

CXOs: How to Successfully Implement AI Initiatives at Your Organization

A Guide to Executing Short-Term Wins *and*
Long-Term Transformation



Introduction

The State of AI Implementation in the Enterprise

The race is on for organizations to transform themselves with AI, and the stakes have never been higher in a changing world that demands constant agility and adaptability from businesses across all industries.

→ **By 2024, 75% of organizations will have established a centralized data and analytics (D&A) center of excellence to support federated D&A initiatives and prevent enterprise failure.**

Source: Gartner Predicts 2021: Data and Analytics Leaders Are Poised for Success but Risk an Uncertain Future, 1 December 2020

However, urgency doesn't equal ease. Companies still feel overwhelmed by the decisions and challenges that stand in the way of implementing an AI roadmap.

For example, many executives feel that they need to choose between empowering analysts or even day-to-day business users and empowering the data scientist group or a centralized center of excellence. Others are paralyzed by the need to build compliance and processes around AI deployment, requiring central control, and the push for more agility at the edges of the organization.

On top of it all, leaders are forced to navigate between buying off-the-shelf AI solutions, which come with the risk of losing their competitive edge and their IP assets, and a build approach, with the risk of costly failure. Given all this uncertainty, what is the path forward?

➔ **By 2025, 50% of large enterprises will have deployed artificial intelligence orchestration platforms to operationalize AI, up from fewer than 10% in 2020.**

Source: Gartner Cool Vendors in Enterprise AI Operationalization and Engineering, 11 October 2021

Successful AI Is Everyday AI

Market trends beyond data and AI — from game platforms to writing and more — show the promise of a creator economy, where a large number of people contribute in different ways to making intellectual goods.

It's easy to see why, when it comes to AI, organizations stand to benefit from the contribution and perspective of various types of users from all parts of the business (whether traditional data roles or not) to operationalize AI initiatives.

In fact, industry analysts identify the enablement of new and varied roles with AI capabilities as one of the top trends for 2022 and beyond. Here at Dataiku, we've seen that companies (including GE, Pfizer, NXP, and more) taking this approach and embracing Everyday AI are the most successful in their transformation efforts.

What Is Everyday AI?

Everyday AI is about widening the scope of AI initiatives to more people across the organization — that means more people both creating and using AI projects.

But beyond that, Everyday AI is about making the use of data almost pedestrian — AI that is so ingrained and intertwined with the workings of the day-to-day that it's just part of the business (not only being used or developed by one central team).



*We've surveyed thousands of executives about how their companies use and organize for AI and advanced analytics, and our data shows that only 8% of firms engage in core practices that support widespread adoption. Most firms have run only ad hoc pilots or are applying AI in just a single business process. **Why the slow progress? At the highest level, it's a reflection of a failure to rewire the organization.***

— Harvard Business Review, Building the AI-Powered Organization¹

Needless to say, transformation to Everyday AI, including embedding AI deeply into a company's operations, isn't something that will happen overnight, so a multi-pronged approach is necessary. This ebook will therefore delve into Dataiku's proven blueprint for successful Everyday AI implementation. It is divided into two sections, each with its own practical recommendations for execution:



Short-Term: Delivering Quick, High-Impact AI Wins — because let's face it, you can't afford to wait years to prove the return on investment (ROI) of AI initiatives.



Long-Term: Enabling a Transformative AI Culture — because thinking about AI implementation only on a use-case-by-use-case basis for the long term isn't sustainable (or economical).

Organizations need to do *both* in order to succeed with AI at scale, and this ebook unpacks the “how” on both fronts.

¹ <https://hbr.org/2019/07/building-the-ai-powered-organization>



Short-Term: Delivering Quick, High-Impact AI Wins

Bad news: If you're not able to quickly prove that AI delivers significant and tangible business value (we're talking about a timeline of months, not years), you will struggle to find success with your long-term AI ambitions. Unfortunately, there's a lot of negativity and scepticism around AI, from ideas about people being replaced by machines to trust issues and more. At Dataiku, we believe these ideas are unfounded — but that doesn't make them any less of an obstacle.

33% of U.S. CEOs cite employee trust as one of the greatest barriers to AI adoption.²

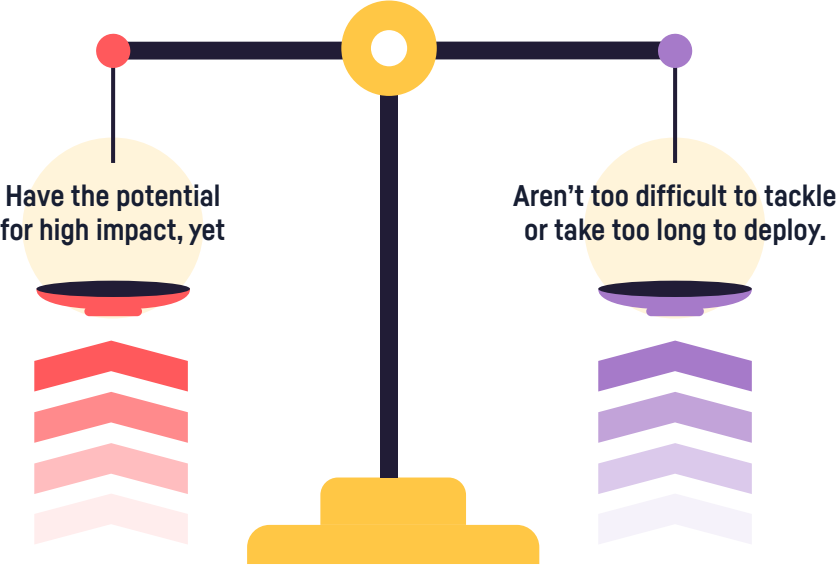
That means CXOs must be equipped with a solid short-term strategy for proving the value of AI, including showing that it's not just hype and that it can help make processes more efficient as well as improve decision making at all levels of the organization. This strategy should include the following elements.



² https://assets.ey.com/content/dam/ey-sites/ey-com/en_gl/news/2019/05/ey-the-ai-race-barriers-opportunities-and-optimism.pdf

Choose an Initial Set of Use Cases

The first step to short-term impact is defining an initial set of use cases that:



Of course, that’s often easier said than done, especially as “classic” AI use cases (think fraud, customer churn, etc.) have become more tablestakes than differentiators, not having the impact they might have had five years ago.



That's where a time-honored framework — the five Ws and an H — can come into play. An ideal AI project will have clear and compelling answers to each of these questions:

- **WHO will this project benefit?**

Collaboration is key — those developing AI projects need to work closely with business stakeholders or other subject matter experts to not just choose, but truly understand the audience. Who are they? What are their needs? What are their current processes and habits?

- **HOW will it specifically improve experience or outcomes, and HOW can this be measured?**

AI use cases (especially early ones when garnering organizational-wide support for AI efforts) should focus on opportunities with real and measurable business results.

- **WHY is using AI for this purpose better than existing processes?**

Avoid doing AI for the sake of doing AI and instead focus on heavy-hitting applications that really provide value at scale or some other tangible value that wasn't previously possible. If using AI wouldn't provide any larger value or ROI than existing processes, move on: choose another use case.

- **WHAT is the upside if it succeeds, and WHAT are the consequences if it fails?**

Choose an AI project that won't have devastating consequences in case of failure.

- **WHERE will the data come from, and does it already exist?**

Minimize external dependencies for a first AI project. This not only reduces complexity, but also steers teams toward early projects with a high speed to value, which is important for gaining organizational traction. As a part of this point, it's important to also evaluate essentials like project scope, access to subject matter experts, and availability of analytical resources (both human and infrastructure).

- **WHEN should an initial working prototype and, subsequently, a final solution in production be delivered?**

A “progress before perfection” and *minimum viable product (MVP) mindset* enables stakeholders to get excited about the long-term vision of an AI project and give feedback while it's still early enough for major course corrections. AI project timelines should include not only development to prototype, but development to operationalization.

For initial use case ideas, at Dataiku, we always recommend that our customers look at parts of their organizations that use a lot of processes and tools but that are not very sophisticated. A good example is pricing, which is often both an underinvested area on top of not being owned by a specific business unit or function. Pricing is a perfect candidate for AI, enabling teams to test and model more assumptions and better predict the impact of potential changes.

How Dataiku Can Help

Dataiku is the platform for Everyday AI, and more than 450 companies worldwide use Dataiku to systemize their use of data and AI to address a huge range of use cases.

But it's not just about technology alone — Dataiku helps customers build and deploy their first flagship use cases more quickly, either via support from our in-house team of experienced data scientists or by leveraging our pre-packaged business solutions.



A small snapshot of Dataiku customers' tens of thousands of use cases

Our industry solutions are Dataiku add-ons that accelerate advanced and basic industry-specific use cases. They are an operational shortcut to achieve real-world use cases designed with the purpose of business value generation. Taking advantage of Dataiku's core features, they are built to be fully customizable and entirely editable.

They come with:



A user-friendly interface that enables fine tuning to match with specific business requirements



Ready-to-use dashboards that can be customized



Documentation and training materials

Dataiku industry specialists develop solutions for every vertical, some of which include:



Financial services and insurance: Interactive document intelligence for ESG, news-based stock alert system




Health and pharmaceuticals: Accelerating drug repurposing, omnichannel marketing



Retail and CPG: RFM-enhanced customer lifetime value, market basket analysis

Identify Ambassadors & Early Adopters

When AI is widespread, there can be reuse and capitalization (note that, in this context, capitalization refers to building on something to create even more value). While tackling larger, high-priority use cases, the organization or even other teams in the business line can also take on lots of other smaller use cases by reusing bits and pieces, eliminating the need to reinvent the wheel with data cleaning and prep, operationalization, monitoring, and more.



AI overall becomes less costly when more groups are involved and pieces can be reused among them.

This is the crux of the long-term transformation component of AI initiatives. It’s not about just one successful use case (or even a handful of use cases), but about creating lots of successful, business-impacting use cases throughout the company. However, this longer-term goal starts with ambassadors and early adopters talking about their triumphs — what use cases worked, how they were executed, and what the results ultimately were for the business.

That means for short-term success with AI, never underestimate the importance of having a group of people at the organization that can both:

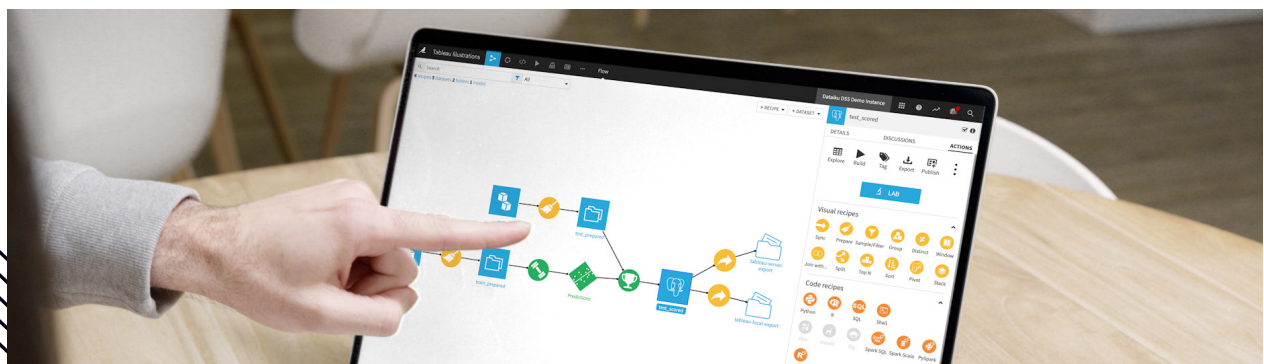
- Communicate the business value and results the team has seen on specific use cases, and
- Be cheerleaders for AI initiatives at the company at a larger scale.

Keep in mind that these ambassadors and early adopters of AI initiatives don't necessarily need to be the most technical users. While technical experts can communicate effectively on specific metrics (like model accuracy, for example), they likely don't have the proper resources to flesh out business value.

Only business teams can evaluate and quantify the value — both implicit and explicit — that they're seeing in ways that will resonate with other teams in the business unit as well as other business units around the company (and it's this word-of-mouth support that boosts short-term success).

Some other key stakeholders to get on board early with AI initiatives include:

- **A STRONG GROUP OF INDIVIDUAL USERS**, usually power users (though again, not necessarily technical) who can both evangelize and recruit other users plus solidify perception of AI initiatives as a positive force in the company, not just a top-down demand. This comes back to the idea of instilling trust from the bottom up.
- **TEAM LEADS AND MANAGERS**, who are often helpful when it comes to driving upskilling programs not just within their own teams but across the entire organization, so getting a few on board as early adopters is key. Upskilling is critical in the move to Everyday AI, and it's also good for business — for example, Accenture released a report in May estimating banks could increase revenues by 34% by investing in AI and upskilling, and evidence points to this benefit in other industries as well.
- **IT MANAGER(S)**, essential for effective, smooth roll-out of any technology as well as — from a more philosophical perspective — critical for instilling a culture of access to data balanced with the proper governance and control.



Of course, short-term success also means getting executive teams themselves on board as a strong front of ambassadors. Executives who invest in AI just because they feel they have to are not the ones that will come out ahead with creative or productive implementations. Those who are positive about what AI can bring to the table will have the potential not just to defend, but to disrupt.

How Dataiku Can Help

Because Dataiku is for both people that are technical and working in code as well as on the business side and low- or no-code, there's no need to focus short-term on enabling one or the other.

Instead, with Dataiku, identify ambassadors that:




Have pain points to be solved with AI (e.g., poor time to value, repetitive work, limited access to data, etc.) and/or ...



Have opportunities for new and exciting capabilities when armed with the right tooling (e.g., smoother access to EKS clusters that enables the processing of many more images, the ability to process exponentially more rows of data, ability to create machine learning models for non-coders, etc.).

Dataiku helps customers identify and empower these key early users by running lunch and learn sessions, hackathons, and more, on top of more formalized training via the Dataiku Academy and more tailored, bespoke training programs.



Long-Term: Enabling a Transformative AI Culture

It's not enough for organizations to attain Everyday AI at any price. Eventually, in order for AI initiatives to be sustainable, one must consider the economics: not just gains, but cost. That's why moving from the short-term, use case, quick win-driven approach to a longer-term transformation is paramount to success.



Attaining significant return on investment for discrete AI or a small collection of business area use cases is difficult because of the tactical focus of AI design and the accumulation of technical debt.

— Gartner “What is the True Return on AI Investment?” February 4, 2021

Unfortunately, there will come a point in time where the economic value of short-term AI initiatives decreases (see Figures 1 and 2) because:

- The marginal cost of the supplemental use cases is not decreasing.
- The marginal value of the supplemental use cases is decreasing (i.e., the first use case has more value than the Nth use case).
- The marginal profit of supplemental use cases quickly becomes negative.



So it's a mistake to think that the organization will be able to achieve Everyday AI by simply taking on increasingly more AI projects throughout the company. Instead, there needs to be a shift to a more holistic vision of value creation, driven by a plan rooted in the following principles.

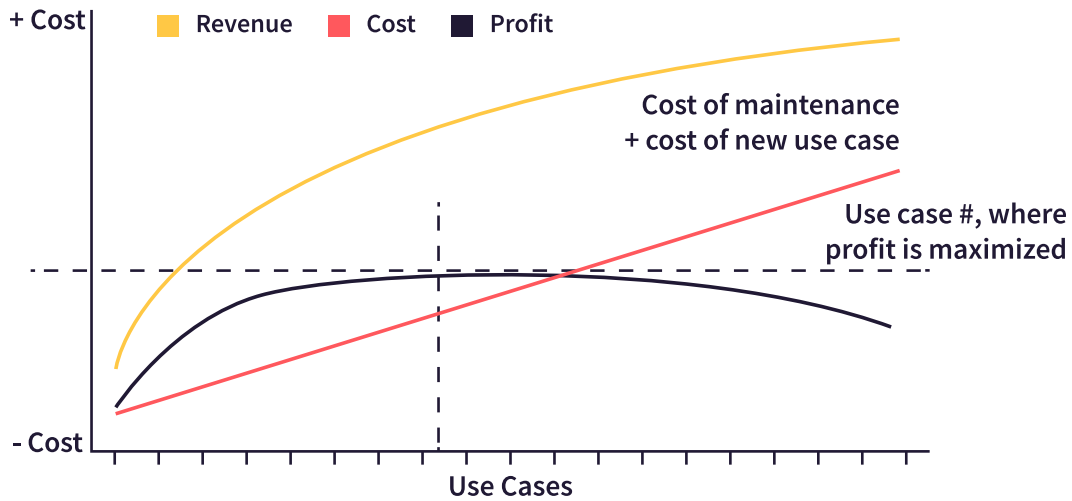


Figure 1: Cumulative revenues, costs, and profits over time (# of use cases); note that after use case #, profit is decreasing due to increased costs and stagnation of revenue




Figure 2: Marginal revenue, costs, and profit over time (# of use cases)

Identify All Opportunities for Business Impact

Ultimately, scaling with AI means aligning AI efforts to strategic objectives, getting past experimentation and moving beyond initial successes with low-hanging fruit use cases.

As discussed in the previous section, it's important to first tackle the “usual suspects” as a starting point in order to test and monitor operational frameworks while simultaneously building trust on AI impact. But the real transformation happens when teams start building beyond the first few use cases and discover that AI is not an isolated topic but rather the catalyst behind the development of core business.

At Dataiku, we've seen organizations find long-term value and business impact by finding an organized way to put people — and data — together that can highlight potential unique issues that fall in the middle of teams' responsibilities and are thus not being solved by anyone.



For example, Dataiku was working with a truck manufacturer that wanted to do a proof of concept (POC) on a relatively classic (but nonetheless challenging) use case: leveraging internet of things (IoT) sensors for advanced predictive maintenance. The project was extremely ambitious, and the ultimate end goal was to develop a system that would trigger alarms for drivers when maintenance was critical. In working on the project, the team put together data from different teams across the organization that they had never blended before, including classic sources like truck movement but also less obvious data such as warranty information.

Ultimately, in putting these disparate data sources together, they found some oddities: in particular, that there were some trucks that were supposed to be out of service, but instead, they were traveling around. In investigating further with business teams, they uncovered cases of warranty fraud — in other words, people were sending parts for repair for trucks under warranty, but actually using those parts in other trucks not under warranty. Ultimately, the POC was a success because the teams uncovered a problem that was even more important to solve than their initial business need.

This isn't to say that it's always a good idea to dig around aimlessly looking for problems to solve — it's still important to address real business concerns to get value from AI initiatives. However, exploring the possibilities of new, less obvious problems by bringing different people together is invaluable.

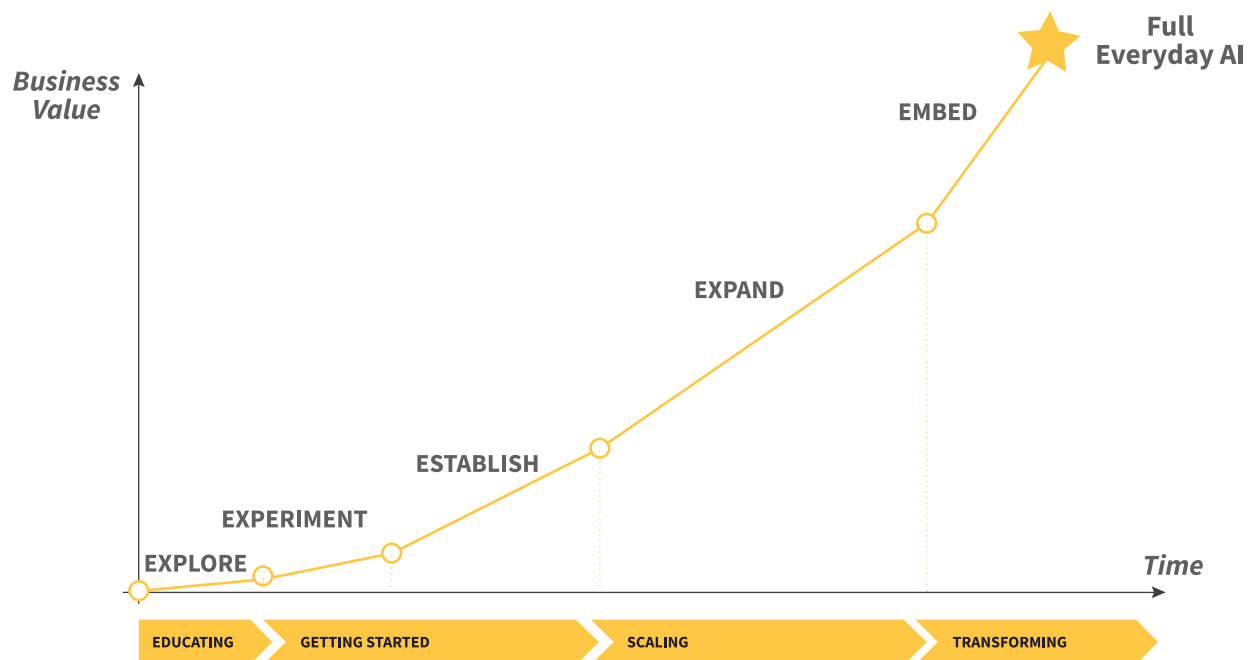
Continually generating ideas for increasingly more value requires a prioritization heuristic that allows tactical and strategic projects to emerge from shared efforts and understanding of AI's capabilities. However, it's important not to prioritize too prescriptively, either. Given the innovative and often disruptive nature of AI, there should always be room for experimentation. Adopt a three-lane approach to the organization's AI project portfolio: experiment, pilot, and productionize.



How Dataiku Can Help

Dataiku has an entire team of people devoted to helping customers derive both short- and long-term business value through Everyday AI. The team can help facilitate “Art of the Possible” workshops to explore company-level vision for AI and possible moonshot initiatives.

Dataiku has also developed a five-step journey toward the path to Everyday AI and works with organizations to benchmark as well as map progress. The maturity model can be used to demonstrate the value of analytics and build momentum internally, a key piece of the puzzle when trying to transform an organization at scale.



See where your organization stands — [TAKE THE AI MATURITY SURVEY](#)

Build the Right Capabilities and Processes

As detailed in the introduction to this section, scaling to achieve Everyday AI requires a fundamental shift in company culture, adopting processes and capabilities that will reduce the cost of each incremental AI use case. This includes, but is not limited to:

- **AI GOVERNANCE PRACTICES** for strategic alignment and steering.

Traditionally, data governance includes the policies, roles, standards, and metrics to continuously improve the use of information that ultimately enables a company to achieve its business goals. Data governance ensures the quality and security of an organization's data by clearly defining who is responsible for what data as well as what actions they can take (using what methods).

A traditional data governance program oversees a range of activities, including data security, reference and master data management, data quality, data architecture, and metadata management. With the rise of data science, machine learning, and AI, the opportunities for leveraging the mass amounts of data at the company's disposal have exploded, and it's tempting to think that existing data governance strategies are sufficient to sustain this increased activity.

But this thinking is flawed; in fact, the need for data governance is greater than ever as organizations worldwide make more decisions with more data. AI governance goes a step beyond data governance and should include critical new components, especially around Responsible AI.

However, another component of a modern AI governance strategy is finding a balance between governance and enablement that will allow Everyday AI to flourish. Put simply, governance should not — and cannot — be a blocker to innovation.

That means in many cases, teams need to make distinctions between proof-of-concepts, self-service data initiatives, and industrialized data products, as well as the governance needs surrounding each. Space needs to be given for exploration and experimentation, but teams also need to make a clear decision about when self-service projects or proof-of-concepts should have the funding, testing, and assurance to become an industrialized, operationalized solution.

● **MLOPS PROCESSES** for reduced risk and frictionless project delivery.

Closely related to AI governance is MLOps. Just as the use of data is governed by a data governance program, the development and use of machine learning models in production requires clear, unambiguous policies, roles, standards, and metrics.

Predictive models are the product of data, algorithms, and tuning. When deploying data pipelines and models into production, you might not expect your artifacts to change, but everything around them can and will. So, the data in production can be different from the data the model trained on, the infrastructure can break down, and more. The point is that organizations need systems to monitor pipelines, models, infrastructure, and services to make sure they are doing what they are supposed to.

A robust machine learning model management program would aim to answer questions such as:



Who is responsible for the performance and maintenance of production machine learning models?



How are machine learning models updated and/or refreshed to account for model drift (deterioration in the model's performance)?



What performance metrics are measured when developing and selecting models, and what level of performance and risk is acceptable to the business?



How are models monitored over time to detect model deterioration or unexpected, anomalous data and predictions?



How are models audited, and are they explainable to those outside of the team developing them?


Having MLOps processes in place is critical to long-term success with AI. MLOps in mature organizations include a plan for when things go wrong (because these organizations know they will). They have tools for troubleshooting, like access and event logs, plus failover and fallback options, like a failover model or rules. This type of robust process gives teams the time they need to figure out what is going on and fix issues or while they bring a newer version of the model through testing into production.

- **FORMALIZED EDUCATION & UPSKILLING PROGRAMS** for user adoption.

Everyday AI is such a shift that it requires a lot of time, energy, and resources via personalized, multi-step training, ingrained in the company strategy and culture and inclusive of hard and soft skills. The goal of training programs should be for non-experts to become autonomous on data projects up to the use of standard AI techniques.


However, upskilling people across an organization is a huge challenge. Diverse skill sets and needs mean one-size-fits-all training may not be the most efficient solution; however, the more specialized the training, the more time and effort required.

Dataiku customer NXP, one of the largest semiconductor suppliers in the world, has seen great success with its citizen data science program. Available to anyone at the company to elevate his or her competencies and skills around data science, the four-month program drives collaboration, upskilling, and self-service analytics at NXP by improving advanced analytics competency and data literacy among non-data professionals, addressing the challenge of solving business problems which have increasing complexity not served by legacy BI tools/methods, and positioning their business leaders to make better, more informed decisions.



88% of companies that see positive ROI from AI train and enable non-data-scientists to leverage AI.

— ESI ThoughtLab, *Driving ROI Through AI*, November 2020



In working with customers to support their upskilling programs, Dataiku has seen a few best practices to address this challenge:



Start with an assessment of skills to learn where gaps lie. The assessment should focus not just on evaluating the tech skills of business people for upskilling, but also the business skills of technical people.



From there, create a training plan based on long-term upskilling goals and existing skills. For example, if the goal is for 80% of the company to be proficient in Python, the solution probably isn't holding a three-day Python training for 80% of the company. Instead, it's evaluating who already has these skills (maybe you're already at 50%) and then identifying which people will be best suited to fill in the gaps. For example, maybe those who already have SQL skills will most benefit — and the organization will have the most to gain — from upskilling.



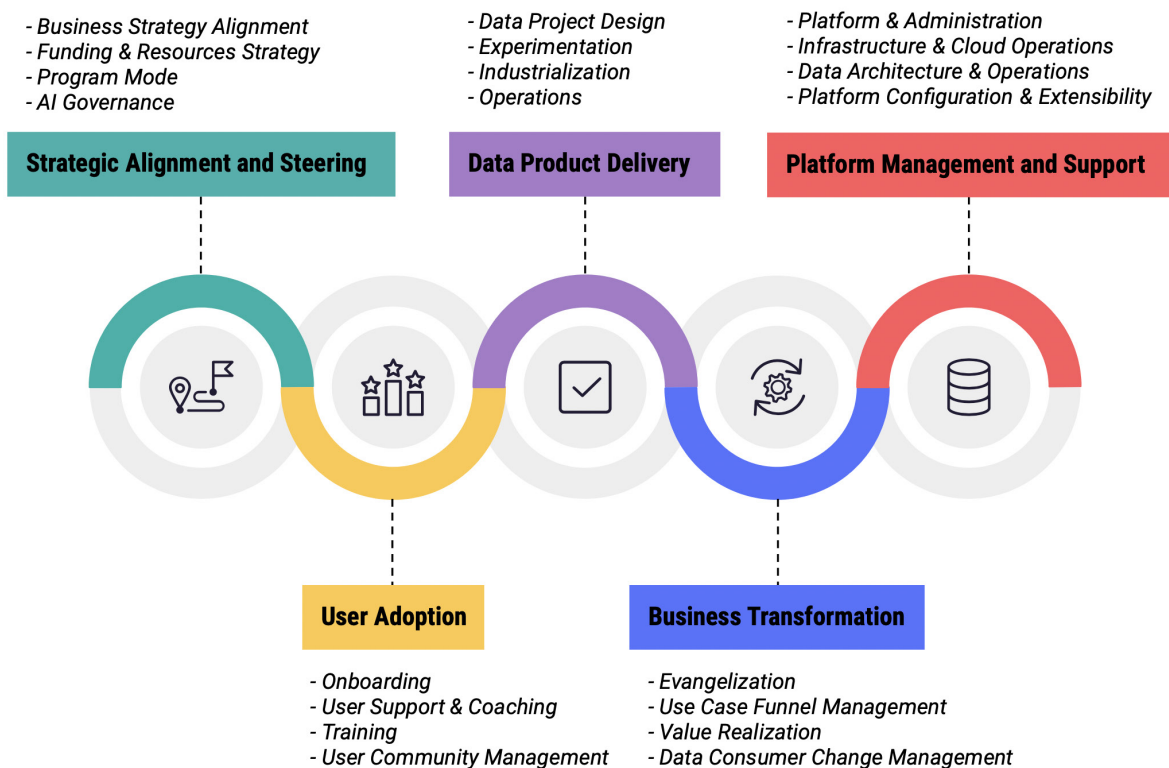
Create tracks based on existing skills or responsibilities. For example, Dataiku customer GE Aviation offers in-depth 100- 200- and 300-level courses to upskill and onboard end users to their self-serve data efforts in addition to a full-day executive training that is more focused to suit their needs.

In addition to these programs, building a community of engaged AI users is a proven way to drive transformative business impact. A strong community notably helps spread usage, core values, and overcome resistance to change. It also helps strengthen and refine a community of practices, leading to standardization — which will ultimately help with overall efficiency and ROI.

How Dataiku Can Help

Dataiku has developed a capability framework that articulates the 20 bricks organizations ultimately need to build to scale Everyday AI effectively. There is of course an order in which to proceed, and not everything needs to be perfect from the beginning. Dataiku helps its customers put a roadmap together to plan the organization's ramp up and make sure the right capabilities and processes are in place for long-term success.

Capability Framework



Conclusion: Dataiku for Everyday AI

We've seen in this ebook how success with AI is driven by both well-structured plans for short-term value combined with a path forward toward long-term change at all levels of the organization. From choosing use cases to identifying ambassadors and building opportunities for business impact on top of streamlined procedures, most of the focus has been on change management from the people and process perspective.

And for good reason — these are hands down the hardest components to get right. However, technology (though never a magic bullet) also has a role to play. In particular, choosing the right tools can usher in success by streamlining some of these people- and process-centered changes.

At Dataiku, for example, our biggest customers rely on us to drive analytics speed and agility, but with an overarching organizational control and governance. Having a platform that can facilitate and address both is invaluable to progressing more quickly toward the goal of Everyday AI.

Case in point:



Upskilling becomes easier with technology that facilitates the development of AI skills in a centralized, controlled, yet creative environment. Dataiku is the best of both worlds: fit for expert users who want maximum flexibility, but also suited for beginners who need guardrails and guidance.



Developing robust AI governance or MLOps processes is within reach by leveraging a tool that has built-in capabilities to facilitate and enforce their implementation. Dataiku allows AI creators and consumers alike to understand model outputs and increase trust as well as have high-level views of governance and MLOps processes.

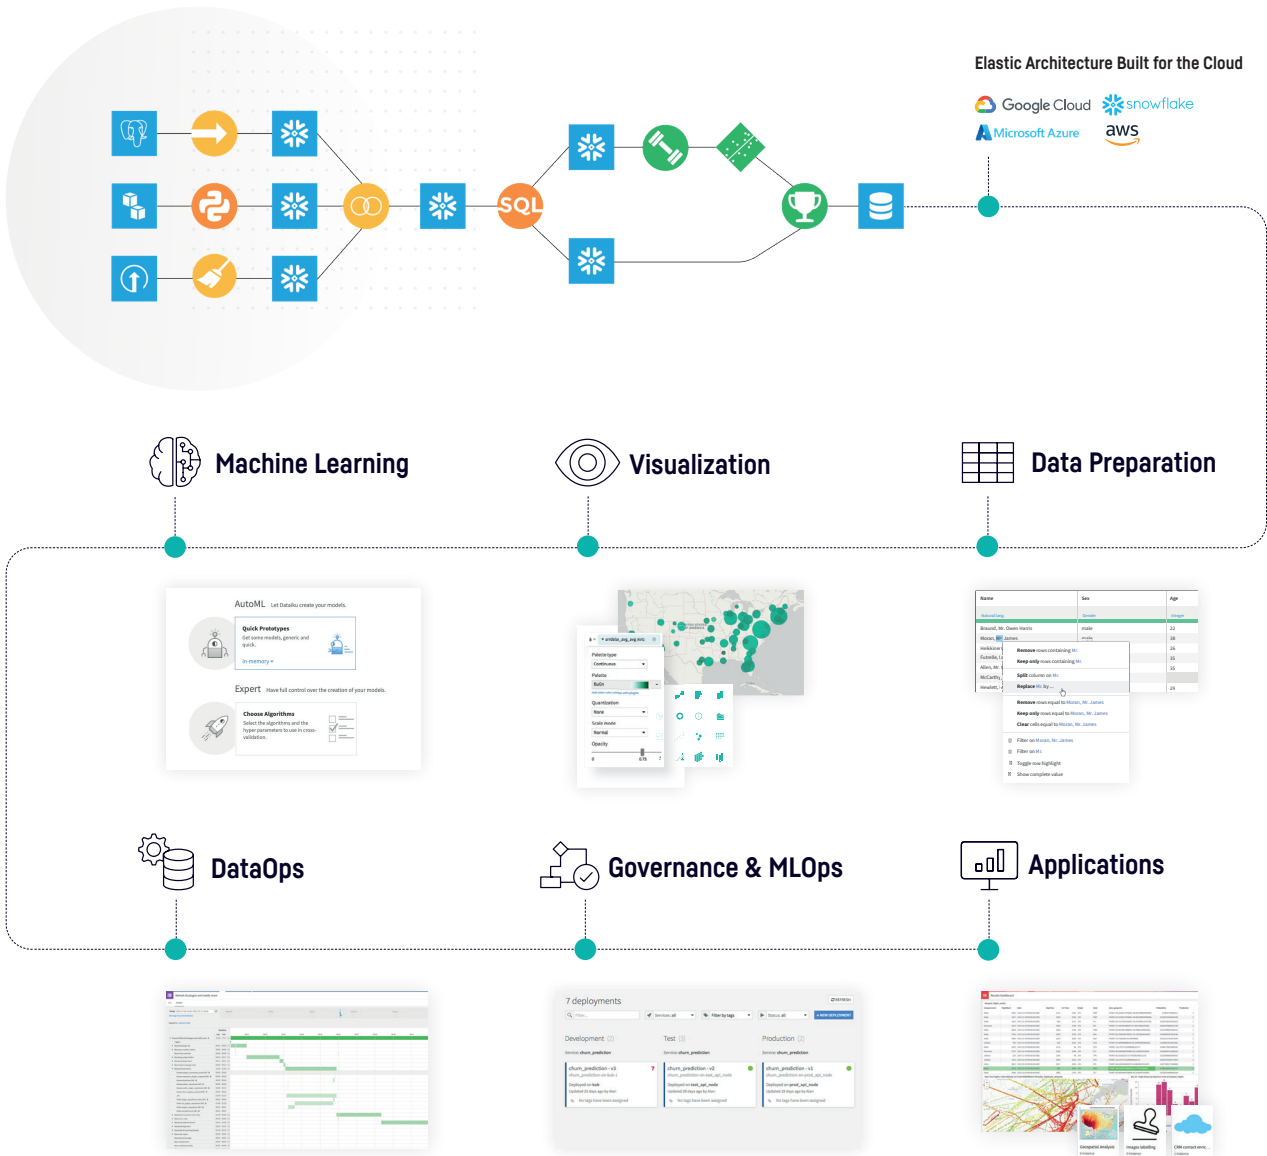


Short-term time-to-value can be accelerated with out-of-the-box solutions for your industry's most pressing use cases, which Dataiku offers and continues to release a robust catalog.

Learn more about systemizing the use of data with Dataiku to successfully implement Everyday AI.

Let's Go!

Everyday AI, Extraordinary People



450+
CUSTOMERS

45,000+
ACTIVE USERS

Dataiku is the world's leading platform for Everyday AI, systemizing the use of data for exceptional business results. Organizations that use Dataiku elevate their people (whether technical and working in code or on the business side and low- or no-code) to extraordinary, arming them with the ability to make better day-to-day decisions with data.

