nevertheless, a copilot, assistant, or companion can improve access modeling and administration with some complex queries. Access request, approval, and certification can be elevated by informing the users what type of access they can request, telling the managers what access they can approve, and supporting them with the reporting and auditing.

Audit scenarios can be simulated. An agent acting as an auditor looks at everything in the environment and identifies potential audit findings before the audit does.

All these script languages and codes are text-based, so these generative AI solutions can quickly operate on them.

Agents can be trained on commands, API references, and some of these languages' fundamentals. In that case, they can be very effective by giving developers and administrators a starting point to write up different types of

code. For example, fine-tuning an IAM copilot for code generation can improve how user profiles are configured and orchestrated. The customization process can be enhanced by extending and integrating the system with adjacent IAM systems, with their adjacent systems, upstream and downstream, and by reconciling and synchronizing data [16].

Another call to action is to consider establishing the foundation of identity and access intelligence. That happens by focusing on a practical project as proof of concepts, where there is a reasonable amount of data to spin up some of these capabilities. At least become familiar with it, understand how it works, what the potential issues are, and, more importantly, create a list of requirements.

Adopting GenAI starts by re-energizing identity and access data engineering and establishing some identity data platform. This could be a data lake or any alternative. Still, a place is needed to bring the data, consolidate it there, and establish some data engineering lifecycle processes that include the steps to defining the use cases for data, defining the requirements and architecture, designing pipelines, making sure the pipeline has the proper access controls in place as data will move through it to make sure confidential data are not leaked. We need to determine how we orchestrate data and ensure its observability. Then, we need to engage with the AI team in the organization to adopt a GenAI model for our identity and access intelligence.

Usually, we start with the foundational model and customize it for the identity and access management environment. Then, we fine-tune that model with the data we have in our identity data platform. As part of that, we must have a set of safety controls to ensure privacy.

We must comply with all the privacy rules and avoid liability by not leaking any intellectual property in our model. On the output side, we need another set of safety controls to ensure the answers that come out of these systems are correct, and we also need explainability.

We can ground those answers to some accurate references. It is not just the hallucination of these models. Then, we have some other AI cases, like reinforcement learning with the human factor.

Finally, we need to evaluate options for orchestration. Orchestration is when we apply these models to a specific application.

To make an IAM copilot, we must ensure that these models get an enriched prompt that gives a meaningful answer. The system must be trained on enterprise data and context. As most of these models are public, we must do more than ship our data to OpenAI, Google, or all the different models there; we must then fine-tune these models.

Techniques like retrieval and augmented generation exist. That layer is usually made by a vector database, which operates and augments the actual generative AI model.

When the user sends a prompt to an automation agent or intelligent agent, it is enriched with the context that we store in the vector database, similar to long-term memory for the AI system. Once that is done, the prompt is executed, generating the answer. This is followed by a back and-forth with some of the IAM tools to fill some of the answers' gaps. Finally, the answer is sent back to the user. This whole process is called orchestration. Several tools are available to help with orchestration, like LangChain or Microsoft Semantic Kernel.

VII. Conclusion

In conclusion, the integration of Generative AI into IAM presents both significant opportunities and challenges. On the one hand, Generative AI democratizes the use of Large Language Models (LLMs), enabling various IAM use cases such as adaptive access and identity threat detection and response1. This technology can enhance security and operational usability by automating processes and providing more efficient user experiences. On the other hand, it also introduces new risks, such as creating deepfakes and sophisticated phishing attacks. Organizations must carefully select vendors, maintain language models, and ensure privacy and security to navigate these challenges. The future of IAM lies in intelligent automation, where AI-driven tools can self-configure and adapt to evolving threats. Continuous learning and adaptation will be crucial to keep pace with the rapidly changing landscape of AI and IAM.

Acknowledgment

I would like to express my gratitude to Professors Tasho Tashev, Milena Lazarova and Ivo Draganov for their support during my PhD journey.

References

[1] S. Mohamadi, G. Mujtaba, N. Le, G. Doretto, D. Adjeroh, "ChatGPT in the age of generative AI and large language models: a concise survey," *arXiv preprint arXiv:2307.04251*, 2023, https://doi.org/10.48550/arXiv.2307.04251.

[2] Identity and Access Intelligence Innovation with Generative AI, Gartner, 2024, https://www.gartner.com/document-reader/document/4625999?ref=solrAll&refval=437208791 [Online].

[3] M. Mallick, R. Nath, "Navigating the cyber security landscape: a comprehensive review of cyber-attacks, emerging trends, and recent developments," *World Scientific News*, vol. *190*, no. 1, pp. *1–69*, 2024, doi: 10.4236/jis.2024.153019.

[4] M. Virvou, "Artificial intelligence and user experience in reciprocity: contributions and state of the art," *Intelligent Decision Technologies*, vol. *17*, no. *1*, pp. *73–125*, 2023, doi: 10.3233/IDT-230092.

[5] G. Zhao, P. Zhang, Y. Shen, X. Jiang, "Passive user authentication utilizing behavioral biometrics for IIoT systems," *IEEE Internet of Things Journal*, vol. *9*, no. *14*, pp. *12783–12798*, 2021, doi: 10.1109/JIOT.2021.3138454.

[6] D. Saini, K. Kumar, P. Gupta, "Security issues in IoT and cloud computing service models with suggested solutions," *Security and Communication Networks*, vol. *1*, 2022, https://doi.org/10.1155/2022/4943225.

[7] M. Rusia, D. Singh, "A comprehensive survey on techniques to handle face identity threats: challenges and opportunities," *Multimedia Tools and Applications*, vol. *82*, no. *2*, pp. *1669–1748*, 2023, https://doi.org/10.1007/s11042-022-13248-6.

[8] N. Tendikov, L. Rzayeva, B. Saoud, I. Shayea, M. Azmi, A. Myrzatay, M. Alnakhli, "Security information event management data acquisition and analysis methods with machine learning principles, " *Results in Engineering*, vol. *22*, 2024, https://doi.org/10.1016/j.rineng.2024.102254.

[9] W. Hilal, S. Gadsden, J. Yawney, "Financial fraud: a review of anomaly detection techniques and recent advances," *Expert Systems with Applications*, vol. *193*, 2022, https://doi.org/10.1016/j.eswa.2021.116429.

[10] A. Paleyes, R. Urma, N. Lawrence, "Challenges in deploying machine learning: a survey of case studies," *ACM Computing Surveys*, vol. *55*, no. *6*, pp. *1–29*, 2022, https://doi.org/10.1145/3533378.

[11] V. Chamola, et. al, "Beyond reality: the pivotal role of generative AI in the metaverse," *IEEE Internet of Things Magazine*, vol. *7*, no. *4*, pp. *126–135*, 2024, doi: 10.1109/IOTM.001.2300174.

[12] B. Gunes, G. Kayisoglu, P. Bolat, "Cyber security risk assessment for seaports: A case study of a container port," *Computers & Security*, vol. *103*, 2021, https://doi.org/10.1016/j.cose.2021.102196.

[13] E. Egho-Promise, M. Sitti, "Big data security management in digital environment," *American Journal of Multidisciplinary Research & Development (AJMRD)*, vol. *6*, no. *2*, pp. 1–34, 2024.

[14] M. Wu-Gehbauer, C. Rosenkranz, "Unlocking the potential of generative artificial intelligence: a case study in software development," *Proc. of International Conference on Information Systems (ICIS 2024)*, 2024, https://aisel.aisnet.org/icis2024/aiinbus/aiinbus/25.

[15] M. Takaffoli, S. Li, V. Mäkelä, "Generative AI in user experience design and research: how do UX practitioners, teams, and companies use genai in industry?," *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS'24)*, Association for Computing Machinery, New York, NY, USA, pp. 1579–1593, 2024, https://doi.org/10.1145/3643834.3660720.

[16] Z. Xi, et. al, The rise and potential of large language model based agents: a survey," *arXiv preprint arXiv:2309.07864*, 2023, https://doi.org/10.48550/arXiv.2309.07864.