Thrive in AI disruption

Getting Started with AI: Proven Best Practices for AI Adoption

by Daniel Faggella
and Raghav Bharadwaj

Research | Strategy | Competitive Intelligence

emerj

Table of Contents

Artificial intelligence adoption is challenging. When we see our clients applying AI in an existing business, whether for AI strategy development or vendor selection, we start at the same place:

Expectations.

While there is no guarantee for AI ROI, most of the wasted money on failed AI projects can be traced back to one of the following errors:

- Incorrect expectations about what AI is capable of
- Incorrect expectations about the requirements
- An incorrect understanding about the long-term role of AI in their organization

This report is intended to handle all three of these errors in the most succinct, direct manner, preparing leaders and teams to adopt AI successfully, and to drive company strategy forward.

On top of the last five years of research and interviews, it has taken over eight months of surveys and case study analysis to put this report together. We've taken the best advice from technical AI experts, PhD-level AI consultants, and non-technical business leaders, and layered them on top of a set of principles and frameworks for effective decision-making.

Because we write for an executive audience, brevity counts, and we've boiled down what could be lengthy topics into the meaningful and critical tenets that matter. This is the same information we share with team leaders and boards in our research and advisory services, but distilled into a concrete set of guidelines for fast reference.

After reading this report, you'll be able to:

- Foresee potential AI challenges, and develop a realistic expectation of what it will take to bring an AI initiative to life
- Select ideal AI projects that match your resources, time constraints, and capabilities
- Assess AI vendor solutions quickly and effectively and make critical buy-vs-build decisions with confidence

We'll begin with framing expectations for success.

1. What to Expect When Adopting AI

One of the biggest hurdles to AI adoption and integration is a lack of proper expectations about applying AI in an existing business. Executives (and their teams) often go into the process blind because so few companies have learned these important lessons and challenges and because even fewer have successful adopted AI in a way that delivers ROI.

In the five points below, we'll explore critical insights about the challenges and opportunities of AI adoption, as well as some of the hurdles and necessary steps to expect in the process of adoption and integration.

1.1. Data

Data is essential for any AI application, but data quality and data access are often challenging hurdles to AI adoption. For established businesses, particularly businesses built before the popular use of the internet, data storage is siloed, and little thought is given to data formats. For many companies, data is stored so that it can be manually found or potentially so that it can be used for dashboards and basic business intelligence purposes.

Artificial intelligence requires that data formatting be unified and that data be accessible by AI systems, sometimes in real time. This often involves not just "cleaning" (properly formatting the data, getting rid of inaccurate or unusable data files) data but often overhauling data infrastructures and changing legacy systems altogether. For this reason, small AI pilot programs can be relatively easy to test, but sustained AI use as part of a business function implies greater changes to core IT systems, which can be costly, and uncertain in ROI.

What to expect about data:

- It's harder to access than usually anticipated
- It requires heavy organizing and sorting in order to format it, and purge unusable files
- It will often require entirely new infrastructure to implement AI sustainably

Related article: ***Fundamentals of Data Literacy***

1.2. Data Science Talent

Data science talent is critical for bringing a true AI initiative to life inside of an existing business. Relying entirely on consulting firms (many of whom are extremely understaffed in real data science talent of their own) is risky and unsustainable for true adoption of AI into any core business process.

Data science talent will be necessary for evaluating potential AI initiatives, for setting up and handling data infrastructure, and for iterating and adjusting algorithms and data to achieve a business result. This involves regular interfacing with business leadership or functional leaders in a given business department.

Data science is in high demand, and prospective employees with these skills will want more than money. They want a rewarding career working on problems that really matter in the world and really matter to their employer. They want support from leadership to get access to the tools they need, and they want functional leaders around them who understand the basic concepts of data science (not technical, just conceptual) and can work with them to solve problems.

What to expect about AI talent:

- Hiring AI talent is expensive and will require a compelling reason as to why your AI initiatives are interesting and valuable for the company. Expect to experiment with different ways of framing your data science role so that it captures attention

- Retaining AI talent means finding a productive role and tangible, important data science projects

- It is challenging to determine how a data science expert will work with other members of the IT team, and with business leadership

Related podcast: ***Why Executives Should Follow AI Trends in Business***

1.3. Teams and Leadership

While data science is critical, so is the "connective tissue" that surrounds that data science talent. We use "connective tissue" to imply the IT team, the functional business leaders, and the relevant subject-matter experts who will need to engage with the data science talent in order to solve problems.

Throughout this report, "subject-matter experts" will refer to business team members with deep expertise in the business processes (i.e. accounting) or field (i.e. petroleum engineering) that a given AI initiative is related to.

An AI project team will need to be led and/or approved by an in-house "champion", an executive leader who understands the scope and goals of the project, the estimated costs of the project. The champions must believe in the strategic value and necessity of the initiative, and be willing to support the project with guidance, resources, and enthusiasm.

These teams will need to work together in order to:

- **Determine the business problem an AI solution can solve:** Data scientists need context on what problems mean, on what data is necessary, and on what data is accessible, and IT/business team members will need to help.

- **Determine how AI can solve that business problem:** Data scientists will have ideas about how AI might detect fraud or might predict events for business leaders, but business/IT team members will need to be present to guide or contribute to the project, including a clarification of the output of the interface. This clarification might include answering questions such as, what will it "look like" when it's done? How will the system act? What actions will the system take automatically, and which need human oversight?

- **Maintaining and updating the system once it's built:** Constant iteration is needed in order to make AI systems deliver results. Calibrating a system over time not only involves technically tweaking the data and algorithms but using business context to determine whether the system is working. This is not a job for data scientists alone.

It should be noted that while data scientists and executives are more frequently involved in strategic planning conversations, it is often hard to pull in-house subject-matter experts from the core tasks that they are employed to do. Leadership must either (a) restructure the schedule and incentives of the involved subject-matter experts, or (b) hire outside subject-matter expert consultants to work in a dedicated fashion for the length of the project.

What to expect about teams and leadership when adopting AI:

- Initially, IT and business employees will subtly resist taking time away from their regular work to interact with data scientists
- Expect to reinforce the importance of these interactions and to make up for the lost time in core IT and business work
- Expect to develop new playbooks and workflows for how AI projects are initiated, measured, and kept up over time

Related article: ***The Critical Role of Subject-Matter Experts***

1.4. Iterating

Getting a machine learning system to work in business isn't easy. After accessing the necessary data, building a team, and selecting the right projects, there is more work to be done. AI is not like "doing IT;" it is like "doing science." There is no plug-and-play method to make AI come to life: only tests.

Some applications involve wholly trained systems, such as security AI applications that detect people or vehicles in a live video stream. Such an application involves a pre-trained model that handles a problem that won't change much (the appearance of humans and vehicles).

Most business problems involve handling bespoke business data in messy formats and taking educated guesses about what data and what algorithm might get a result. Months of testing and tweaking are often required to get an AI system to perform on par with humans or previous software systems.

Even when a system is built, there must be ongoing attention paid to the outputs of the system, ensuring that odd trends of faulty data haven't thrown off the system's results. This involves both measurement of the system's results and tinkering with the data and the algorithms. This ongoing process is usually nowhere near as challenging as the initial process of setting up an AI application.

What to expect in terms of iteration and calibration of AI systems:

- Expect that systems will require significant trial-and-error in order to approach human-level performance or an acceptable level of performance
- Expect this even after using new data, new methods, or new hypotheses
- Expect some level of ongoing measurement and maintenance of the live machine learning system

Related interview: [Practical AI Ethics - Impacts on the Bottom Line]()

1.5. Deployment Challenges

AI is new, and bringing it to life in an existing enterprise often involves a learning curve. A variety of issues might come up to prevent or delay an AI system from being deployed, including:

- **Legal or regulatory concerns:** Heavily regulated sectors (like healthcare and finance) must ensure that their AI applications aren't in violation of any laws. In-house legal staff should sign off on an AI initiative early on, rather than waiting until deployment. Other applications such as recommendation engines or targeted marketing applications might stir a PR nightmare, and last-minute executive concerns on these matters can often be prevented by planning for and considering these impacts from the outset of an AI initiative before sinking resources into it.

- **Data access:** Getting a pilot AI application to work is often done with stagnant older data, just enough to test the concept. Deploying such an application often involves ongoing access to data (in real time or in relatively regular intervals). This sometimes requires manual and monotonous human effort (not ideal, but sometimes necessary) or an overhaul of data infrastructure built for data access and compatibility with future AI plans. This makes deployment especially challenging in older enterprises.

- **Losing a champion:** Not every C-level executive or functional business leader will be excited to experiment with AI. In addition, even those who are excited often don't have the proper context on AI's capabilities or understanding of its requirements. If the right champion of a project moves to another department or another company, the funding and resources needed for an AI initiative may disappear as quickly as the champion themselves. Securing a committed, long-term connection to the C-suite (or whichever party holds the budget) is critical.

What to expect when it comes to deploying AI:

- Expect hurdles to arrive between starting and deploying an AI initiative
- Prepare to record these hurdles and determine a playbook to move beyond them, starting with some of the simple advice in this report

emerj

The Difference Between "IT" and "AI"

	Information Technologies (IT)	Artificial Intelligence (AI)	Insight to Reinforce
Time to Deliverable	Hard to predict with new projects. Relatively easy to predict with repeat or well understood projects.	Nearly impossible to predict with new projects. Still very challenging to predict with repeat or well understood projects.	Only apply AI where it can be viewed as a long-term investment in the company's skills and capabilities -not a short-term ROI.
Product Capabilities	Relatively easy to deliver on. Can be hard- coded so that the system behaves predictavly, and can do anything we program into it.	Challenging to know ahead of time. The data and algorithms may not be able to deliver on the capabilities we originally wanted.	AI systems are probablistic, and most applications - especially novel applications - must be considered "experiments."
Current In-House Competence	Relatively strong. Most companies have competent in-house IT skills.	Weak. Most companies lack data science talent, and their business leaders lack an understanding of data science concepts.	AI often requires a significant effort in data science talent, and in familiarizing business leaders with data science concepts and use-cases.
Monetary Resources Required	Hard to predict with new projects. Relatively easy to predict with repeat or well understood projects.	Nearly impossible to predict with new projects. Still very challenging to predict with repeat or well understood projects.	Only apply AI where it can be viewed as a long-term investment in the company's skills and capabilities -not a short-term ROI.
Role of Subject-Matter Experts	Needed primarily to define the functions or capabilities of the IT solution being developed.	Needed to define the business problem, to determine the data required to solve the problem, and to determine the outcome.	AI often requires a significant effort data science talent, and in familiarizing business leaders with data science concepts and use-cases.
Data Requirements	Minimal or none.	Potentially robust data requirements. May involve overhauling data infrastructure or data formatting. May require months of prep.	AI adoption often requires systematic change in the way data is handled and managed.
Depth of Integration	Can often be layered on top of existing IT systems.	Often must be integrated into multiple IT systems, accessing data in specific formats, and drawing inputs from many sources.	Only apply AI where it can be viewed as a long-term investment in the company's skills and capabilities -not a short-term ROI.
Strategic Value	Critical in the short term.	Critical in the long term.	Only apply AI where it can be viewed as a long-term investment in the company's skills and capabilities -not a short-term ROI.

2. Viewing AI as an Investment

There are very few quick wins in artificial intelligence. "Plug and play" applications do exist in some sectors and for some very controlled use-cases, but these simple, pre-trained systems do not deliver one of the most important benefits of an AI project: the critical skills and capabilities that a company enables for its future.

In the four sections below, we'll explore the importance of viewing AI as an investment in long-term capabilities and practical steps to prepare teams for creating both near-term and long-term benefit in critical AI initiatives.

2.1. Proper and Improper Views of AI

In our work with public and private sector executives, we find that AI is almost always viewed through one of four perspectives:

- **An automation upgrade:** AI viewed as just another kind of IT that can be "plugged in" to improve efficiency and save on margins
- **A silver bullet:** AI viewed as an instant upgrade to any business, improving any business metric
- **A hype bubble:** AI viewed as a skill for startups or silicon valley but as being wholly irrelevant to a business' processes and sector in general
- **An investment:** AI viewed as a new set of capabilities to be enabled over time to not only improve efficiencies but to rethink core processes and win long-term market share

Businesses should develop an "investment" perspective within their executive teams and departments. The "investment" perspective involves a number of critical insights:

- Enabling capabilities with artificial intelligence requires new ways of managing data and talent, and it will take most businesses years to develop the processes and habits required to enable these capabilities
- Data access and team agility will allow for new kinds of efficiencies and new ways to service customers, most of which are not yet discovered. Being ready to adopt these new capabilities as they emerge and to develop new capabilities ahead of one's peers is a strategic investment
- Business leaders should not exclusively focus on near-term AI ROI, though they should budget carefully and responsibly. Leaders should see value in the improved understanding of AI concepts, experience working with data, and the organization of an AI project team. Improved familiarity with applying AI concepts is an investment in and of itself

Practical steps to encourage an "investment" AI culture include:

- **Set clear expectations about how success will be measured in an AI initiative.** Teams should have an understanding of the challenges ahead and the way a project will be measured. They should also be aligned with the view that familiarity with data and AI projects is a benefit (in essence, a "team upgrade," not just a "tech/process upgrade").

Without this long-term perspective from the top, team members (especially data scientists) are more likely to look for short-term, surface-level improvements than the genuine transformation that is often required to bring a functioning AI application to life.

- **Record lessons learned in working with data and in team collaboration on AI projects.** This will make a tangible example for the team that learning is a "win" to be excited about. In addition, these lessons will often be as (if not more) valuable than the initial results of AI initiatives and experiments. Building on this skillset and AI familiarity is critical.

Related article: ***[Enterprise Adoption of Artificial Intelligence – When it Does and Doesn't Make Sense]***

2.2. ROI Uncertainty

The vast majority of AI applications today are without strong evidence of ROI. In-house and vendor solutions often fail. The challenge of applying AI to an existing enterprise is often compounded by a lack of data access, lack of data quality, lack of data science talent, and lack of experience working with AI.

While some leading companies or particularly innovative firms will have the resources and time to experiment with wholly new methods of applying AI in business, most companies will not have the extensive budgets or risk tolerance to do this. Most firms will allow their Fortune 500 counterparts to take up much of the novel and experimental work and instead act as fast followers on AI applications in their sector that seem to be garnering strong evidence of ROI.

While young, digitally-native sectors like eCommerce and online media will proliferate AI use-cases rather swiftly, it may be 5 or more years until AI is genuinely accessible and easy to use in older sectors like finance, healthcare, and manufacturing. It shouldn't be suspected that AI will always be a loss, but near-term ROI isn't something businesses should count on. The accrued benefits of organizing data and building AI-capable teams is the larger payoff at present. Properly vetting vendor solutions is one step to having the best chance at AI ROI.

Practical steps for acting in ROI uncertainty and how to vet AI vendors:

- Do not assume that a vendor application that worked in one sector, such as retail, will work in another, such as banking. Ask vendors for examples of when their solution worked for a client like your business with a problem like yours. Do not assume transferability.

- Look for vendor companies with strong AI talent on their leadership teams, either (a) advanced degrees in computer science or hard sciences from reputable schools, or (b) robust technical experience with AI at marquee tech companies, such as Google and

Facebook. "AI" firms with neither kind of talent on their leadership team are often lying about AI being a part of their product at all.

- Buyers with an interest in ROI should work with vendors that have raised more money. These firms are more likely to have a strong solution with existing traction, and they're less likely to "pivot" their application wildly in three or six months, which many young companies are forced to do as they find a fit for their product in the enterprise.

Related article: ***[Predicting the ROI of AI – Pitfalls to AI Adoption in the Enterprise]***

2.3. Ease of Deployment

Different AI applications can vary vastly in terms of the cost and time it takes to get an application up, running, and delivering an ROI.

The Ease of Deployment Score is a tool that we use to compare vendor product offerings side-by-side to determine applications that are relatively easy and relatively difficult to deploy.

This tool can be used to assess limited amounts of information about a variety of vendors and develop an imperfect but extremely useful indicator. We recommend using this Score when deciding on which vendors to shortlist for demos.

We grade vendor product offerings on 3 factors to generate their Ease of Deployment Scores:

- The amount of data required to use the product
- The difficulty of the product integration
- The level of data science talent required to use or maintain the product

Data Requirements

Begin by answering the following questions:

- What kind of data does the software require to "learn?"
- What challenges do you face in data management when implementing the software?

We grade vendor product offerings on the following scale to generate their Data Requirement Scores:

- **5:** The product requires no data at all

- **4:** The product requires data to train the algorithm behind it, but the client doesn't need to provide that data

- **3:** The client needs to provide some enterprise data, such as transaction and customer data, to train the algorithm
- **2:** The client needs to provide some enterprise data and external data to train the algorithm

- **1:** The client needs to provide a large amount of enterprise and/or external data for the initial integration of the product

- **0:** The client needs to provide a large amount of enterprise and/or external data for the initial integration of the product and on an ongoing basis to maintain the product

Integration Requirements

Begin by answering the following questions:

- Does the implementation involve a complicated deployment or a complicated "instrumenting" of your data environment?

- How is your data extracted and managed from your internal databases?

- How is the output from the AI system consumed?

We grade vendor product offerings on the following scale to generate their Integration Requirement Scores:

- **5:** The client can use the product "out of the box" or the product is "plug and play."

- **4:** The product requires minimal setup via, APIs, HTML5, SDKs, etc.

- **3:** The product does some amount of data preprocessing by itself

- **2:** The product is an automated machine learning platform

- **1:** The client requires in-house data scientists for the initial integration of the product

- **0:** The client's in-house data scientists and subject-matter experts need to collaborate extensively to integrate the product

Data Science Requirements

Begin by answering the following questions:

- How much in-house talent will you need to integrate a system like this?
- What would be the amount of ongoing data science upkeep needed for this application?

We grade vendor product offerings on the following scale to generate their Data Science Requirement Scores:

- **5:** The client doesn't require in-house data scientists to integrate the product
- **4:** The client requires in-house data scientists for small parts of the integration
- **3:** The client requires 1 - 2 data scientists for the full integration
- **2:** The client requires a small team of data scientists for the full integration
- **1:** The client requires a large team of data scientists for the full integration
- **0:** The client requires a large team of data scientists for the initial integration and the continued use of the product

2.4. "Quick Wins" vs. Critical Capabilities

Some AI applications require essentially none of the skills and infrastructure that most robust AI applications involve. For example:

- A natural language processing (NLP) vendor asks for the patterns you'd like to discover in your customer emails or in online forums and social media. The vendor then uses their software to analyze connections between terms and phrases, delivering a report to the client. This doesn't require real-time data access nor in-house NLP skills. As a result, very little improved familiarity with AI will result from the project.
- A computer vision vendor detects theft behaviors at the checkout counter of a grocery chain. The camera streams are simply fed into the vendor's algorithm. The client may have an interface to analyze the instances of theft detected by the camera, but the client doesn't need to manually train the algorithm or have any in-house data scientists to manage or use the product itself.

This varies drastically from more robust AI implementations. For example:

- A vendor offers a system to detect money laundering at banks. Banks that adopt the solution must integrate their client data and financial data from across georegions. The solution requires calibrating this unique mix of data to find unique patterns of fraud that can be reliably found over time. This training requires context from the vendor, from the client's fraud experts, and from the client's own in-house data science staff. An application of this kind cannot be successful without the client's in-house data science and without robust integration into legacy banking systems. This is a very difficult challenge.

The "quick win" applications have the benefit of not requiring in-house data science talent and not requiring much integration or complex use of data. The downside of these applications is that they don't genuinely develop AI skills and abilities within the business. These applications are more likely to be successful because they are consistent use-cases.

Robust, complex AI solutions are just the opposite. They require expensive in-house data science talent, integration with existing IT systems, and potentially an overhaul of how data is stored and organized. Companies learn many lessons about AI capabilities by applying such complex solutions, but these solutions are more likely to fail to produce ROI because they are much more experimental and bespoke.

How to balance self-contained applications and critical capabilities:

- Businesses with limited budgets or without in-house AI talent may choose to go with self-contained applications. Even in these cases, AI should be the right tool for the job; AI should not be adopted for its own sake.

- Businesses interested in building long-term capabilities within their teams will need more budget, more patience, and more risk tolerance in order to genuinely transform key business processes. Only businesses with budgets, talent, and risk tolerance should make serious steps towards robust AI overhauls of business functions.

- These criteria may mean that a company chooses not to apply AI in the very near-term. There is nothing wrong with this decision if it is well-grounded in the business's needs and in realistic expectations of artificial intelligence.

2.5. Defining Critical Capabilities

Part of the challenge of AI is that much of its greatest benefits and advantages are unknown to us now. Much like the internet in the 90's, AI holds the promise of transforming businesses drastically but in ways that we cannot now predict.

Critical capabilities are the skills, resources, and culture involved in successfully leveraging artificial intelligence to its full potential. To go a bit deeper:

- **Skills**
    - Data science skills and talent
    - Using teams of data scientists and subject-matter experts to solve problems
    - An understanding of data science concepts and basic AI applications in non-technical team members, not just data scientists

- **Resources**
    - Harmonized standards for required data
    - Access to required data
    - In-house playbooks and guides for handling different data and AI challenges, retained knowledge

- **Culture**
    - A willingness to accept risk, to iterate
    - Valuing data
    - Valuing cross-team collaboration

The list of factors above will help a company adopt current AI use-cases and will prepare them to take advantage of future use-cases and innovations as well. It is these skills, resources, and elements of culture that should be viewed as an investment, as future-enabling factors for AI, and for team and tech agility.

Related interview: Table Stakes AI Insights for the Enterprise - with Vlad Sejnoha, former CTO of Nuance Communications

3. AI Adoption Motives

3.1. Poor Motives

The following adoption motives are not advisable, but are remarkably common:

- **Following press releases:** Many businesses simply aim to adopt the AI applications that they see amongst their competitors. This is generally a bad idea because competitor companies only reveal the AI applications that they want to reveal. Revealed applications are not necessarily high ROI or high priority. Rather, they are often revealed (in the form of an announcement or press release) simply to impress customers or investors, without any consideration for the strategic value of the technology.

- **Toy applications (AI for AI's sake):** "What can we do with AI?" This is a well-intended question, but it steers companies in the wrong direction. It frames AI like an IT solution that can be "plugged in." AI is a long-term investment that should be aligned with company goals and with an understanding of AI's capabilities broadly, not as a knee-jerk effort to "do AI" for its own sake.

3.2. Proper Motives

PhD AI consultations and AI experts we've interviewed often highlight the following AI adoption motives as appropriate:

- **AI is the best tool for the job for a critical business process:** Assess all existing solutions to a problem. From in-house manual effort, to outsourcing, to software, to AI. If a job is genuinely best-suited to be solved by artificial intelligence and you have the budget, team, and data to do it, then AI can be a right choice.

- **AI will be critical for future strategy:** If your own company strategy will be greatly enhanced by an application of artificial intelligence, then it may be a viable reason to adopt. For example, an eCommerce business whose 5-year strategy involves becoming the best at right-time marketing campaigns through email and social media may require AI to achieve that strategic aim.

- **AI capabilities are inevitable in a specific business function:** Some business functions will inevitably involve artificial intelligence or machine learning. For example, payment fraud and cybersecurity are business functions that cannot possibly be handled by hard-coded rules, and there is a consensus that these applications will shift inevitably more and more towards AI solutions. A business with the budget, team, and data infrastructure to make the leap into AI now may do so in order to stay ahead of the curve.

4. Timing and Purchase

4.1. When to adopt AI

A variety of factors come together to contribute to the timing of AI adoption from budgets, to data, to proper expectations of the technology.

Here is a short checklist for AI adoption:

- Ensure that functional leaders have a grasp of AI fundamental concepts and use-cases
- Ensure that functional leadership can sustain the risk of the program failing. Ensure that they can view AI skill building and problem solving as an investment in their teams
- Ensure that AI is an appropriate tool for the job. Ask data scientists if the use-case in question is viable and realistic given the data and resource constraints at play
- As a rule of thumb, ask yourself: have your largest competitors begun experimenting with AI? If the answer is a clear "yes," then AI is more likely to be realistic in your sector or for your business. If the answer is "no," AI likely needs to gain gain initial traction and find a proven use-case in your sector

4.2. Pre-adoption Considerations

Ask the following questions as a way to frame expectations and properly prepare for contingencies:

- **How will success be measured? Cost per 'X'?** Speed of completing a process? Comparison against current human or software performance?
- **How will the business process look if this implementation goes well?** It is best to confirm this forward-looking vision not just with subject-matter experts in your company, but with data scientists who help set expectations for a realistic resulting solution.
- **How will existing talent and resources be used?** What will you do with the legacy system that you no longer use or use less of because of this AI implementation? How will daily workflows change for your team members? If team members are displaced, where can you still use their skills within the company?
- **What will the PoC, incubation, and live phases look like?** We discuss this further below.

4.3. Build-or-Buy Considerations

In this section, we address the relative advantages and disadvantages of buying from an AI vendor and building AI solutions in-house.

None of the factors listed below should be used in isolation but rather, should be seen as factors worth considering when leaning in the direction of "build" or "buy":

Do you have the in-house data science talent to execute on this AI initiative?

- **No:** You may be more likely to require a vendor for help.
- **Yes:** You may be able to build the solution yourself.

How mission-critical is this AI application to your competitive strategy?

- **No:** You may be more likely to use a vendor application.
- **Yes:** You may be heavily incentivized to build the application in-house. Owning the IP may be essential for any AI application that will be pivotal to your strategy moving forward. Any B2B AI vendor that sells AI products or services cannot possibly rely on another vendor to provide it's key IP, it's product. Similarly, a business initiative that holds up a company's entire future competitive strategy will likely have to be built in house.

Do you have the data required to execute on this AI initiative?

- **No:** You must either begin collecting the data now in order to build the solution in-house, or, if you cannot wait or don't have access to the right kind of data, you should work with a vendor firm. Some vendor firms have pre-trained models that might bolster your lacking data stores and provide functionality that a self-trained system cannot offer.
- **Yes:** You are more likely to successfully build an AI product in-house.

Do existing vendor solutions address this outcome directly?

- **No:** It is more likely that you will have to build your solution in-house. Vendors are very likely to say that they can stretch their system to suit any problem, but this is a dangerous proposition for a buyer. Being an experimental, unique use-case for a vendor can be costly and may well result in a vendor deciding that such an edge-case is not something they can support in the long run, even if it is critical for your business operations.
- **Yes:** It may be worth considering the use of a vendor to solve your specific business problem.

Do you need this application deployed swiftly?

- **No:** You may have more time to develop and test a robust solution in-house. Alternatively, you also have more time to find a vendor.
- **Yes:** It may be best to find a vendor with a strong track-record in solving your exact problem, in your exact position.

Might this application require greater flexibility in the future?

- **No:** You have the ability to either build in-house or buy a solution.
- **Yes:** You have a strong motive to build in-house if possible because you're unlikely to get a vendor company to bend their solution to your specific needs.

Artificial Intelligence Buy / Build Decision Guide

Critical Question	More Likely to Build If	More Likely to Buy If
Do existing vendor solutions address this outcome directly?	No	Yes
Do we need this application deployed swiftly?	No	Yes
Do we have the data required to execute on this AI initiative?	No	Yes
Might this application require great flexibility in the future?	Yes	No
How mission-critical is this AI application to your competitive strategy?	Yes	No
Do we have the in-house data science talent to execute on this AI initiative?	Yes	No

Related interview: [A Framework for Long-Term and Near-Term AI ROI, with Microsoft GM of AI David Carmona]()

5. The 3 Phases of AI Adoption

5.1. Proof of Concept Phase

A proof of concept is intended to determine whether an AI application could potentially deliver a specific goal or outcome. It is not intended to replicate an actual production environment. It is merely intended to prove that the AI approach is potentially valuable.

This phase is important because, frankly, many AI approaches are purely experimental. Determining whether such approaches, whether for fraud detection, for recommendation engines, or whatever else, could be used on the specific data and specific business problem of a specific business is uncertain.

Often a proof of concept is run in a "sandbox," a controlled environment, with historical data, without any actual integration into core business systems, and without any actual security concerns.

The vendor, if one is involved, is often doing the project themselves without significant involvement from the in-house data science talent of their client.

There are no end-users to please, no day-to-day operations that can be interrupted or challenged. If the vendor has run many such sandbox tests, they may be able to confidently set up and run this very controlled test. If the vendor is solving a new problem or working with new kinds of messy data, they may be learning as they go. There may be no walk-through of the user interface for the client; the interface and work is often being done by a vendor.

The aim of a PoC is to determine some measurement of performance, for example:

- False-positive, false-negative, true-positive, and true-negative rates
- Time to results (per the processing of one loan application, for example)
- Cost-per-item (for a chatbot, this might be per reply or per satisfied user query)
- Ease of use

By coming up with some kind of standard measure, a company can compare vendors (or a vendor versus in-house approach) side by side. A company can also look at the time it takes to complete the process.

Goals to achieve in the proof of concept phase:

- Determine ways to measure the success of an AI application

- Determine whether your data and the AI approach you have in mind could viably solve a problem

- Compare vendor applications or in-house approaches to determine the application with the greatest promise

5.2. Incubation Phase

With traditional IT, a proof of concept can often provide enough insight to determine whether or not a solution should be integrated fully into a business. Because of the iterative nature of AI (calibrating data and algorithms to achieve a goal) and because of the novelty of AI workflows and skills, an incubation phase is often required in order to flesh out the challenges and potential of the technology.

Goals to achieve in the incubation phase:

- Determine how to properly train teams on the workflows needed to manage the system properly

- Determine how to manage this AI process (daily, weekly, monthly), including the team members and resources required

- Determine a success criterion of how the system will be measured. There should be a threshold of agreed-upon metrics that take the project from incubation to integration

- Security-test the system thoroughly to ensure that it is safe to go live

- Agree to a support model with the vendor if you worked with one

Often the goals of the incubations will be tailored to your unique business and business problem. Relatively mature vendor companies will have proficient strategies for proving a proof of concept, but they are unlikely to have a succinct process for the incubation phase. In many cases, the vendor companies will be learning as much in this phase as the client.

Questions to ask in the incubation phase:

- How will this work into the workflows of your teams? Who will do what to keep the system running regularly?

- How do you collect new data? Who tags it? Where do you put it?

- What do you do if the model seems to be wrong?

This incubation period will fade as vendors and use-cases become more capable and as teams become more used to data science and ML systems, but for the foreseeable future, incubation is necessary for the integration of most robust AI systems.

5.3. Integration Phase (Going Live)

Integrating artificial intelligence into an existing business environment is challenging but can be improved by having the right set of expectations, as we've tried to cover throughout this report, and by going in with the right set of questions.

Questions to ask when moving an AI application from incubation to integration:

- Do you have a team leader or champion who understands the basics of artificial intelligence and is willing to take risks and iterate?
- What team members will need to be involved in keeping this AI systems working?
- What additional in-house talent or consulting help do you need in order to implement and run this system?
- How do you need to prepare these team members to work together and change their current workflows?
- What are the additional security considerations you should take into account before going live?

The Three Phases of Artificial Intelligence Adoption

	(1) Proof of Concept	(2) Incubation	(3) Integration
Goals	Determine a way (or ways) to measure the success of an AI application. Determine whether our data, and the AI approach we have in mind, could viably solve thebusiness problem. Compare vendor applications or in-house approaches to determine the application with the greatest promise.	Determine how to properly train teams on the workflows needed to manage the system properly. Develop a "run Book", determine how to manage this AI process (daily, monthly, etc) – including the team members and resources required. Determine a success criterion of how the systemwill be measured. There should be a threshold of agreed-upon metrics that take the project from incubation to integration. Security test the system thoroughly to ensure that it is safe to go live.	Upgrade teams to data science talent as necessary to achieve our desired result. Upgrade or overhaul data infrastructure in order to consistently provide our application with data in the right formats. Update "run book" of processes and procedures (expect a large number of updates within the first 3-6 months post-integration).
Criterion for Progression to Next Phase	Measurement(s) of success are determined. The application proves itself to perform at our determined standard within an isolated, "sandbox" environment.	Processes and workflows are determined for our teams, with a plan on how to get them started. The ongoing processes seem viable and cost-effective for the result we expect to achieve. Buy-in from users/operators of the AI solution.	(Not Applicable)
Vendor Considerations	Explore the vendor's understanding of your business problem. Compare vendor solutions directly via the determined success criterion of the PoC.	Agree to a support model in place with the vendor, make ongoing needs and support clear, and tailored to your unique business needs.	Stay in touch with vendors as necessary. There maybe benefits in keeping a strong communication between vendor AI experts and in-house AI experts, to transfer knowledge and glean lessons learned from the vendor's other client integrations.

Emerj Artificial Intelligence Research

Survive and thrive in an era of AI disruption

We work with business and technology leaders to develop artificial intelligence strategies that win in the competitive marketplace.

Core Services

Market Sizing and Growth Rate Forecasts

Measuring the impact of AI on specific business sectors and lines of business. The sizing and forecasting data that leaders need to confidently make strategic decisions. Learn more.

Vendor Landscape - AI Capability Space

Assess the full range of AI capabilities for your business or sector with a robust scoring of the full AI vendor landscape, separating hype from reality by determining the credible applications with real evidence of ROI. Learn more.

Competitor Analysis - AI Capability Space

Quickly measure your firm against industry leaders and direct competitors by assessing the AI initiatives of companies in your sector, with a clear view on the strengths and weaknesses of their strategies. Learn more.

AI Business Strategy Process

An immersive two-day on-site process to determine a coherent, powerful strategy around the AI capabilities most aligned to company goals and strengths. Learn more.

Other Products and Services

- Syndicated research reports
- Advisory and consulting
- AI vendor selection services
- More…

Contact Research Services

services@emerj.com
1-617-945-8567

Daniel Faggella, CEO at Emerj

Called upon by the United Nations, World Bank, INTERPOL, and many global enterprises, Daniel is a globally sought-after expert on the competitive strategy implications of AI for business and government leaders.

Daniel helps organizations navigate the competitive landscape of AI capabilities, determine high-ROI applications that match and organization’s strengths, and build AI strategies that win.

In addition to his advisory work with leaders, Daniel has interviewed thousands of AI researchers and founders, and his research and reports are cited by Harvard Business Review, the World Economic Forum, and other leading publications.

Daniel has been devoted to studying the consequences and applications of AI since graduating from UPENN with a master’s degree in cognitive science. He lives in Boston.