

How to Thrive in the Age of Data Dominance

DataRobot **BITanium**

HOW TO THRIVE IN THE AGE OF DATA DOMINANCE 10 Keys to AI Success in 2022

TABLE OF CONTENTS

Introduction	
1 Decision Intelligence	3
2 Trust in Al	6
3 Fair, Unbiased, and Ethical Al	11
4 Decluttering Your Tech Stack: End-to-End Al	14
5 The Rise of Machine Learning Operations (MLOps)	
6 Democratize to Scale	
7 Regulatory Compliance	25
8 Continuous Learning	
9 Scalable AI: Build vs. Buy	
10 Al Cloud	
Summary	

DataRobot

DATA CONTINUES TO BE SOCIETY'S MOST VITAL ASSET.

The amount of data created over the next three years will far exceed the amount created over the past 30 years (IDC, <u>Worldwide Global DataSphere Forecast</u>, 2020-2024). It is also estimated that the global GDP could be 14 percent higher in 2030 as a result of AI – the equivalent of an additional \$15.7 trillion – making it the biggest commercial opportunity in today's economy (PwC: <u>Sizing the prize</u> What's the real value of AI for your business and how can you capitalise?, 2017).

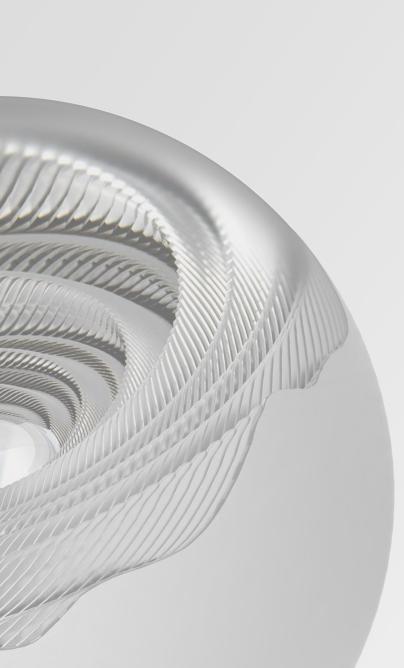
How do businesses capture this massive potential? With the enormous value of AI that is yet to be realized, every organization is under growing pressure to transform this sea of data collected by their systems into valuable insights that drive impact at all levels.

But the reality is that success with AI often falls short. Experts estimate that as much as 85 percent of AI projects fail to move past the proof of concept phase (DataRobot/Algorithmia, <u>5 Latest Trends in Enterprise Machine Learning</u>, 2021). And those that do fail are plagued by months-long delays — setbacks that can be very costly. This gap between demand and the industry's ability to deliver results is unsustainable.

The world of global business has never been more complex. With 81 percent of organizations now working across multiple public cloud providers (Gartner, <u>Why organizations choose a multi-cloud strategy</u>, 2019), multi-cloud is the de facto model for the enterprise, creating a complex model for data and infrastructure. The <u>rise of edge computing</u> will only accelerate the distribution of data in order to improve performance and reduce latency across distributed organizations. All of this creates new complexity in how we transform distributed sets of data into business insights, while working to protect the security, performance, and operations of every organization.

If all of this leaves your head spinning, you are not alone. Many organizations are daunted by the prospect of understanding AI and what they need to do to deploy AI successfully. That is why we are drawing from our own experience working with global enterprises to lay out ten key points that you need to consider when implementing your AI program. Every day, we help organizations solve AI challenges, so we are eager to share with you the steps to help you realize the enormous benefits that AI can offer your business in 2022.

DECISION INTELLIGENCE



How does your business make decisions? It is such a simple question, but for many organizations, decisionmaking can be a process wrapped in uncertainty and bogged down by complex processes. And it is getting more complicated by the day. In fact, in a recent survey, Gartner found that 65 percent of decisions are more complex and involve more stakeholders or choices than just two years ago (Gartner, <u>How to</u> <u>Make Better Business Decisions</u>, 2021). The current pace of decision-making cannot keep up with the fast-changing context in which business decisions are being made. Something needs to change.

No decision is made in isolation. Every decision in one area of a business affects other areas of the business. In addition, the decision-making apparatus is a constantly evolving process that needs to be responsive to changing conditions in the marketplace. Because of all of these factors, Gartner notes that **"effective decision making — that is connected, contextual and continuous — results in a host of business benefits, including greater transparency, accuracy, scalability and speed"** (Gartner, How to. <u>Make Better Business Decisions</u>, 2021). This is where decision intelligence comes in to ensure that the decisions that businesses make from their AI are clear, actionable, and have a positive effect across the organization. For example, a grocery store may predict that it needs 1,000 avocados at a particular store. This prediction is powerful information, but ultimately, the buyer at the store needs to make a decision about how many avocados to purchase. The prediction intelligence does not do that for them. If the forecast underestimates demand, that could lead to disappointed customers. As a result, most stores carry safety stock to ensure that doesn't happen. But how much extra stock should they carry? This decision is where the **"intelligence"** portion comes into play.

Organizations using AI most effectively put a lot of thought into decisions like these, creating decision rules that act on the output of their predictions to standardize the decision-making process across the organization. Often, companies start with simple decision rules, such as "carry a safety stock of an extra 10 percent" and over time, they layer on more intelligence. Unfortunately, implementing decision rules can be difficult to do.

So, where do you start? First, you need to decide whether the process has a human in the loop or whether the process will be fully automated. When a person is making the final decision, AI is a great decision-making aid, because it helps evaluate a lot of information very quickly. On the other hand, many decisions need to be made very quickly or at a scale





where human intervention is impractical, such as swiping your credit card. Other cases need to be automated because regulatory requirements dictate a consistency and transparency that you cannot achieve only with human judgement.

Good decision intelligence solutions incorporate both the human in the loop and automated decision processes. Frontline decision makers can streamline this process by building an application that gives them decisions right at their fingertips. But app development requires web programming and security skills that most data scientists and analysts just do not have.

In order to make app development easier, more scalable, and more transparent, we recommend incorporating the following rules:

Complex rules should be easy to create. Your decision logic should be easy to implement, understand, and update, embedding both your own business knowledge and your AI models.

Decisions should be made with speed and scale in mind. After creating a decision flow, you can apply it to datasets easily and make new decisions in real time or as a batch process to make millions of decisions each day.

Decisions should be transparent and trusted. In any app, you should be able to see which decision flow was used for any decision. This transparency should be available for everyone from non-technical stakeholders to experienced data scientists.

In the end, an AI project is only successful if it leads to outcomes that the whole business is excited about. Organizations applying decision intelligence can close the gap between AI predictions and business outcomes, making smart decisions by transforming the predictions from machine learning models into optimal decisions.

TRUST IN AI





While AI is a powerful tool that can deliver tremendous value to a company, there are numerous stories about lost public trust in AI and the concerns that people have around bias and fairness. Once AI users lose confidence in a system, it can be very hard for them to trust the models they deploy and the decisions they make. As an AI program expands, it can be even harder to scale the system and achieve better adoption across the enterprise. According to Gartner, 42 percent of respondents to a recent survey said that fear of the unknown was a top concern (Gartner, <u>3 Barriers to AI Adoption</u>, 2019).

For all of these reasons, it is important to establish trust between an AI user and the system so that users have the confidence to rely on the models they deploy and the decisions they make.

Establishing trust starts with better transparency and accountability. At DataRobot, we sort the components of trust in an AI system into three main categories:



Trust in the **performance** of an Al model.



Trust in the **operations** of an Al system.

Trust in the **ethics** of the workflow, both to design the AI system and how it is used to inform your business process.

TRUST IN AI 10 Keys to Al Success in 2022



Al Perfomance

The assessment of a model's performance is based, in part, on its ability to make predictions based on new data. We break down performance assessment into the following categories:

Data Quality – The performance of any machine learning model is directly tied to the data it was trained on. It is important to understand the provenance of this data, including how it was collected and curated.

Accuracy – You should be able to evaluate the accuracy of your model throughout the development process. Selecting an appropriate accuracy metric for the use case and communicating the results are both critical.

Robustness and Stability – A model in production must be reliable and stable in order to continue delivering value. Continuously monitoring a model in production for data drift, changes in accuracy, and system responsiveness are crucial to building a trustworthy system.

Speed – The speed of model scoring directly impacts how it can be used in your business processes.



TRUST IN AI 10 Keys to Al Success in 2022



Al Operations

A trustworthy AI system is one that has sound processes and operating procedures. Well-constructed operating procedures do not constrain people but instead build confidence and clear up ambiguity so that your AI systems align to organizational goals and ethics. The following are components to consider for a reliable system:

Compliance – Depending on your industry, it may be essential to align model performance to regulatory requirements when putting a model into production.

Security – Sensitive information about a company, such as revenue, employee performance, salary, and sales leads may be part of your training data. It is essential to implement best practices to keep your data and your models secure.

Humility – Not all model predictions are made with the same level of confidence. A trustworthy system knows when to be humble and how to respond to uncertainty.

Governance and Monitoring – Without good governance, the best models might still result in undesired or unintended behavior. Good governance and monitoring is essential to ensuring that your AI system delivers the value you need in production.

Business Rules – Al delivers the most value to your business when it is integrated into existing processes, making it a valued member of your team and a seamless part of your business process.



TRUST IN AI 10 Keys to Al Success in 2022





There is no one-size-fits-all approach to AI ethics, especially since AI systems and the data they use can span borders and cultures. Ethical components to consider for your AI system include the following:

Privacy – Individual privacy is a fundamental right, but it's also complicated by the use and exchange of data. It is important to define what role AI plays in the management of your sensitive data.

Bias and Fairness – Al bias is a difficult topic to navigate, due to both its complexity and the social implications of what it means to be fair in decision-making. Look at what tools can be used to help define the values you want reflected in your Al system and how Al systems can be used to promote fairness and equity in your decision-making.

Explainability – Explainability is one of the most powerful ways to build trust between a user and an AI model and should be a central feature of your AI solution.

Impact – An impact assessment tool helps you think systematically about how to identify the desired behavior of your system that will reflect your organization's values.

These dimensions of trust are not a checklist. Rather, they are critical guideposts that your organization can rely on to ensure trust from the moment you implement your AI system.



FAIR, UNBIASED, AND ETHICAL AI





While we covered trust in AI in the previous section, it is important to call out ethics in AI specifically. Why? Because as organizations scale their use of AI, it has the power to impact people in either a positive or negative way. In other words, when you have AI making vital decisions that impact real people in real time, the risks associated with AI become serious very quickly, particularly those around AI bias. It is no wonder, then, that bias and fairness in AI has become the cause of much concern and debate.

Far from being a neutral technology, AI can often reflect the bias that humans unwittingly encode in its data, leading to unintended consequences. And the risks can be far-reaching. According to the Capgemini Research Institute, Ethics in AI customer survey, 2020, almost 60 percent of organizations have attracted legal scrutiny and 22 percent have faced a customer backlash because of decisions made by their AI systems (Capgemini, AI and ethics – why is it important?, 2020). In addition, 70 percent of customers expect organizations to provide AI interactions that are transparent and fair (Capgemini, AI and the ethical conundrum, 2020).

So, where do you begin to address this crucial issue? To many companies, it may feel like the goalposts around Al bias and fairness are constantly moving. But there are steps you can take to encode your values in your Al, make better and fairer decisions, and reduce your organization's risks.

Start with doing an impact assessment.

Impact assessment should be conducted before modeling begins, as initial data sources are identified and evaluated. However, it is a valuable tool to revisit at different junctures in the process. Your assessment should be threefold:

- Include multiple stakeholders.
- Ensure that your fairness metric is aligned with organizational values.
- Ensure that your fairness metric is compliant with regulations.

Implementing ethical AI starts with educating people.

It is crucial to understand the perspectives and needs of each stakeholder group. Al innovators need to focus on problem formulation, understanding how to frame the problem at hand for an Al system to prevent unintended bias. For people creating Al, Al education can revolve around data collection and data processing best practices, the technical definition of bias, and how to implement these measures within the system.



As you build your bias and fairness metric, your team can ask some key questions that will yield compelling answers:

What best practices will keep private information safe while building and using an AI model?

Understand what kind of data can be defined as personally identifiable information (PII). Generally, this is data that should not be used in the training of a machine learning model in the first place, which aims to learn relevant patterns from a dataset.

What unique risks to privacy do AI systems pose?

Gathering sensitive information may be possible and advantageous for an organization, but there is a huge risk that your brand may receive negative attention in the marketplace, even beyond the legal concerns.

What does it mean for an AI model to be biased?

A family of bias and fairness metrics in modeling describe the ways in which a model can perform differently for distinct groups within your data. Those groups, when they designate groups of people, may be identified by protected or sensitive characteristics, like race, gender, age, or veteran status.

Where does AI bias come from?

The largest source of bias in an AI system is the data it was trained on. Ultimately, machine learning learns from data, but that data comes from us – our decisions and systems.

How do I measure AI bias?

You may want to ensure fairness of error: this means that the model performs with the same accuracy across groups so that no one group is subject to significantly more error in predictions than another. Alternatively, you may want to ensure fairness of representation: this means that the model is equitable in its assignment of favorable outcomes to each group.

How can I make a model explainable?

Making sure that you can explain your model's outcomes and why it arrived at a conclusion is essential to trust in AI and ensuring fair outcomes. No one trusts a black box, so make sure that your AI is not one.

DECLUTTERING YOUR TECH STACK: END-TO-END AI



With the heavy lift involved in getting AI technology and teams up to speed, why do companies make such significant, upfront investments? The answer is that the positive returns are impossible to ignore. Analysis shows that AI leaders generate five times the amount of ROI than firms that are further behind on their AI journey (DataRobot, Driving ROI. with AI, 2020). In fact, 22 percent of firms state that at least 5 percent of their earnings in 2019 were attributable to AI (McKinsey & Company, The State of AI in 2020). Organizations that adopt AI now will survive and thrive while those that lag will fall behind.

To generate those kinds of returns on investment and not spend time playing catch-up, it is essential that companies adopt an end-to-end AI solution that is a single ecosystem rather than separate tools that were not designed to work together. It also means adopting a solution that allows businesses to be independent



the amount of ROI that AI leaders generate comparing to firms that are further behind on their AI journey

from the underlying technology infrastructure. This approach allows you to own your Al, deploying it to your cloud platform or platforms of choice, with the freedom to move your Al artifacts around, regardless of whether it's hosted on AWS, Azure, Google Cloud, or your own infrastructure on-premise.





DECLUTTERING YOUR TECH STACK: END-TO-END AI 10 Keys to AI Success in 2022

An end-to-end platform delivers AI success in numerous ways. For one thing, it takes you from data to value much more easily. In addition, data will drift, requirements will shift, and models will need to be retrained or replaced. It is much easier to handle these changes when you have a product that works throughout the full lifecycle. Some of the other reasons include the following:

Simplifies the tech stack.

An end-to-end platform allows your team to focus on strategic work, rather than time-consuming manual work that drains resources. Instead of spending time getting all of your existing technical tools to work together and managing update schedules and product lifecycles, your team can focus on doing their jobs rather than wrestling with non-integrated technologies. The efficiencies gained by this approach will contribute greatly to your return on investment (ROI).

Increases AI adoption.

Having a single end-to-end platform improves Al adoption across the enterprise. Rather than learning the features of multiple systems and keeping track of the updates and schedules of multiple tools, the whole team will spend far less time and resources if it only has to learn one solution and can communicate far more effectively and efficiently.

Standardizes processes.

Data management can be a challenge in any Al project. If you are using separate tools across different teams, it can be a major headache. From data input to model deployment, one endto-end platform enables Al creators, operators, and consumers to manage your data, models, governance, and all aspects of the Al lifecycle under a single pane of glass. This standardization ensures repeatability and can help you to scale your Al initiatives.

Simplifies life for IT.

With one integrated platform, your IT support team can address issues much more quickly than if you have numerous, disparate AI systems. Because of this efficiency, it's a lot more likely that you can scale your AI more effectively and achieve AI success.





Many organizations have a number of teams using a variety of platforms and languages across various departments to solve data science problems — a fact that runs up against the reality of production run-time environments. Typically, this is a complex variety of environments and systems managed by the IT Operations team. This is where models are to be deployed before many of them can provide meaningful value. But therein lies the problem.

Often, there is an immense gap between IT and data science teams. Deploying and running models in production is very different from building them, and it requires a brand new set of skills and tools due to the unique nature of machine learning. Until you bridge this chasm, your machine learning will not fulfill its promise. It is no surprise that among organizations that are actively engaged in machine learning, **55 percent have not deployed a machine learning model**, according to Algorithmia's 2020 State of Enterprise Machine Learning report.

Enter Machine Learning Operations (MLOps), which brings together people, processes, and practices to automate the deployment, monitoring, and management of machine learning models in a way that is scalable and governable.



55% have not deployed a machine learning model

THE RISE OF MACHINE LEARNING OPERATIONS (MLOPS) 10 Keys to AI Success in 2022



Why is MLOps so essential? We all know that we derive value by having models in production. Your competitive advantage comes from how quickly you can deploy and iterate on models in production. In other words, speed and scale are key. This is where MLOps provides the answers to several key challenges:

Challenge #1:

Iteration Speed

The machine learning lifecycle is an evolving ecosystem where new data is introduced continually, and models are constantly retrained and monitored, deployed, and managed. If this step is manual, it is impossible to keep up with the speed and scale with the changes in data and deliver results.

Solution:

Breaking out of your traditional app development and speeding up the entire process of model deployment, training, and management will allow you to scale your Al.

Challenge #2:

Training and Production Are Very Different Architectures

In most organizations, there will be more than one existing training tool. There is a long compute cycle, and it is usually run by a single user.

Solution:

MLOps helps your organization put models into production in shorter compute bursts that can accommodate many different users. It is also much more elastic and stateless and, therefore, works with the constantly evolving landscape.





Challenge #3:

Heterogenous Tooling and Dependencies

Organizations usually work with dozens of language and framework combinations, and new frameworks emerge every year or so. Frameworks and languages evolve constantly and require ongoing maintenance and testing. In addition, hardware dependencies usually require a substantial investment in architecture.

Solution:

With all of this consistently changing, you must have an MLOps system that is flexible enough to build heterogeneous tooling for dependencies and languages.

Challenge #4:

Composability Compounds the Challenge

In most organizations, there will be more than one existing training tool. There is a long compute cycle, and it is usually run by a single user.

Solution:

MLOps helps your organization put models into production in shorter compute bursts that can accommodate many different users. It is also much more elastic and stateless and, therefore, works with the constantly evolving landscape.



Challenge #5:

Diversity of Models and Frameworks Complicates A uditability and Governance

In order to scale out machine learning, you need to know who called out what data, when, and why. This is a very important step, since this will be a requirement for providing that traceability for auditing and governance.

Solution:

Being able to provide traceability is key for doing machine learning at scale successfully. MLOps provides governance that delivers model version control, automated documentation, and complete, searchable lineage tracking for all models in production. All of this allows companies to minimize risk and maintain transparency.

Challenge #6:

Lack of Reusability

Teams are reinventing the wheel unnecessarily. Models and other assets often exist only on laptops or local servers, and the use of multiple languages and frameworks introduce incompatibility.

Solution:

How can we provide a platform that can run all these models no matter where they are or where they were created? MLOps provides reusability of models so that you do not have to waste the time of data scientists doing the same thing over and over again, allowing you to package, control, and scale your models.

Most organizations cannot build a machine learning operations systems internally in a way that is resourceful and cost-effective. The solution is an external MLOps system that provides a single environment for continuous integration and deployment of AI, thereby significantly reducing the complexity that most organizations face with their AI projects. This solution allows you to monitor everything, everywhere within a centralized location, regardless of where that model is hosted or how that model is deployed.

DEMOCRATIZE TO SCALE



Many organizations find that democratization is an issue that comes up continually and is an essential component in implementing effective AI.

Democratization means making Al accessible to everyone in your organization. In many companies, the perception is that the Al belongs to the IT and data science teams and that only those with technical backgrounds can build, manage, and understand its complexities. The reality is that Al should never be separated from outcomes that the business wants to produce, so it is essential that the rest of the business is invested in every aspect of your Al. To accomplish this, you have to include the non-technical business stakeholders, including business analysts, leaders, and executives, because when everyone is invested in your Al, you can build an Al-driven enterprise.

2

3

4

5

But where do you start? There are a number of key steps you take to make sure that your AI is democratized across your enterprise:

Assess your current situation. Before getting started, it is essential to know where you are starting from. Take an honest look at your situation from a technical and business perspective and assess what your goals are. What existing processes can be automated or enhanced with AI? What gaps do you need to fill in? Where can you make the most impact?

Build a solid education program about AI. Many companies hold sessions with all levels of employees to educate them on how AI and machine learning work and what it can do for the business. While data scientists are key players in your AI program, the route to success is through those who are solving the business problems, helping them to feel comfortable and letting them ask the right questions.

Align with critical business goals and values. Many businesses mistakenly believe that AI is about technology when it's really about business outcomes. Therefore, it is important that you choose AI use cases that will have the biggest impact in the fastest amount of time. Look for those areas with maximum impact.

Determine everyone's role in your Al project. Any project will go more smoothly when people know what their roles and responsibilities are. That is why it's key to set up an effective organizational chart that tells everyone what their goals are and how they fit into the project. Many companies have success with establishing a Center of Excellence to provide consistency, encourage rapid development, and share best practices.

Get leadership support. Perhaps the most important aspect of building your AI program successfully is getting support from your company's leadership. Beyond simply signing off on the budget, your leadership team needs to help determine what targets you want to hit, how the organizational chart will be drawn up and, most important, how leadership will encourage and empower people while removing blockers.

REGULATORY COMPLIANCE

Organizations putting machine learning models into production are dealing with regulatory, compliance, and corporate risk minefields that can undermine their

production are dealing with regulatory, compliance, and corporate risk minefields that can undermine their efforts to become AI-driven if they do not proceed carefully. 56 percent of organizations struggle with IT governance, compliance, and auditability (DataRobot/ Algorithmia, <u>2021 Enterprise Trends in Machine</u> Learning). With the introduction of privacy regulations like the California Consumer Privacy Act (CCPA) or the General Data Protection Regulation (GDPR), it is a guarantee that companies in every industry will be impacted by privacy obligations at some level. Even if your industry is not facing compliance rules now, it should expect to see them on the horizon soon.

Governments have stepped into the discussion, too, instituting regulations that companies may need to follow closely in order to comply with the law. <u>The United States House of Representatives</u>. <u>Resolution 2231 seeks</u> to establish a requirement of algorithmic accountability, addressing bias and discrimination, a risk-benefit analysis and impact assessment, and issues of security and privacy. The government of Canada has put out its own online <u>Algorithmic Impact Assessment tool</u>.

While most governments have taken a "wait and see" approach to AI regulations, the European Union has

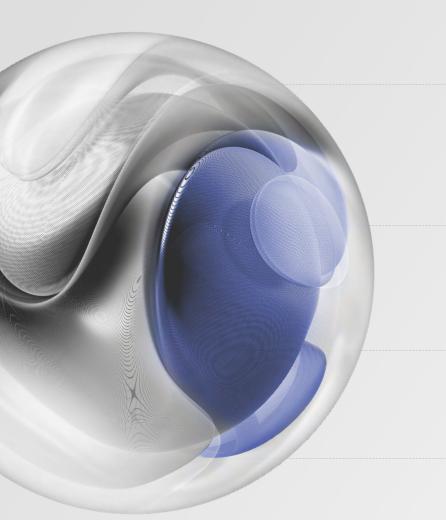
56%

percentage of organizations that struggle with IT governance, compliance, and auditability

taken a proactive stance. In April 2021, <u>the European</u> <u>Commission issued a proposed AI regulation</u>. Their group of experts recommended that in order for companies to operate high-risk AI systems within the European Union (EU), they will have to meet certain requirements, such as strict data protection obligations in addition to the GDPR and mandatory logs of all activity. Failure to comply could result in hefty fines.

This issue becomes especially critical for organizations operating on a global scale, where the maze of rules and laws becomes very difficult to navigate. In these situations, organizations need to maintain complete model lineage tracking (approvals, model interactions, versions deployed, updates, etc.), something that is practically impossible to perform manually.

26



REGULATORY COMPLIANCE 10 Keys to Al Success in 2022

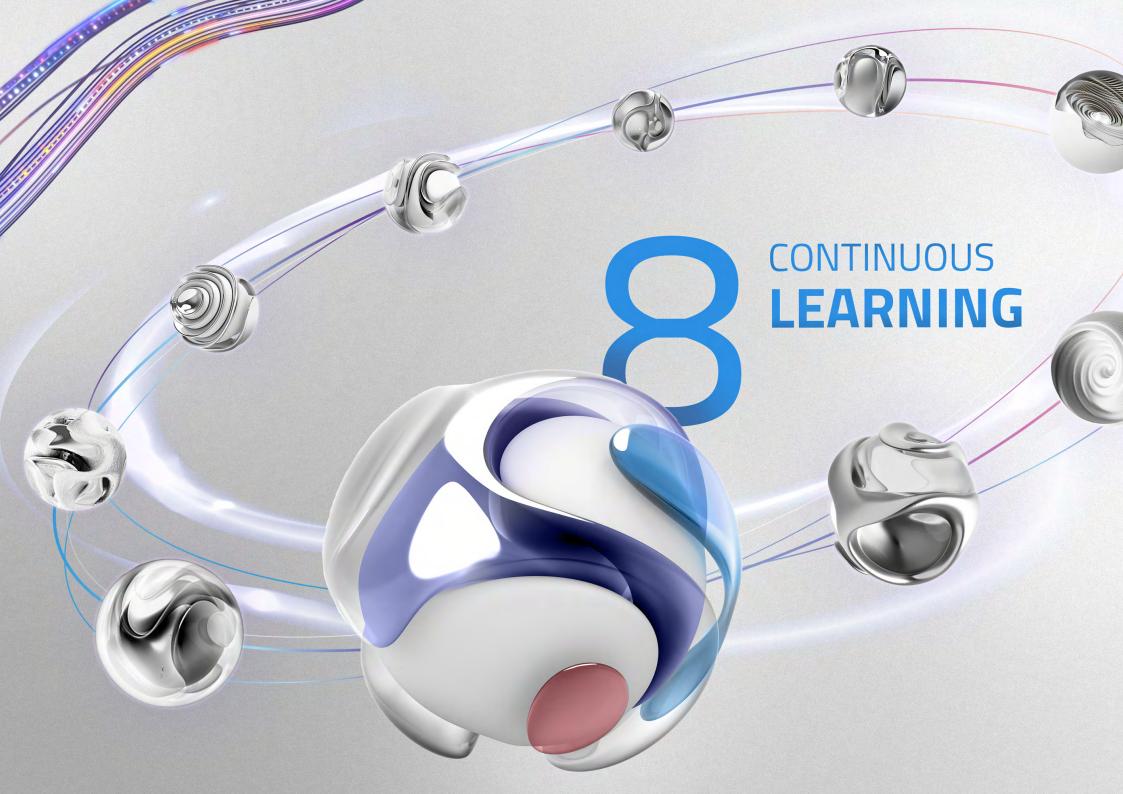
But with a robust, comprehensive AI solution to support your business, you can navigate these new and often confusing waters with confidence. There are a few key hurdles that companies need to overcome to get to that point, and they all involve ensuring that proper model deployment, monitoring, lifecycle management, and governance are in place.

Continuous Monitoring. Models can degrade over time for a variety of reasons. Customer behavior, economic conditions, and the weather can all change quickly, rendering your dataset obsolete. When models begin to decay, in the form of data drift or a loss in accuracy, this may signal it's time for a training refresh or a new model altogether. Without dedicated production monitoring that is designed for machine learning, your business could be exposed to risks.

Scalability. As organizations build out robust machine learning initiatives, the number of models actually in production may grow exponentially, making managing the lifecycle of these models a cumbersome task. Bringing order to this process means ensuring that the phases any model goes through are streamlined, approved via a flexible workflow, and automated. A production model lifecycle management system ensures efficient oversight of these processes so that you can scale quickly and efficiently.

Model Lineage and Audit Trail. Organizations putting machine learning models into production need to maintain complete model lineage tracking so that they can track model history and performance and show the path they took to making their decisions.

Service Health. The AI system must be consistently available and operate at a high latency. Users need to know that they can access their models and have consistent API access for application teams on the other end, regardless of deployment environments or cloud service providers.







Things change, and businesses that embrace change succeed. But the question is never simply, what has changed? Businesses must also ask, why did things change? What is causing it? Which data is affected? How do we adapt to these changes? When the marketplace is shifting – and your data along with it – what processes can you put in place to adapt, analyze information, and act quickly?

In order to embrace change, you need an AI solution that is built on a foundation of agility, efficiency, and nimbleness. Continuous learning should be a key component of any AI platform with these capabilities, which means:



Automated strategies that keep your production **models at peak performance** regardless of external conditions.



The ability to **refresh your models** based on the schedule of your choosing or when data drift occurs.



The ability to automatically create and recommend new **challenger models** for you.



How Can You Deliver Continuous Learning?

You should work to simplify your models through a process of constant iteration and experimentation. The pre-training and tuning that takes place before deployment is important, but it is fine tuning after deployment that increases accuracy. Many data science teams do not spend as much time refining a model after deployment, but this step is just as important. If you have a backlog of use cases, it can be incredibly laborious for data scientists to spend their time working on each problem. With high-quality automated machine learning (AutoML) with continuous learning, you can break this cycle, intensifying your experimentation and evaluation efforts when the models go live without actually investing more of your data scientists time - an efficiency that enables you to scale your AI efforts further.

Continuous Learning and Challenger Models

When you deploy a model for the first time, you are usually trying to get a good, solid model that works well with the goal of iterating after the fact. But experimenting and iterating as data changes is a key part of getting the most optimal predictions from your models. With an AutoML system running in the background, you are able to run experimental, challenger models all the time, which can suggest to you that something has changed and needs to be reevaluated. Through this challenger-model scenario, you can change your model without changing anything about the data engineering pipeline that feeds into it. Essentially, it is a data science improvement without requiring the investment in data engineering on the back end.

Adherence to Compliance and Regulation

With continuous learning, you can establish a governance framework that ensures that your production models always pass through existing governance policies and gated approvals first. This framework gives your company the assurance that models are never updated or replaced without passing through existing governance policies and gated approval workflows that protect your data and your organization.

The Competitive Picture

Keep in mind that for every model you have deployed, your competitors probably have a similar model actively competing against yours. So, the ability to update your models quickly, competitively, and at scale becomes a major challenge. With continuous learning, you can meet that challenge with an efficiency of scale that helps your company keep its competitive edge.

SCALABLE AI: BUILD VS. BUY





When it comes to building an Al-driven organization, one of the first questions many teams encounter is: should we build our own Al solution, or buy one? It's a complex question with many factors to take into consideration.

In the early days of AI, any organization that wanted to deploy models at scale had to build an AI system from scratch. Many organizations still gravitate towards this approach, but today, it is largely seen in organizations with a large number of models. Among organizations with more than 100 models, 60 percent choose to build and maintain their own systems from scratch (DataRobot/Algorithmia, <u>2021 Enterprise Trends in</u> <u>Machine Learning</u>). Many may have done this before other third-party options were available, but today, there are platforms that can take away the costs and headaches associated with developing an AI solution on your own. So, how do you decide which option is right for your organization?

One of the first factors to take into consideration is the resources you have available to devote to customization of your solution. As your business works to solve its own unique problems, you will probably need to optimize your system to address your specific needs. This process often requires a huge amount of your data science resources and time. Since 38 percent of organizations spend more than 50 percent of their data scientists' time on deployment (DataRobot/Algorithmia, 2021. Enterprise Trends in Machine Learning), it is worth taking into account how much additional bandwidth your IT and data science resources have to spend on customization on top of their already stretched resources. Keep in mind that this is also a problem that gets harder when trying to scale your AI.

Organizations must also consider the fact that premier AI talent is in incredibly high demand. According to the TalentSeer 2020 AI Talent Report: Current Landscape and Market Trends, companies that hire leading AI engineers must be willing to pay a median base pay of \$200,000 in compensation to be competitive with top technology companies. Not only is top AI talent able to command high salaries, but they are also on the hunt for exciting challenges that they can solve with their skills. Often, this does not include building a platform from scratch — a fact that your company should definitely consider when weighing all the factors involved in building a system and scaling it.



KEY TO SUCCESS



Finally, it takes time to establish trust in an AI system throughout your organization. Many people still feel skeptical about AI and how much autonomy to give it over a business's most critical decisions. The time it takes to create an infrastructure of people, training, and systems that people feel they can rely on will factor into how guickly you can put models into production.

When considering the costs of buying an AI platform, many leadership teams focus on the costs of adding one more solution to their technology stack. Here is where it is worth noting that, when we look at MLOps technology in particular, buying a third-party solution costs 19-21 percent less than building your own (DataRobot/Algorithmia, 2021 Enterprise Trends in Machine Learning). In addition to up-front costs, organizations that buy a third-party MLOps solution also tend to spend less time on model deployment. After calculating the average percentage of data scientist time spent on model deployment, the report noted that it was lower for organizations that buy a third-party solution as opposed to building one from scratch.

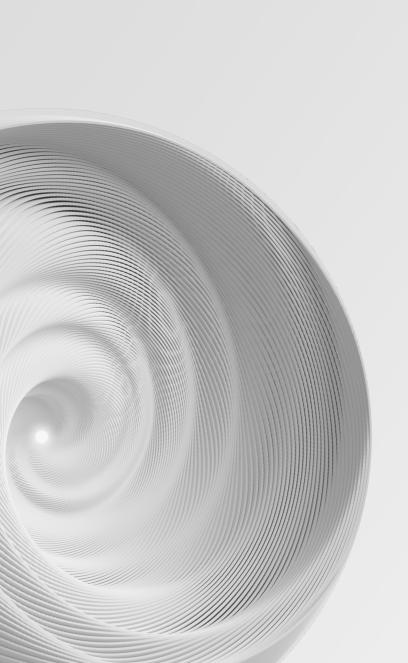
All of these elements together can make the difference between scaling your AI and realizing ROI or slowing down your process and risking falling behind your competitors. When considering the combined factors of customization, trust in AI, and hiring AI talent, many organizations realize that buying an AI solution or building a hybrid solution that gives them both customization options and out-of-the-box readiness saves time and helps them solve their most pressing business problems faster.

Build vs. Buy cost

Building an MLOps solution in-house is 19-21% more expensive than buying one



ALCLOUD



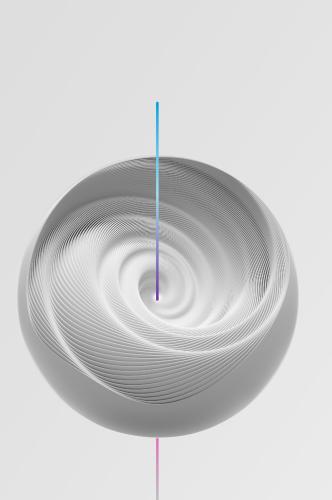
 KEY TO SUCCESS
 1
 AI CLOUD 10 Keys to AI Success in 2022

The marketplace is reaching an inflection point. All is becoming the essential technology to accelerate business growth, drive efficiency, and help the most successful organizations across the private and public sectors differentiate themselves. Businesses can go beyond simply predicting market behaviors to accelerating innovation and jump starting economic growth. All has the potential to help erase historical bias in otherwise entrenched institutional settings. And in medicine, All can accelerate breakthrough medical treatments and even predict and prevent the next global pandemic. In short, the possibilities of All are immense.

But realizing Al's potential is often hampered by a unique set of challenges. These challenges include the following:



While organizations need to move faster to realize value from their AI investments, the reality is that 87 percent of organizations struggle with long deployment cycles (DataRobot/Algorithmia, <u>5 Latest Trends in Enterprise</u>. <u>Machine Learning</u>, 2021). Delays mean missed opportunities that can lead to lost revenue, damaged customer relationships, and an advantage for the competition.





Enter AI Cloud and augmented intelligence to help companies realize the potential of their AI initiatives.

Augmented Intelligence

means bringing together the best of machine intelligence with human intelligence to realize better outcomes from predictive analytics. It helps organizations to merge the benefits of human intuition and experience with the efficiency and power of machine learning.

AI Cloud

is a new approach built for the demands, challenges, and opportunities of AI today. It is a single system of record, accelerating the delivery of AI to production for every organization. By bringing together data scientists, data engineers, and IT and business users, AI Cloud is a unified environment that allows for collaboration and continuous optimization across the entire AI lifecycle.

Al Cloud and augmented intelligence help to close the gap between how Al is developed and deployed and how it is delivered. By bringing together and harnessing the expertise of a broad set of teams, Al Cloud can:

- Transform diverse data from various sources.
- Unify diverse teams of data scientists, analytics experts, business users, and IT professionals.
- ✓ Drive trusted predictions from an end-to-end platform.
- Seamlessly deploy to any of the following major cloud infrastructure providers, to operate at scale: Amazon Web Services, Microsoft Azure, and Google Cloud Platform.
- Extend to business apps and integrate with current systems.

The effect is to close the gaps that have plagued AI projects where teams are disconnected, data is disorganized and distributed, and security and governance are an afterthought. AI Cloud and augmented intelligence allow businesses to deliver clear and powerful business insights accurately and at scale to meet the increasing demands of the marketplace.

HOW TO THRIVE IN THE AGE OF DATA DOMINANCE 10 Keys to Al Success in 2022

SUMMARY



KEY TO SUCCESS #1: Decision Intelligence

With decision intelligence flows, you can close the gap between your Al and the actions you take and make smart decisions by transforming the predictions from your machine learning models into optimal decisions.



KEY TO SUCCESS #2: Trust in Al

It is important to establish trust between an AI user and the system so that users have the confidence to rely on the models they deploy and the decisions they make. Establishing trust starts with better transparency and accountability. At DataRobot, we sort the components of trust in an AI system into three main categories:

- Trust in the performance of an AI model.
- Trust in the operations of an AI system.
- Trust in the ethics of the workflow, both to design the AI system and how it is used to inform your business process.



KEY TO SUCCESS #3: Fair, Unbiased, and Ethical AI

To many companies, it may feel like the goalposts around AI bias and fairness are constantly moving. But there are steps you can take to encode your values in your AI, make better and fairer decisions, and reduce your organization's risks. Do an impact assessment. Educate your people and identify your bias and fairness metrics.



KEY TO SUCCESS #4:

Decluttering Your Tech Stack: End-to-End AI

It is essential that companies adopt an end-to-end AI solution that is a single ecosystem rather than separate tools that were not designed to work together. This means adopting a solution that allows businesses to be independent from the underlying technology infrastructure. This approach allows you to manage your own AI, deploying it to one of the three major cloud platforms, with the freedom to move your AI artifacts around, regardless of whether it is hosted on AWS, Azure, Google Cloud, or your own infrastructure on-premise.

SUMMARY



KEY TO SUCCESS #5: The Rise of Machine Learning Operations (MLOps)

MLOps provides a single environment for continuous integration and deployment of AI, thereby significantly reducing the complexity that most organizations face with their AI projects. This solution allows you to monitor everything, everywhere within a centralized location, regardless of where that model is hosted or how that model is deployed.



KEY TO SUCCESS #6:

Democratize to Scale

Democratization means making AI accessible to everyone in your organization. In many companies, the perception is that the AI belongs to the IT and data science teams and that only the technical people can build, manage, and understand its complexities. The reality is that AI should never be separated from outcomes that the business wants to produce. To accomplish this, you have to include the non-technical business stakeholders and ensure that the rest of the business is invested in every aspect of your AI.



KEY TO SUCCESS #7: Regulatory Compliance

Organizations putting machine learning models into production are dealing with regulatory, compliance, and corporate risk minefields that can undermine their efforts to become AI-driven if they don't proceed carefully. Even if your industry is not facing compliance rules now, it should expect to see them on the horizon soon. With a robust, comprehensive AI solution to support your business, you can navigate these new and often confusing waters with confidence.



KEY TO SUCCESS #8: Continuous Learning

In order to embrace change, you need an AI solution that is built on a foundation of agility, efficiency, and nimbleness. Continuous learning is a key component of any AI platform, helping you to keep models at peak performance regardless of external conditions, giving you the ability to refresh your models, and creating and recommending new challenger models.



SUMMARY



KEY TO SUCCESS #9: SCALABLE AI: Build vs. Buy

When considering the combined factors of customization, trust in AI, and hiring AI talent, many organizations realize that buying an AI solution or building a hybrid solution that gives them both customization options and out-of-the-box readiness saves time and helps them solve their most pressing business problems faster.



KEY TO SUCCESS #10: Al Cloud

Enter AI Cloud and augmented intelligence to help companies realize the potential of their AI initiatives. Augmented intelligence means merging the benefits of human intuition and experience with the efficiency and power of machine learning to realize better business outcomes. AI Cloud is a single system of record, accelerating the delivery of AI to production for every organization. By bringing together data scientists, data engineers, and IT and business users, AI Cloud is a unified environment that allows for collaboration and continuous optimization across the entire AI lifecycle.



As our CEO Dan Wright said in his Al Cloud launch keynote address, **"Al is a generation-defining technology** with the potential to reshape every industry, every business service, every customer interaction." We hope these points help you to realize the incredible potential of Al and serve as guiding principles for your future Al success.

DataRobot was founded in 2012 to democratize access to AI. Today, DataRobot is the AI Cloud leader, delivering a unified platform for all users, all data types, and all environments to accelerate delivery of AI to production for every organization. DataRobot is trusted by global customers across industries and verticals, including a third of the Fortune 50, delivering over a trillion predictions for leading companies around the world.

Learn more at datarobot.com