Vendor Profile

Credo AI: An AI Governance and Risk Management Platform Opens New Frontiers in AI Revolution

Ritu Jyoti

IDC OPINION

The fast-emerging field of artificial intelligence (AI) is everywhere and is starting to improve our lives in myriad ways, from simplifying our shopping to enhancing our healthcare experiences. Their value to businesses also has become undeniable. According to IDC's Worldwide Semiannual Artificial Intelligence Systems Spending Guide, February 2022 (V1) – which tracks artificial intelligence software, hardware, and services across industries and use cases – enterprises worldwide are expected to invest $112.9 billion on AI solutions in 2022. This spending is expected to grow to $221.8 billion at a compound annual growth rate (CAGR) of 26.2% for 2020-2025. Furthermore:

- Even as AI generates value, it is giving rise to a host of unwanted, and sometimes serious, concerns. Artificial intelligence is often perceived as a black box technology, with a lingering fear of unintended negative consequences including the ones that are not yet known or experienced. Disastrous repercussions – including the loss of human life, if an AI medical algorithm goes wrong, or the compromise of national security, if an adversary feeds disinformation to a military AI system – are possible.
- Organizations are confronted with potential negative business impact (e.g., damage to brand reputation, reduced public trust, revenue loss, criminal investigations, regulatory backlash, customer privacy loss, hidden costs) if the AI/machine learning (ML) business risks (e.g., data privacy, discrimination, black box, noncompliance, and safety) are not mitigated appropriately.
- Trustworthiness is emerging as a dominant prerequisite for AI, and businesses must take a proactive stance.
- Responsible implementation of AI is up to us. Enterprises need to have a holistic governance framework to safeguard themselves across the AI life cycle from design and development to production and ongoing. According to IDC's *Future Enterprise Resiliency and Spending (FERS) Survey – Wave 8* conducted in September 2021, at the worldwide level, organization-wide assurance efforts are rapidly maturing. Increasing government actions related to fair use of AI will drive more investment in Responsible AI software and governance. According to the same survey, at the worldwide level, enterprises whose IT spending will be higher than originally budgeted have higher assurance maturity. Overall, IDC research highlights that enterprises are looking for continuous AI governance to scale AI, compliant AI/ML deployments, and access to compliant data. However, they are unable to do so for the following reasons: lack of gold standard and policies, inability to audit for compliant AI, large gaps in AI expertise, lack of tools for continuous governance, and misalignment of incentives.
- IDC believes that Credo AI's governance and risk management platform opens new frontiers in AI revolution, as it addresses these pain points/gaps by providing a foundational offering to support an enterprise with holistic AI governance and risk management. It is an open and multi-stakeholder collaboration platform that helps enterprises operationalize Responsible AI at scale.

March 2022, IDC #US48903722

IN THIS VENDOR PROFILE

This IDC Vendor Profile examines and reviews offerings from Credo AI, a private enterprise software company, and how Credo AI's Responsible AI governance platform empowers businesses worldwide to build and embrace AI responsibly – aligned with human values and putting humans first.

SITUATION OVERVIEW

Company Overview

Credo AI is a venture-backed start-up founded in 2020 with the mission to help organizations build AI aligned with human values. The company is headquartered in Palo Alto, California. The company was founded with $5.5 million in seed funding led by Decibel Partners, along with Village Global and Andrew Ng's AI Fund, by founder and CEO Navrina Singh and her cofounder and CTO Eli Chen.

Navrina Singh is an ex-Microsoft and Qualcomm employee with deep understanding of the market needs. With expertise in AI/ML product development, her faith in AI governance was bolstered when she saw the AI regulations and policy initiatives expand globally as a young global leader with the World Economic Forum. Credo AI's cofounder and CTO Eli Chen has deep expertise in security, compliance, neural networks, and large-scale distributed systems with his ex-employers Twitter, Netflix, and other start-ups.

The executive team is rapidly expanding with industry experts. Susannah Shattuck, head of product, brings in extensive experience in AI Governance and monitoring as an ex-Arthur AI, Google, IBM, and Stanford employee. Kyle Ledbetter, VP of Design, brings leadership in trusted user experiences and big data systems from his time at Teradata, eBay, and MicroStrategy.

Although it's early days, the company is gaining traction for its Responsible AI governance SaaS platform. According to Credo AI, the company is on track to multimillion dollars in revenue by the end of 2022. Credo AI's customers are primarily among the top 2000 global companies in finance, retail, HR, government, and technology. According to Credo AI, a Fortune 50 financial services firm is one of its biggest customers, as is one of the largest cloud providers, along with a leading defense contractor.

Company Strategy

Product Strategy

Credo AI positions itself as empowering organizations to create AI with the highest ethical standards by allowing business and technical stakeholders to measure, manage, and monitor AI-introduced risks across data, models, and processes to ensure responsible, auditable, and compliant AI at scale.

Credo AI would like to make the world a better place with Responsible AI deployments. Credo AI's vision is to build a generational company, where Trust with AI is existential. However, the company is starting with "AI Risk Management and Governance" where businesses can *align* on the right guardrails based on business and regulatory context; *assess* the risks and compliance of their AI models, data, and processes as they embark on a Responsible AI journey; and *report* on how they are bringing oversight and accountability to the AI life cycle. This is followed by Trust as a Competitive Advantage, where Credo AI becomes a trusted sherpa guiding enterprises and regulators to comply with global regulatory requirements. Then comes Trust as Existential where Credo AI leverages the

vendor-enterprise-customer-regulator ecosystem to expand across verticals and other frontier technologies to become the seal of trusted technology.

Credo AI is a Responsible AI platform that serves both business and technical stakeholders and provides real-time, context-driven, comprehensive, and continuous governance and risk management of AI. Credo AI is unlike machine learning operations (MLOps) platforms that focus on end-to-end model management, from data collection to operationalization for only technical stakeholders in data science, ML, and product. Credo AI is not an ML monitoring tool to monitor data and/or concept drift either. However, it does integrate into an organization's existing machine learning operations and monitoring tools to extract evidence against risk controls. Credo AI takes the signals and conducts the overall risk analysis and provides the translations needed to provide a holistic risk and compliance view of everything that is happening in production. Accordingly, when viewed as a component of an enterprise data architecture, Credo AI complements the adjacent market segments of MLOps and governance, risk, and compliance (GRC).

Credo AI serves the buying decision-making C-suite executives (e.g., chief data officer, chief AI officer, chief compliance officer, and chief risk officer), influencers/users (e.g., product managers, data scientists), and oversight professionals (e.g., compliance manager/associate, internal/external auditors, and data architects tasked with security focus).

Credo AI is purpose-built for industries ranging from finance and banking, retail, insurance, HR, and talent management to government. The product supports a broad set of use cases like fraud detection, risk scoring, underwriting, employment decision tools, marketing, facial recognition, and speech recognition.

Credo AI seamlessly fits into an enterprise's existing ML (AI platforms, data management, MLOps/Model Ops tools, ML monitoring tools and infrastructure) and risk management (GRC, MRM, or other enterprise risk management) foundation. Credo AI offers flexible deployment options, either as a standalone full suite deployment or as the Responsible AI engine that integrates into an enterprise's existing AI and risk management stack. Credo AI plans to offer the choice between SaaS or on premises.

Product/Service Offerings

Credo AI is a SaaS platform that enables end-to-end Responsible AI governance. It is a multi-stakeholder collaboration platform and single source of truth for all the model audit artifacts and governance needs, bringing the business and technical stakeholders together. Credo AI is preparing to offer an on-premises option of its solution in 2022.

Credo AI Responsible AI platform brings contextual risk management and governance to data and AI life cycle. As previously mentioned, the platform is focused on bringing continuous and real-time oversight and accountability by helping stakeholders:

- **Align** on the right guardrails based on business and regulatory context.
- **Assess** the risks and compliance of their AI models, data, and processes.
- **Report** on business impact through the risk and compliance outcomes.

Figure 1 shows Credo AI's Responsible AI governance platform.

FIGURE 1

Credo AI: Responsible AI Governance Platform

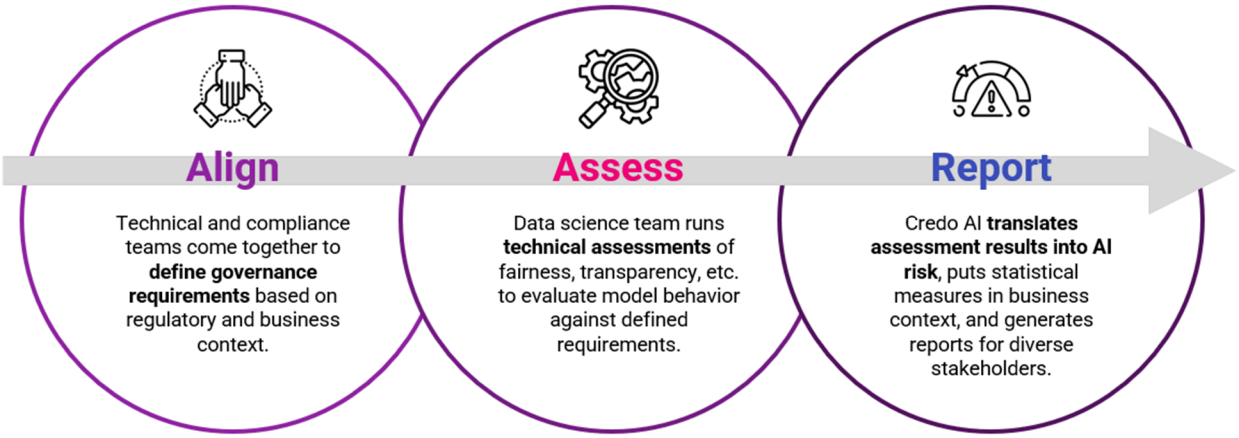

Source: Credo AI, 2022

The first step in using Credo AI starts with registering an AI use case for governance in the Credo AI Use Case Registry. An AI use case can be made up of any number of machine learning models, but it's really this use case-driven governance that Credo AI encourages organizations to start with. Once a use case is registered for governance, the organization is led through a scoping process where businesses determine the contexts that the system is going to be operating in and then they select governance requirements based on that context. Scoping gets organizations to define structured metadata about the use case, which includes things like deployment region, whether this use case is going to impact people, and whether it's related to any particularly high-risk areas as defined by emerging regulations like the EU AI Act and the Algorithmic Accountability Act.

Each AI use case has an AI use case card, which is an overview of the current risk level of the AI solution based on the current state of governance. This is based on the evidence that has been pulled in by data scientists, product managers, and the compliance team that's working together on the solution.

Once the context of the AI solution has been defined and businesses have registered the ML models that are going to be part of the system, either through integrations or manually, then the businesses apply the governance requirements, which in Credo AI platform are called "policy packs."

Policy packs are modular sets of policies and controls that are designed to address objective technical risks and process-based compliance. For instance, a policy pack designed to help organizations develop fair AI systems includes a required technical fairness assessment of the models and data sets, along with specific actions that the development team should take during the development life cycle. For technical assessments, like fairness or performance, an additional layer of alignment and definition is required. Credo AI offers a valuable tool for product managers, data scientists, and compliance roles to come together to discuss, assess, and define their measurement metrics.

The Credo AI platform includes a broad set of policy packs based on not just the company's understanding of industry best practices but also existing and emerging regulatory frameworks and

standards from organizations like the IEEE and NIST. Credo AI can also help enterprise customers create modular sets of governance policies so that an organization can codify its values into governance requirements. For example, if a customer has an AI use case that's being deployed in the EU, it can pull in the EU AI Act Readiness Pack, and if the customer has a system that's deployed in financial services, then the customer could use the policy packs informed by organizations like the Federal Reserve.

Once the right guardrails have been put in place for an AI use case, Credo AI enables enterprises to check model compliance against enterprise policies and regulations with its assessment framework, Credo AI Lens. Credo AI Lens is a Python toolkit that provides standardized model and data set assessment capabilities for fairness and performance, with transparency and security assessments coming soon. With Credo AI Lens, organizations can integrate model assessments into their development and deployment pipelines to allow for continuous risk management. Credo AI Lens automatically sends assessments results back to the Credo AI platform, reducing the burden of assessment and reporting on technical teams and data scientists. Credo AI's governance dashboard summarizes key compliance information for Models and AI use cases in development and production across an organization (see Figure 2).

FIGURE 2

Credo AI: Dashboards

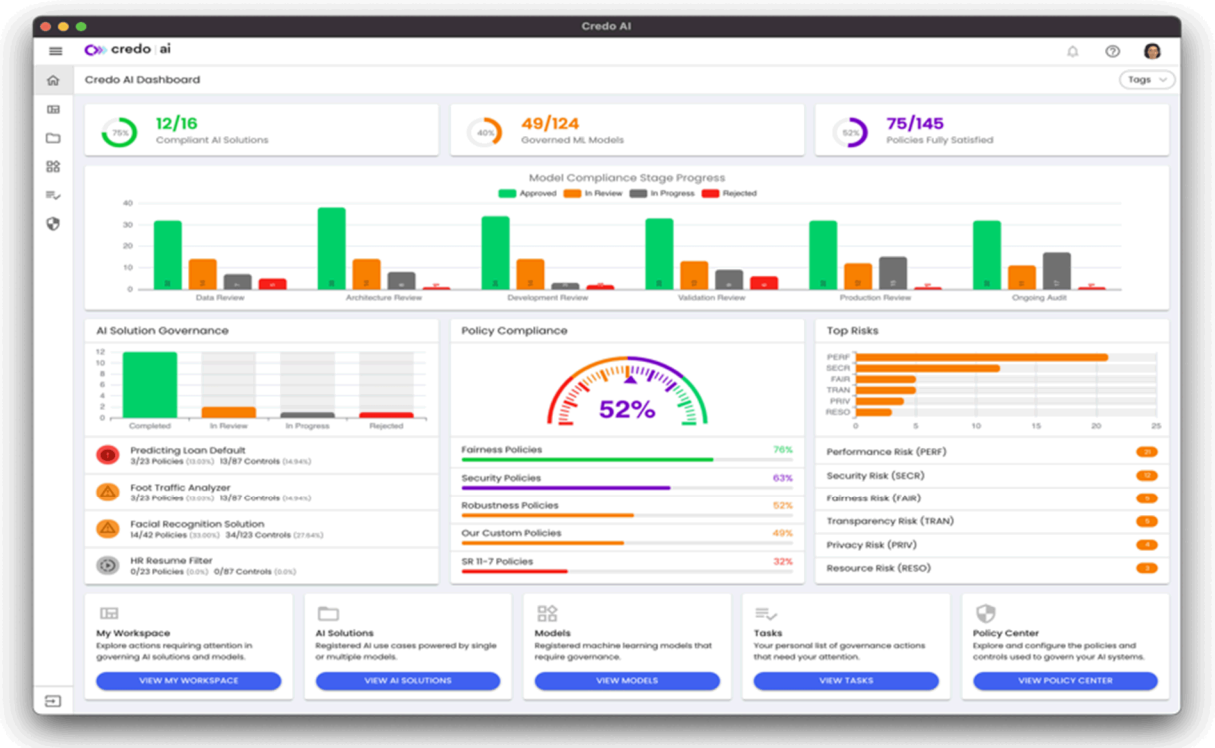

Source: Credo AI, 2022

The Credo AI platform provides a variety of reporting capabilities to share assessment results with internal and external stakeholders (see Figure 3). The platform automatically generates risk and compliance reports, along with Model Cards and an audit trail of governance actions and decisions taken during the development life cycle.

FIGURE 3

Credo AI: Reporting

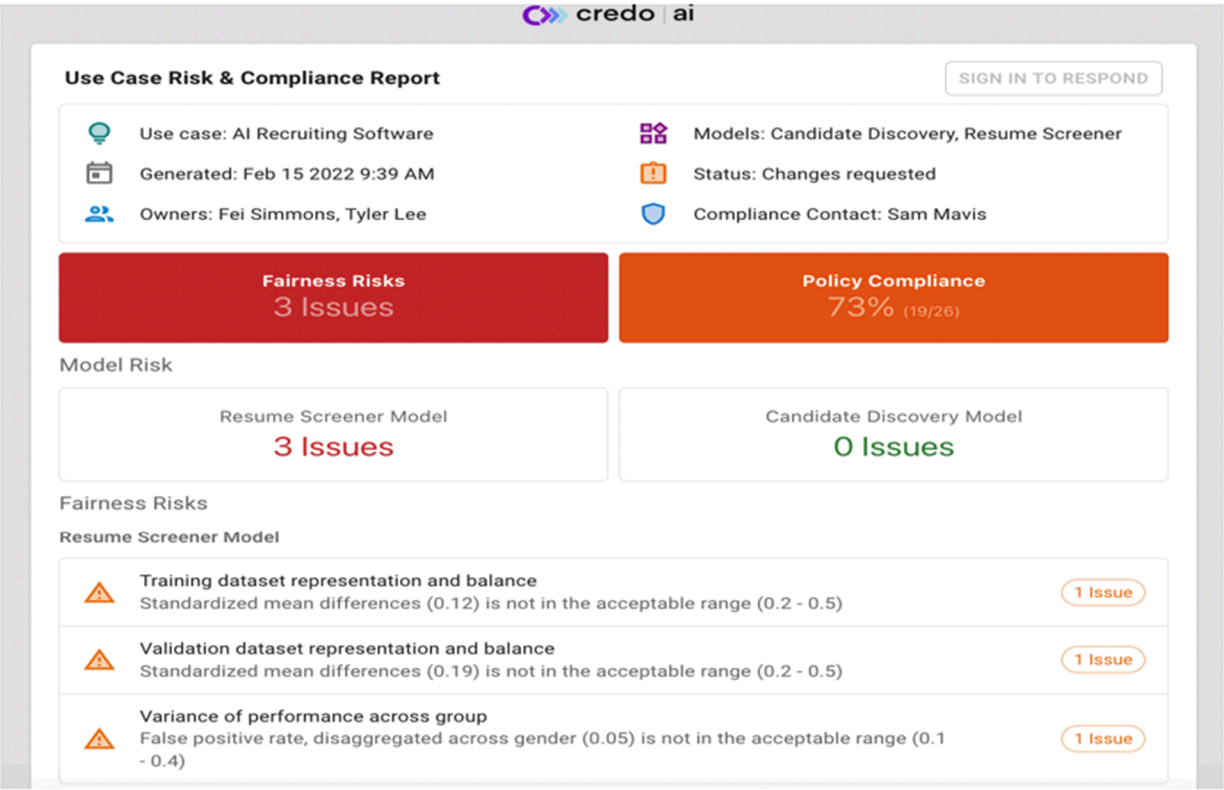

Source: Credo AI, 2022

Overall, Credo offers end-to-end governance workflow for both ML models and AI solutions and out-of-the-box policies and controls for AI governance. It supports centralized system of record for evidence collection at every stage of the AI development life-cycle features and collaboration features to enable multi-stakeholder governance.

Business Strategy

The Responsible AI market is still in its infancy and Credo AI's business strategy has been focused on businesses with AI expertise and large-scale deployments. Credo AI's target market consists of Global 2000 enterprises, including verticals (e.g., finance and banking, retail, insurance, HR and talent management, high tech, and government). It is currently offered as a SaaS offering and will be available as an on-premises solution in 2022.

Credo AI has a balanced go-to-market strategy, which includes a combination of an inbound funnel of direct client prospects and the leverage of a network of partners to jointly accelerate innovation. With the growing momentum in the Responsible AI ecosystem, Credo AI has an opportunity to extend its channel partnerships in 2022. Ease of deployment and time to value (governance) are core to Credo AI's approach to scale business' ROI.

According to Credo AI, it offers sandbox/POC environments to help businesses quickly experiment, gain confidence, and accelerate deployment. Customers are provided 24 x 7 support through web and phone. Credo AI offers flexible pricing options.

Early Adopter — Customer Scenarios

- **A Global 2000 publicly traded financial services corporation** is using AI for a broad set of use cases — with fraud detection and risk scoring being the top ones. Managing AI risk and providing oversight to ML models was a manual process and required expensive (~$50 million annually) investment in staff and professional services. By using Credo AI's Responsible AI platform, the corporation's oversight professionals and technical stakeholders can collaborate effectively to govern and manage AI-introduced risks. With Credo AI, they have a single source of truth and audit artifacts for all their AI/ML governance requirements, which can be shared effectively internally and externally. They have been able to deploy critical practices and policies to manage AI risks and are regulation ready. With Credo AI, they can critically assess the quality of their ML for fairness (and other risk areas) to gain confidence in their AI deployments. Through use of Credo AI, they also have reduced the burden of governance activities on data science teams, who are now able to run and report on fairness and performance assessments in minutes instead of days. Overall, using Credo AI has helped them increase customer trust, thereby unlocking more sales. Centralized governance and effective management of risks has resulted in cost savings and has increased ROI on AI enabling top-line growth.
- **A Fortune 500 global cloud services provider's** top AI use cases are facial recognition and speech recognition. The provider needs to show proof of good governance to its customers to gain consumer trust, unlock more sales opportunities, and accelerate procurement. Consumers (nontechnical) are demanding visibility into the building of AI with appropriate disclosures, which are resulting in massive brand and reputation risks that may cost millions of dollars as well as forcing a long procurement process (9–10 months) of their AI technologies. Use of Credo AI's Responsible AI platform has enabled faster procurement of their products by customers, along with unlocking more sales with increased customer trust and has resulted in increased ROI on AI enabling top-line growth.
- **A large defense contractor** needed help evaluating third-party AI risk and managing governance of vendor AI applications. The contractor needed to understand the risk of third-party commercial AI vendors to its organization and align them with new government frameworks (e.g., Defense Innovation Unit Responsible AI, NIST, DoD Responsible AI, GAO framework). By using Credo AI's Responsible AI platform, it is critically assessing the quality of third-party commercial AI vendor models and attain confidence in their deployment across the organization. The contractor has been able to make informed, risk-based procurement of AI-based systems, aligned with its frameworks and standards and thereby increase ROI on AI enabling top-line growth.

FUTURE OUTLOOK

When something goes wrong with AI, and the root cause of the problem comes to light, there is often a great deal of headshaking. With the benefit of hindsight, it seems unimaginable that no one saw it coming. Enterprises hoping to shift their posture from hindsight to foresight need to better understand the types of risks they are taking on, their interdependencies, and their underlying causes. They can embrace risk management core principles to reduce their exposure to AI/ML risks and potential negative impact:

- Enterprises need to use a structured identification approach to pinpoint the most critical risks by assembling leaders from business, IT, security, and risk management to evaluate and prioritize their greatest risks. Inputs to this exercise could include a clear-eyed look at the company's existing risks and how they might be exacerbated by AI-driven efforts under consideration and at new risks that AI enablers, or the AI itself, could create.
- Enterprises need to institute companywide controls to guide the development and use of AI systems, ensure proper oversight, and put into place strong policies, procedures, worker training, and contingency plans.
- While enterprisewide controls are important, they are rarely sufficient to counteract every possible risk. Another level of rigor and nuance is often needed, and the requisite controls will depend on factors such as the complexity of the algorithms, their data requirements, the nature of human-to-machine (or machine-to-machine) interaction, the potential for exploitation by bad actors, and the extent to which AI is embedded into a business process. Conceptual controls, starting with a use case charter, sometimes are necessary. So are specific data and analytics controls, including transparency requirements, as well as controls for feedback and monitoring, such as performance analysis to detect degradation or bias.

Overall, from the perspective of AI/ML risks and challenges, it is easy to see that there is a need to address these pain points from a more ethical point of view. What is right and what is wrong? Independent of a business' answer to this question, the decision process should include not only information and knowledge but also, to some extent, their values and preferences. AI should reflect the company's values. This means that the answer might not be the same for two businesses, yet AI-enabled decision support systems might be used globally across the world — adding to the complexity of judging the correctness of the outcome from a value perspective. How do we design Responsible AI applications that are truly global?

The reason that ethics is so important is that now we have machine intelligence that sits between us and the organizations that we are dealing with. AI algorithms aren't neutral and are not programmed like traditional software. Rather, AI is trained on historical data, which means that AI may learn patterns of behaviors that might be surprising or unintentional.

Lawyers, activists, and researchers emphasize the need for ethics, responsibility, and accountability in the design and implementation of AI systems. But this often ignores a couple of tricky questions: Who gets to define those ethics, and who should enforce them?

Philip Alston, an international legal scholar at NYU's School of Law, proposes a solution to the ambiguous and unaccountable nature of ethics: reframing AI-driven consequences in terms of human rights. "[Human rights are] in the constitution," Alston said at an AI symposium. "They're in the bill of rights; they've been interpreted by courts," he said. If an AI system takes away people's basic rights, then it should not be acceptable, he said.

Other pioneers believe that using AI responsibly comes down to a combination of things. Algorithms produced by different companies must be constantly benchmarked and refined so that they are as accurate as possible. There should be clarity on how the usage of algorithms are recommended to the end users.

For example, if an NLP model is used to provide superior document search, it may be acceptable to have a confidence level or threshold that is around 80%. If that same NLP model, however, is used to analyze and filter resumes from job candidates, the confidence threshold target should be much higher, and even then, it shouldn't be the sole determinant in the decision making. There should be a human involved who reviews the model scoring in addition to other sources of information to make a holistic decision about a candidate.

Individual nations may have to decide what standards, regulations, or guidance they are going to give the companies that use these types of technologies. What is the moral responsibility of a data team today? As artificial intelligence and machine learning technologies become part of our everyday life and as data and big data insights become accessible to everyone, CDOs and data teams are taking on a very important moral role as the conscience of the corporation.

Every enterprise needs to define and articulate its Responsible AI mission and principles while also establishing a transparent governance structure across the organization that builds confidence and trust in AI technologies. The enterprise needs to strengthen compliance with current laws and regulations while monitoring future ones, develop policies to mitigate risk, and operationalize those policies through a risk management framework with regular reporting and monitoring. While it is early days, extending their existing risk quantification methodology can help organizations identify and prioritize AI/ML risks to inform decisions of deploying AI/ML

ESSENTIAL GUIDANCE

Advice for Credo AI

At IDC, when we look at the current state of AI adoption, drivers, and inhibitors, we believe Credo AI can benefit from the following:

- Understanding that an enterprise AI landscape will include multiple vendor models and applications, Credo AI should look to provide third-party model governance.
- While the current set of policy packs supported by Credo AI is impressive, Credo should look to support policy pack versioning with compliance alerts and update flows. This should help enterprises keep up with changing regulations and ensure compliance is up to date with policy pack update alerts and tools.
- Even though the current MLOps and monitoring tools market is quite fragmented. Credo AI should explore either integration with a couple of them or support additional tools for data science teams to manage and support performance, explainability, and robustness assessments run.
- As more and more business applications are embedding AI/ML capabilities, Credo AI should explore collaboration with leading ERM and CRM technology suppliers to help them not only use the platform during their development life cycle but also once it is deployed in an enterprise environment to support ongoing governance and risk management.

- Last, Credo AI should collaborate with their early adopters to drive awareness on the best practices for AI governance and risk management. Credo should continue to educate the enterprises through deep dive end-user webinars and Q&A sessions.

LEARN MORE

Related Research

- *IDC FutureScape: Worldwide Artificial Intelligence and Automation 2022 Predictions* (IDC #US48298421, October 2021)
- *Market Analysis Perspective: Worldwide Artificial Intelligence Software, 2021* (IDC #US48243221, September 2021)
- *Manage AI/ML Business Risks and Thrive with Trustworthy AI* (IDC #US48235521, September 2021)
- *Worldwide Artificial Intelligence Software Forecast, 2021-2025* (IDC #US48125621, August 2021)
- *AI StrategiesView 2021 Premium: Banner Tables* (IDC #US47638621, April 2021)

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com