



Enterprise AI for the Enterprise Minded

How Small Data Teams Can Harness the Power of AI





AI technology has never had more buzz or been more accessible than right now. Which is exciting, or intimidating, or tantalizing, or terrifying... or maybe all the above.

Where you stand on the prospect of AI in your organization will depend on many factors, including how far along you are on the Data Maturity Curve. But don't let the size of your organization impact your interest in advanced data applications – Enterprise AI isn't only for the large enterprise.

In this article, we'll explore how smaller, enterprise-minded organizations can approach AI and ML adoption and answer some frequently asked questions about AI and ML. You don't need a huge data team, and you don't need to double your IT budget to implement use cases that can present real value to your organization. With the right approach, AI and ML technologies can become an integral part of your organization's data strategy.

AI and ML FAQs

What is AI, ML and GenAI?

To start, let's define AI, ML and GenAI and identify common use cases associated with each technology.

Artificial Intelligence (AI) – The Umbrella

Artificial Intelligence (AI) refers to any technology that enables machines to mimic human cognitive abilities. AI tools are already embedded in nearly every website browser and modern word processing and productivity tool used today. Common use cases for AI include customer service chatbots that process customer queries and provide responses and fraud detection tools that collate transaction data, user behavior patterns and historic fraud cases to identify suspicious activities in real-time.

Machine Learning (ML) – The Algorithm

Machine Learning (ML) is a subset of AI that trains algorithms to learn from experience, making decisions or predictions based on data. ML has been in use for decades, but what was once solely the domain of data scientists and mathematics PhDs. New technology advancements have opened ML to a wider audience of data practitioners and business line users, with common applications including predictive maintenance using sensor data, historical maintenance records and usage logs, and next-best product recommendations based on historic user preferences, purchase history, and item characteristics.

You don't need a huge data team, and you don't need to double your IT budget to implement use cases that can present real value to your organization.

Generative Artificial Intelligence (GenAI) – The Brain

Generative AI supercharges traditional AI with deep learning capabilities that allow it to create new content based on existing media or text. While traditional AI focuses on performing tasks that require human intelligence, GenAI is dedicated to learning from inputs to generate original content. GenAI can speed up software development by generating code snippets based on high-level requirements or provide a more human-like chatbot experience by learning from conversational data, customer interactions and task automation scripts.

How much does it cost to get started with AI, ML, or GenAI?

Assuming you've built the foundation of a centralized version of truth for your data, you can get started on your first AI use case with a \$25,000 investment. Building a modern data warehouse in Snowflake reduces overhead and operational expenses, as you can do most of the AI within your Data Cloud environment. For smaller data teams, Snowflake has an added benefit of building AI competencies from within, using developer tools your engineers already use. These tools include a Python connector that connects to free open-source ML libraries; the Snowpark development environment, which can be used to write and execute custom code; and

While you consider the upfront investment, consider the opportunity cost of being leapfrogged by new technology. The foundation you build for AI today will help you prepare for the data challenges and opportunities of tomorrow.

Snowflake User Defined Functions (UDFs), which can support simple ML tasks. Certain use cases will require additional tooling, and many AI and ML applications can be found on the Snowflake Marketplace and brought directly to your data.

A good rule of thumb is to expect a \$25,000 – \$50,000 investment in man hours or services for each use case depending on complexity. Should your use case require new software to run on top of your data warehouse, you should expect anywhere from \$50,000 – \$500,000+ on the software, depending on the scale and complexity of the use case.

As with other data modernization initiatives, the cost of building

an AI strategy can often pay for itself in added efficiency and revenue. While you consider the upfront investment, consider the opportunity cost of being leapfrogged by new technology. The foundation you build for AI today will help you prepare for the data challenges and opportunities of tomorrow.

What do I need to prepare for AI, ML or GenAI?

Good news for smaller data teams that are hoping to get started with AI and ML – the same work you are doing to develop a strong data quality and data governance foundation is the best preparation for advanced data applications. These include:

- **Support a single version of the truth with a unified data estate.** We like Snowflake, because it simplifies data management with support of structured and semi-structured data, and integrates with data visualization tools, ETL process and BI software your organization might already have in place. By moving data out of silos into a well-governed data lake warehouse, you can create views of your customer or organization that encompass data from across your ecosystem.
- **Prioritize data quality.** Data estate modernization initiatives can quickly lose value if data quality isn't prioritized. Make sure you can measure

data quality and have a way to resolve discrepancies in data accuracy, completeness and timeliness. AI technologies train on the data you have and can compound the impacts of inaccurate data – whether that is duplicated data, incomplete data, or out of date data, by an order of magnitude.

- **Have a data governance plan that defines data ownership and provides observability into the data lifecycle.** AI technologies should support, not supplant, human brilliance on your team. AI depends on data governance practices that identify and protect sensitive data in your organization, while ensuring that humans are ultimately in control of the final Data Product. Core data governance competencies you should establish to prepare the added regulatory and ethical burdens of AI include establishing Active Data Control to monitor data usage and access in real-time, prioritizing data classification that support tracking data lineage, and assigning Data Stewards, Data Owners, and Data Custodians to guide AI implementation.

Getting Started with AI, ML and GenAI – Special Considerations for Smaller Data Teams

How Small is too Small?

If you have a modern data strategy, your organization's size shouldn't determine whether you adopt AI into your technology stack. Read on for information about building your AI Team and choosing a first use case.

Building an AI Team at Smaller Organizations

Too often, organizations think you need to have a data scientist on hand to deliver results and ROI. Emerging AI and ML tools have democratized the playing field, and you might already have the resources on hand to execute your first AI use case. At Camden National Bank, data leaders have been able to reassign team members by automating tasks and creating efficiencies.

At a minimum, your AI team should include:

- **A Data Steward** – a business-line expert who understands the bottom-line impact of the use case and how the data is used.
- **A Data Owner** – probably someone on your data team, who can delineate the policies and procedures to stay compliant.
- **A Data Custodian** – an engineer or technical team member who can

execute the technical requirements of the application in your data environment.

If you have a functional data governance program in place, these individuals should be easy to identify.

AI and ML Use Cases for Immediate ROI

A use case that aligns with your business goals is more likely to show ROI and encourage adoption across your organization. AI and ML use cases can be applied across your organization, from creating operational efficiency to supporting a better customer experience. Below, you'll find use cases that could bring immediate value to your organization. If you'd like help getting started, Passerelle provides consulting and implementation support for all use cases listed below, along with a complimentary AI-readiness assessment.

Customer Service – Prevent Customer Attrition/Churn

Predictive analytics and anomaly detection algorithms can leverage historical data to anticipate future events, including customer churn. For a churn prevention use case, organizations can use machine learning models to analyze data based on features that indicate a customer may be at risk of leaving. The model assigns a score to each customer, which can be integrated into front-end digital systems and customer-facing platforms to support retention efforts. This enables relationship managers or customer service teams to proactively engage at-risk customers with personalized offers or one-on-one interactions to improve retention.

Operational Efficiency – Automate Workflow Processes

Organizations can use AI to automate repetitive tasks or optimize workflows, resulting in significant ROI—not only through productivity gains but also by allowing team members to focus on work that directly contributes to the bottom line. Common automation tasks include chatbot-supported customer service, but AI can also be used to streamline processes such as application approvals or data verification by rapidly assessing information from various sources, thereby reducing the need for manual checks.

Sales & Marketing – Identify Next Best Product and Drive Customer Loyalty

Many are familiar with next best product recommendations on e-commerce platforms, but most B2C organizations can derive significant value from using AI to enhance personalization and customer experience initiatives. By analyzing rich data on customer interactions, behaviors,

and preferences, organizations can use next best product strategies to upsell customers on new offerings. Additionally, this data can help drive increased service usage and customer loyalty by suggesting actions that align with the customer's needs and well-being.

Risk & Compliance – Fraud Prevention

Organizations can leverage AI to detect patterns or anomalies in data, making it particularly effective for fraud detection and quality control. As cybersecurity risks increase, especially with the rise of AI tools accessible to malicious actors, AI-driven fraud prevention measures could become essential for organizations where security is paramount.

Sales & Marketing – Mine Unstructured Data

Natural Language Processing (NLP) can analyze human language to support use cases such as sentiment analysis, automated document processing, and voice recognition. With NLP, organizations can unlock insights hidden in previously unstructured data sources, such as call center transcripts and PDFs. These insights can be invaluable in identifying trends, customer needs, or previously overlooked opportunities within existing customer bases.

Compliance/Regulatory Efficiency

NLP and Large Language Model (LLM) tools can be employed to mine unstructured data for enhancing vendor risk management, compliance, and regulatory validations within your organization's policies and document repositories. These tools can cross-check specific regulatory wording, identify missing elements, and validate document metadata, such as update dates and responsible individuals, ensuring thorough and proactive governance of regulatory requirements.

Would you like help getting started? We help organizations of any size prepare for the AI Frontier with a complimentary AI Readiness Assessment. Learn more at www.passerelledata.com, or through the QR code.

A use case that aligns with your business goals is more likely to show ROI and encourage adoption across your organization. AI and ML use cases can be applied across your organization, from creating operational efficiency to supporting a better customer experience.

LEARN MORE





About Passerelle

Passerelle connects data to action. Our purpose-driven engineering supports Agile Data Governance and AI-Readiness. We leverage partnerships with leading-edge data technology across the data value chain, spurring adoption through our IP and blueprints, technical expertise, and use-case-based deployment.

In addition to engineering and system integration services, Passerelle is the creator of Data Rocket®, an end-to-end acceleration architecture that modernizes data infrastructure and delivers critical business insights – securely and accessibly. Data Rocket unlocks industry-best data technology for businesses of any size, with a focus on data quality, scalability and advanced data applications.

Passerelle works with companies at any stage of technology adoption, from holistic data estate modernization to cloud and cloud hybrid migration, API services, ETL engineering, custom dashboards, ML and AI integration, predictive analytics solutions, and ongoing consulting and support. We work in data-intensive industries, including financial services, healthcare, public health, energy, manufacturing, retail and higher education.

www.passerelledata.com