

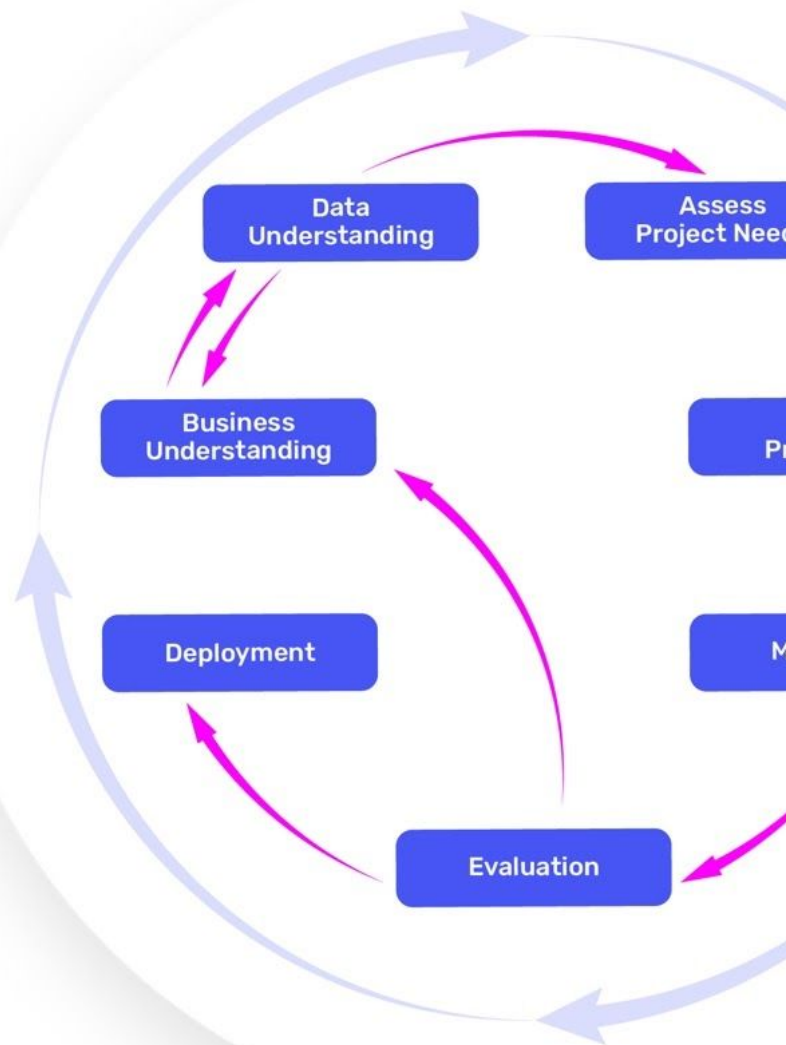


Special Edition for
**IT and Strategy
Consultants**

Thrive in AI disruption

AI Deployment Roadmap

An End-to-End Guide to the AI Lifecycle
in Business



by Daniel Faggella

Emerj Artificial Intelligence Research

Emerj Artificial Intelligence Research helps global organizations develop AI strategies and initiatives that win in the market. We map the capability-space of AI across major sectors, with a finger on the pulse of academia, Fortune 500s, and the global artificial intelligence startup ecosystem.

“We help leaders survive and thrive in an era of artificial intelligence disruption.”

We create cutting-edge AI impact research, inform executive leadership, and make important contributions to important decisions around governance, innovation, and strategic planning. We’re called upon by many of the largest and most reputable organizations in the world:



Our research focuses on three critical aspects of AI capabilities:

- **Applications (“What’s Possible?”)** – Examining the landscape of AI applications, open-source tools, and use-cases that might solve organizational problems.
- **Implications (“What’s Working?”)** – Determining the use-cases with a genuine track-record of ROI, and determining the integration costs or ROI potential.
- **Plans (“What to Do?”)** – Informing strategy by honing in on the AI trends or capabilities most likely to deliver the desired results or the organization.

Through our [Research Services](#) and [AI Business Strategy Process](#), we help clients win market share and make more profitable decisions.

Contact Emerj

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



© 2019 Emerj Artificial Intelligence Research. Unauthorized copying or distributing is a violation of copyright law. - citations@emerj.com or 1-617-945-8567

Table of Contents

1. Prerequisites to AI Deployment
 - 1.1. Critical Capabilities
 - 1.2. In-House Data Science Talent
 - 1.3. Contextual AI Knowledge at Leadership Level
 - 1.4. Strong Digital Infrastructure
 - 1.5. Evidence of Adoption from Larger Counterparts
 - 1.6. IT Services and Consulting Considerations
2. AI Deployment as a Living System
 - 2.1. IT Services and Consulting Considerations
3. The 3 Phases of Deployment
 - 3.1. Proof of Concept (PoC)
 - 3.2. Incubation
 - 3.3. Deployment
 - 3.4. IT Services and Consulting Considerations
4. The 7 Steps of the Data Science Lifecycle
 - 4.1. Business Understanding
 - 4.2. Data Understanding
 - 4.3. Assess Project Needs
 - 4.4. Data Preparation
 - 4.5. Data Modeling
 - 4.6. Evaluation
 - 4.7. Deployment
 - 4.8. IT Services and Consulting Considerations
5. Putting the Steps and Phases in Action
 - 5.1. Preparation
 - 5.2. Ongoing Consideration
 - 5.3. Retaining Enterprise AI Competence
 - 5.4. IT Services and Consulting Considerations



1. Prerequisites to AI Deployment

Before exploring the phases of the AI deployment, or the steps involved in the AI lifecycle, it's important to note that not all companies are prepared to adopt artificial intelligence.

While some readers of this report will be able to apply AI immediately within the companies they work for - others will be able to use this guide to properly prepare for AI before making the leap.

With billions being wasted on fruitless AI initiatives and failed proof-of-concept projects, it is valuable for team members to know when *not* to apply AI - and what to work on instead so that it's value can be realized.

While there may be nearly a dozen such prerequisites we could list, these three are fundamentally among the most important to consider:

1.1. Critical Capabilities

Artificial intelligence requires more than data scientists - it requires skills, resources, and culture that we at Emerj refer to as *Critical Capabilities*.

Below is a short passage from our complete report titled [Getting Started with AI - Best Practices of AI Adoption](#), which highlights the three broad categories on *Critical Capabilities*:

- **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.

For most companies, AI adoption is harried and rushed, and neglects all learning and workflow evolution in the pursuit of short term financial objectives. This is a major cause of AI's failure in established enterprises.

Companies who approach AI intelligently, with an eye to long-term value and avoiding the waste of resources in the near-term, will think about AI adoption in roughly two phases:

1. Assessing / Establishing a Critical Capability Foundation
2. Deploying AI While Continuing to Build Critical Capabilities

Most firms - even firms who are wholly unprepared to adopt artificial intelligence, with essentially no knowledge or in-house talent, will skip the first phase entirely, and waste both time and money on projects bound to fail. While the process of developing capabilities never ends, there is a baseline foundation of Critical Capabilities that a firm needs to achieve before even seriously considering AI adoption at all.

In the sections 1.2 through 1.5, we'll highlight some of these foundational capabilities that should generally be considered as firm prerequisites to considering AI adoption:

1.2. In-House Data Science Talent

Without in-house data science expertise, it is near impossible to get real value from an AI application, and it is also near impossible to rank or assess potential applications and vendors.

AI consultants have a useful place in the AI ecosystem, but they cannot be relied on as a crutch for a lack of talent. AI requires new ways of working with data, and entirely new



skillsets (iterating with algorithms, harmonizing data in new ways, etc) that require in-house skills.

As a rule, any company (potentially, any department) should have a data science leader with (a) a robust background in applying data science to actual business problems, and (b) a strong grounding in the subject-matter expertise of the business or department in which they are stationed.

It is widely understood (but rarely talked about) that business leaders without any data science experience in the room are responsible for the vast majority of failed AI proof-of-concept projects.

Engagements with new vendors, AI-related consultancies, or potential AI talent for hire - should all be screened through an in-house expert who speaks data science as a language - and someone incentivized to have the interest of the business (not the vendor) at the core of their decision-making.

1.3. Contextual AI Knowledge at Leadership Level

Arguably more important than in-house data science leadership - at least in the earliest phases of AI consideration and adoption - is to have business leadership with a strong contextual understanding of AI.

By that I mean an understanding of roughly three areas of knowledge:

How AI “Works”

A business person should have an understanding of what artificial intelligence is, and of how it can be applied to business problems.

Combined with use-case knowledge, this basic understanding of what AI is and how it works allows a nontechnical professional to look at a problem and have an intuition about the following questions:

- Is AI the right tool to apply to this problem?
- Do we have enough data and expertise to apply AI to this problem?

Basic AI Terms and Components

Understanding detailed linear algebra terms or the nuance differences between this and that kind of support vector machine isn't necessary for 99% of nontechnical folks (myself included). That said, some basic terms deserve a strong conceptual understanding, including:



- Artificial intelligence and machine learning
- Supervised vs unsupervised learning
- Computer vision
- Natural language processing

The point isn't to be able to tell data scientists what algorithms to use, or to understand Python code when you look at it. The point is to understand how AI works – and why your own AI needs might be, by understanding fundamental AI approaches and concepts.

Related resource: Visit the [Emerj AI Glossary Terms page](#) for succinct definitions of AI terms, along with business use-cases to showcase their application.

Representative Use-Cases

Use-case familiarity with AI is crucial. While a conceptual understanding is the right foundation, it's important to be aware of a wide range of current AI applications within your sector – or even related sectors.

In order to understand the actual ROI and adoption of these different use-cases (and to separate the hype from reality) – it's important to examine use-cases from a few different vantage points:

AI Vendors

- Who are the leading AI vendors in this sector?
- Which vendors have the most credible and numerous case studies?
- Which vendors have raised the most venture money – and do these well-funded firms have anything in common?

Enterprise AI Applications

- Where are the Fortune 500 firms in this sector investing in AI?
- What applications are being built, and which are being bought?

Capabilities

- Across the AI applications in this sector, what are the new capabilities being enabled (i.e. the new actions or results available to the business, which are enabled primarily by AI)?

Functions

- Which business functions within this sector (sales, customer service, compliance, etc) are being impacted because of new AI functionality – and how?



The most value that I offer to the table as a nontechnical AI strategist and advisor is an understanding of what AI does, and where it's really working, in other similar business situations. This is what we do with our [AI Capability Map services](#) – but it's possible for business people to understand some of these basic ideas with their own independent research on Google.

Related resource: Use the general search bar or the menu navigation on [Emerj.com](#) to search for AI use-cases within your sector. From insurance to pharma to banking and more, our publicly available articles span across industries and geographics, and are a good starting point for business leaders who want to obtain use-case familiarity quickly.

1.4. Strong Digital Infrastructure

If a firm is burdened by inefficient legacy systems and paper processes, AI will be more of a headache than helpful. Get to a reasonably strong position of digital transformation, and worry about AI when company data is accessible and the company is agile enough to embrace it.

Firms whose digital transformation is in it's very early stages are often in no position to take advantage of AI, because:

- a) Data is impossibly hard to access.
- b) There has never been a need for in-house data science talent, contextual knowledge of AI from leadership, nor a culture that values data itself.

Firms that are woefully behind in digital transformation can benefit from understanding the prerequisites to AI, and in building their *Critical Capabilities* to a point whereby AI becomes more accessible.

1.5. Evidence of Adoption from Larger Counterparts

If the companies 5 to 10 times larger than your company has no evidence of AI adoption or traction, it is extremely unlikely that adopting AI in your company will be fruitful. Applications need large, well-resourced companies to be “guinea pigs” and pilot experiments - fleshing out details until an application becomes more viable and reliable.

Trying to be a guinea pig (first mover) in a sector without also being extremely well financed and extremely AI-capable is a recipe for wasted money and lengthy, unpredictable experimentation. With very few exceptions - firms that are novice in artificial intelligence are not benefited by exploring new and open vistas of AI that don't have at least some strong precedence of traction amongst their larger competitors, or larger firms in adjacent sectors.



A filter question might be asked about any potential AI application: *Is there any precedent of a firm larger than us, in our sector or a similar sector, actually getting some value from an application of this kind?*

If the answer is new, and your firm is new to AI, then the application is likely to be a poor place to start.

1.6. IT Services and Consulting Considerations

In the near-term, there will arguably be more money to be made for services firms in helping companies prepare for AI adoption than there will be for actual AI projects. As firms move away from “toy applications” (AI projects done solely for the sake of AI, with no vision of longer-term value), leaders will more and more understand the need to properly set their foundation before embarking on AI.

Consultants and IT services firms involved in digital transformation projects may be closer to helping their clients with AI than they realize. Digital transformation is often an on-ramp towards AI transformation, particularly when those projects involve building Critical Capabilities.

As of today, the services firms who are most likely to “help their clients take advantage of AI” are not the firms who are helping to set up and train novel AI proof of concept projects (most of which will never make a difference in the core business itself), but will be the firms who help establish a strong AI foundation (AI-educated executives, data infrastructure, etc).

Because clients often say that they want to “do AI”, consultants are eager to follow suit - and indeed it’s hard to argue with a client who wants to pay for something. Smart consultants, however, should be able to use these often naive firms project to open up a broader conversation about enabling I in the long-term, which should lead to much better results for the client, and more work for the services firm.



2. AI Deployment as a Living System

More than the vast majority of traditional IT projects, AI projects are “living, breathing” solutions. They are not static or predictable at almost any phase of their deployment - particularly in novel use-cases.

The “3 Phases of Adoption” involves bringing a data science application from idea to production in a business environment, and the “7 Steps of the Data Science Lifecycle” involves the iterative process of finding the right data and model to deliver a specific result.

The 3 Phases of AI Deployment is a linear progression resulting in deployment where the 7 Steps of the Data Science Lifecycle continuously occur during each of the three phases.

The 3 Phases of AI Deployment

- **Description** - From proof of concept, to incubation, to deployment, an AI application “matures” by finding a fit between the data, the context of deployment, and the business requirements. Eventually - after iteration, experimentation and questioning - some (but not all) AI applications will reach maturity (deployment) - allowing it to stand on its own merit and deliver business value.
- **Biological Analogy: Phases of Development** - Not all newborn animals reach adulthood. Competing animals, food scarcity, and changes in the environment pose challenges to survival. Some percentage of animals reach adulthood - still requiring food and shelter but no longer demanding the constant attention of parents.

The 7 Steps of the Data Science Lifecycle

- **Description** - The steps of the data science lifecycle progress and cycle through relatively quickly. There is a constant pulse of referring to business requirements, adjusting data assumptions, adjusting features, and iterating on models and workflows in order to arrive at a promising outcome. These seven steps occur within each of the 3 phases of AI deployment.
- **Biological Analogy: Breathing and Circulation** - The ongoing cycle of steps in the data science lifecycle is like blood circulation, or other homeostatic processes of a biological system. They begin at the very earliest form of the organism’s life, and continue for as long as the organism is alive. If they stop completely, the organism (AI application) is no longer alive (delivering any value or results in the business).



The complexity of occasional unpredictability of the biological world is often an apt analogy for AI - especially for business people who are unfamiliar with the unique challenges and processes involved in actually applying AI.

Building an AI application is not “plug and play”, but implies taking care of a living thing, with growing and varied needs.

Many businesses are not prepared for the responsibility, iteration, and upkeep of a AI system, and would be better served by avoiding costly AI research and development (using traditional IT solutions, or more pre-trained vendor tools which require far less iteration and monitoring). It is important for leaders to know what they are getting themselves into before embarking on an AI project.

In the outline below, we explore the components of both the Phases of AI Deployment and the Data Science Lifecycle in greater depth.

2.1. IT Services and Consulting Considerations

Upside

The complex and evolving nature of AI systems means that services firms can be involved in an almost endless number of projects.

There is no clear “end” to most AI solutions, and the overall AI transformation needed to take advantage of AI (see Critical Capabilities at the beginning of this report) means that any given client interested in AI readiness has many years of work ahead of them. These factors pose an upside to consulting and IT services firms who can present this bigger picture, and become a trusted partner in this long-term transformation.

Downside

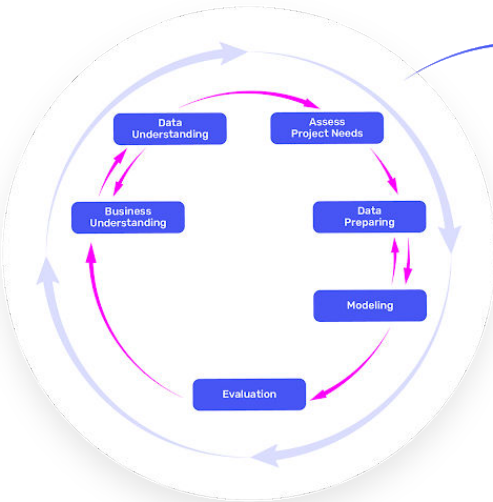
The evolving and dynamic nature of AI projects also presents difficulties for IT services and consulting firms. The serious challenges to obtaining ROI, and challenges of adoption, and the Critical Capabilities needed to take advantage of AI will all often come as a surprise to clients, who expect AI to be like any other IT solutions (“plug it in” and it’s done).

There is a delicate balance between scaring a client away from getting into AI (or scaring them away from your firm, just to do business with a firm who makes unreasonable promises), and allowing a client to go headlong into an unreasonable set of AI projects - destined to lose money and disappoint the buyer.

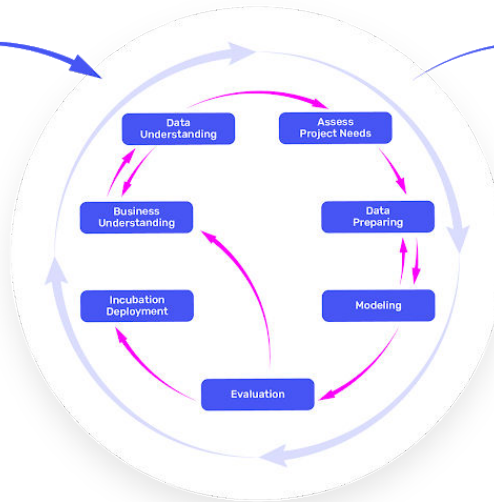


3. The 3 Phases of Deployment

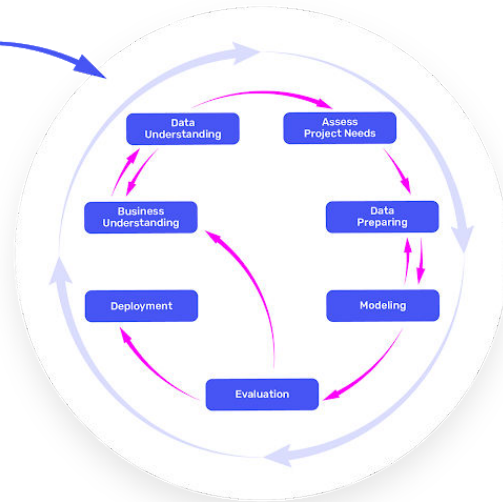
PoC



Incubation



Deployment



@Emerj Artificial Intelligence Research

emerj.com

An AI application moves through three phases. Unlike the data science lifecycle, it is unusual for an application to move backwards in phases of deployment, and - ideally - these phases are not cycled through, but progressed through linearly.

We'll illustrate the phases below with the use of two example companies:

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. The eCommerce firm sees promise in improving its cart value and improving on-site user experience, particularly for existing customers with a history of purchases and activity.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. The manufacturing firm has a strong digital infrastructure and aims to leverage its existing data streams to detect breakdowns and errors in the manufacturing process before they happen.

3.1. Proof of Concept (PoC)

- **Goal** - Determine whether or not AI can deliver a specific benefit within one business function (often in a "sandbox" environment with historical, test data).



- **Challenges** - Determining the right data and features to train on. Selecting problems with a high enough potential business value.
- **Criteria for Moving to Next Phase** - The application proves that it is capable of showing promising results - reaching some predefined criterion of success) in an isolated “sandbox” environment.

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. Use historical purchase, demographic, and behavior data to train a recommendation engine. Have a small subset of users, or in-house staff, assess these recommendations compared to the existing recommendation methods.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. Use historical data and some real time data to try and predict breakdowns of one specific kind of machine in the plant (probably, one that we have collected the most data from). See if this prediction system can be used on a number of live machines to detect issues preemptively.

3.2. Incubation

- **Goal** - Determine whether or not AI can deliver a specific benefit in a limited, live environment (i.e. real users, real-time data, etc).
- **Challenges** - Adjusting “sandbox” assumption to the real world, finding a fit within real business workflows (much different than a testing environment), and with real-time data (much different than using canned historical data, or static sample data).
- **Criteria for Moving to Next Phase** - Predetermined success criterion are reached in incubation environment. A playbook of how to manage and upkeep the AI application is determined.

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. Assuming the initial historical data test worked well - open up the recommendation engine experience to a certain subset of users. This might be done by exposing 15% of logged in users with the new AI-based product recommendation model, while other users see the default recommendations. Take measurements, adjust the algorithm to drive a higher cart value, and adjust and control for issues and errors with this incubation user group.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. Assuming the initial predictive tests showed promise - instrument a portion of all machines of a certain type with sensors and equipment to help with prediction, and



allocate both data scientists and subject-matter experts to determine whether the new predictive methods do a better job of predicting breakdowns than traditional methods.

3.3. Deployment

- **Goal** - Achieve a positive business impact in a real business environment - integrating data flow into the AI application, and integration the AI application into actual company workflows.
- **Challenges** - Training staff and adjusting workflows to account for the new AI application. Maintaining the AI system and vigilantly adjusting the system to account for changing data, a changing environment, and potentially changing or evolving business requirements.

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine.

Assuming the incubation period had fruitful results, and the teams are trained on the new AI-related processes - roll out the recommendation engine as the default user experience, with substantial team dedicated to testing and adjusting the recommendation algorithm, and collecting feedback on its results.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application.

Assuming the incubation period had fruitful results, and the teams are trained on the new AI-related processes - instrument all machines of a certain kind with sensors and connectors, creating a central set of dashboards to monitor the machines, and a dedicated full-time staff to maintaining and improving them.

Note that the PoC version of the data science lifecycle does not include the “Deployment” step, and that the Incubation phase includes an “Incubation Deployment” step. Here is the incubation lifecycle, to see the parts individually:

3.4. IT Services and Consulting Considerations

Clients should absolutely be aware of the phases of deployment in advance of starting an AI proof of concept.

Having an expectation of a separate “incubation” phase will allow the client to understand this phase as “progress”, instead of framing it a kind of hang-up or friction to smooth deployment. Being able to identify the phase a project is in, and to leverage that knowledge to determine clear next steps is something that most clients will be unable to do on their own.



Ideally, a services firm should be able to pull in the expertise of the client's executive leadership, data scientists, and subject matter experts to inform and move through each phase.

Inevitable setbacks (some data sets prove useless, certain kinds of results seem unattainable, the desired business outcomes are reconsidered entirely) should be framed from the outset as part of the learning process - and part of the ROI of early AI projects altogether.



4. The 7 Steps of the Data Science Lifecycle



Unlike the more linear phases of the three phases of deployment, the data science lifecycle steps circulate rather quickly, and there is often jumping from one step to the next in order to push forward toward an initial deployment, or an improved deployment. Steps 1 and 2 (Business Understanding and Data Understanding) and steps 4 and 5 (Data Preparation and Modelling) often happen concurrently, and so have not even been listed linearly.

The data science lifecycle has steps that can be considered in order - but that rough order is not always followed precisely in the heat of a real deployment.



For example, in the midst of data preparation, a team may decide to go “backwards” to business understanding in order to address additional budget needs (ie. data requires intensive and timely cleaning and more staff is needed), or in order to clarify a business outcome. Similarly, a team in the evaluation step might return to data understanding, or to assess project planning, before being able to actually deploy a solution.

As with the 3 phases of deployment, we’ll illustrate the phases below with the use of two example companies:

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. The eCommerce firm sees promise in improving its cart value and improving on-site user experience, particularly for existing customers with a history of purchases and activity.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. The manufacturing firm has a strong digital infrastructure and aims to leverage its existing data streams to detect breakdowns and errors in the manufacturing process before they happen.

4.1. Business Understanding

- **Goal** - Determine the business aim for the project, along with the resources allocated to its achievement. Ask: “What is the result that we are after?” Ask: “Is AI really the right tool for the job?” Ask: “What is the measurable and strategic value of this potential AI initiative?”
- **Challenges** - Finding opportunities that are reasonable and accessible for the company to reach. Do not overreach with assumptions about what AI can do. Accepting the long iteration times and critical skills and competencies that a company must develop in order to bring AI to life in an enterprise.
- **Likely Persons Involved** -
 - Senior leadership
 - Lead data scientist
 - Project manager
 - Functional subject-matter experts

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine.

Discuss the various options that the company has for growth and profitability - is a recommendation engine a priority when compared to the other options? What is understood about our customers and their purchase behavior that should be taken into account with this kind of marketing project?



Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application.

Determine how a predictive model would be measured. Think through which machines require this kind of predictive maintenance - which risks and breakdowns are most costly for the company to endure, and can we focus on those first?

4.2. Data Understanding

- **Goal** - Determine the accessibility and potential value of your data. Ask: "Can we achieve our business aims with our present data assets?" Ask: "Are there challenges with this data, or opportunities to use this data in new ways to achieve our desired business outcomes?"
- **Challenges** - Accessing the value of data, getting subject-matter experts and data scientists to look at data together to determine how it should be accessed, how it should be improved, and which features are likely of the highest value for the business outcomes.
- **Likely Persons Involved** -
 - Lead data scientist
 - Project manager
 - Functional subject-matter experts

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. Assess the quality of customer purchase behavior. Does this data tell a coherent story? Do we feel confident that one customer account is one person, or do multiple family members (different ages, priorities, genders, preferences) shop on one account, making things more complicated?

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. Look at existing data sources from manufacturing equipment. Is this time series and telemetry data from similar machines stored in similar ways, and stored in the same way? Can we ensure that the data is reliable? Where has it been least reliable, and can we reduce the factors that influenced the data this way?

4.3. Assess Project Needs

- **Goal** - Determine the requirements and resources to continue forward with the project. This might include additional budget, additional training for staff, additional subject-matter experts to join the cross-functional project team, or access to new data systems.



- **Challenges** - Getting senior leadership to endure the inevitably complex and changing needs of real AI projects (especially for firms who lack previous practical data science experience).
- **Likely Persons Involved** -
 - Senior leadership
 - Lead data scientist
 - Project manager
 - Functional subject-matter experts

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. The cross-functional team assigned to the project may decide that they need access to more historical data, and the resources to clean and organize it. They may also determine that - given the ROI opportunities in different parts of the business - they will want to apply the recommendation engine to two very specific product categories (as opposed to all products on file), and the team might request access to a dedicated subject-matter expert from that part of the business.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. The team determines the number and type of sensors that they plan to put on their various devices - and the specific subject-matter experts they would need in order to properly set up, interpret, and understand these new data streams in order to run a successful PoC.

4.4. Data Preparation

- **Goal** - Accessing, cleaning, and harmonizing data. Feature engineering to determine and distill meaningful aspects of the data corpus. Determining the feasibility of the project given the data available.
- **Challenges** - Data scientists speaking frankly with business leadership about the challenges and costs of organizing data, which are often substantial (particularly in older firms, or firms with little or no practical data science experience). Admitting that a project is not viable or feasible if the amount or quality of the data is not viable for use.
- **Likely Persons Involved** -
 - Senior leadership
 - Lead data scientist
 - Data science team
 - Functional subject-matter experts



Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. The team cleans and harmonized historical data, and determines the specific format that new data will need to take in order to help feed the recommendation engine. The data scientists and subject matter experts work together to determine the features within the purchase and user behavior data that they believe to be most important for initially training their models.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. The data science team works closely with engineers and machinists to determine the most important telemetry signals (heat, vibration) of the equipment that they are aiming to place sensors on. Then, initial sets of data is collected and analyzed, and combined in a time series with existing data streams coming from central manufacturing software. Sensor and core system data is reformatted or reorganized in a way that will allow it to be used to train models.

4.5. Modeling

- **Goal** - Establish a relationship between inputs and outputs, iterating on the data and algorithm to reach business value.
- **Challenges** - Cycling back on data processing, data understanding, and business understanding in the iteration process. Pulling in subject-matter experts to contribute to the hypotheses and practical training of the models.
- **Likely Persons Involved** -
 - Lead data scientist
 - Data science team
 - Functional subject-matter experts
 - Project manager

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. Bearing in mind the success metrics that the team has decided upon - the data science team tests new product recommendations within the specific product categories of focus. Feedback is used from team members, and (potentially) from a small cohort of users in order to calibrate towards improved cart values and conversion rates. New features in the data are used, or weighted at different levels in order to dial in towards the desired outcomes.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. The team would work together using past repair and breakdown data, and new telemetry data, to predict machines more likely to break down. This may require a relatively long time horizon, or a relatively large number of machines to initially test with, in order to



find more instances of machines in need of repair, as only these events would help to inform the model's predictive ability.

4.6. Evaluation

- **Goal** - Determining whether or not our data assets and models are capable of delivering the desired business result. This often requires many cycles back to steps 1, 2, 3, 4 or 5 - as hypotheses are refuted, and new ideas surface.
- **Challenges** - Handling challenges in evaluation, determining strong, quantifiable criteria for measuring success (where benchmarks are hard to determine). Involving senior leadership and subject-matter experts to contribute to a robust evaluation in order to allow for a confident deployment.
- **Likely Persons Involved** -
 - Senior leadership
 - Lead data scientist
 - Project manager
 - Functional subject-matter experts

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine. Over time, the team would measure their new product recommendations to previous product listings or recommendation methods. In this evaluation phase, data scientists and subject matter experts come together to determine what seems to be working, what is not working, and how to adjust the models, the data, or the user experience of the recommendation model to better drive towards desired outcomes (higher cart value, higher conversion rate of users to customers).

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application. The cross-functional team assesses the predictive models suggestions, determining if they are tangibly better or worse than previous methods. In an early Proof of Concept or Incubation phase, this may be more qualitative (ie. Do we believe that our previous methods would have detected this equipment malfunction?), while in a real deployment, this measurement would be quantitative (ie. How many breakdowns occur per month? How much lost uptime occurs across X category of machines per month? What are the rates of false positives for the predictive maintenance system?).

4.7. Deployment

- **Goal** - To successfully integrate the AI model or application into existing business processes. Ultimately, to deliver a business outcome.



- **Challenges** - Training staff to leverage the new AI application. Ongoing upkeep required to keep the model working, and adapting to change.
- **Likely Persons Involved** -
 - Lead data scientist
 - Data science team
 - Project manager

Example 1 - An eCommerce Firm Adopting a Product Recommendation Engine.

- **Phase 2: Incubation Deployment:** The recommendation engine, which has been tested sufficiently in a sandbox environment, with feedback from internal team members - is integrated into a part of the eCommerce website, and 10% of users are exposed to AI-generated recommendations, rather than previous recommendations.
- **Phase 3: Full Deployment:** The recommendation system is integrated into the website fully - becoming the default experience on all web interfaces where the team believes it will deliver value. A monitoring system is established to calibrate the results and findings from the new system, with a regular pace of meetings and diagnostics to ensure that the system is performing and improving.

Example 2 - A Manufacturing Firm Adopting a Predictive Analytics Application.

- **Phase 2: Incubation Deployment:** The predictive maintenance system is integrated into a part of the workflow on the manufacturing floor. Now, a small cohort of machinists and engineers - some of whom were likely not part of the cross-functional AI team - are able to use and respond to this new system, under the guidance of the AI team.
- **Phase 3: Full Deployment:** The predictive maintenance is integrated into the manufacturing workflow fully - becoming the default process in all the machining functions that the AI team believes it can deliver value (areas that have already been tested in the PoC and incubation phases). A monitoring system is established to calibrate the results and findings from the new system, with a regular pace of meetings and diagnostics to ensure that the system is performing and improving.

4.8. IT Services and Consulting Considerations

The most challenging aspect of the data science lifecycle in a business environment is getting buy-in and participation from the relevant parties involved - from data understanding to evaluation and beyond.

There is no single playbook for how to do this - and initial attempts are bound to be clumsy as the process involves pulling data scientists, subject matter experts, and leadership into conversations they are not used to, during the hours in which they would



normally be doing their usual job (the work they were hired to do, which was not to be on a cross-functional AI team).

As a services firm who established trust with a client, it is important to emphasize the need for cross-functional participation from the start, and to identify where that participation may be lacking (so that it might be resolved) as a project is ongoing. Again, this is bound to be clumsy in the early days of AI adoption - especially for firms with little experience with enabling AI in their business in the past.

From the standpoint of a services firm, the most important stakeholder to involve is the economic buyer (usually some kind of senior leader), who - more than any other stakeholder group - should have an understanding for the cross-functional needs of a project. It is this buyer who will likely have the strongest “pull” to get such team members involved in a project, and it is this buyer who is most likely to be able to veto a project if they don’t understand how it is progressing, or feel that it is out of control.



5. Putting the Steps and Phases into Action

5.1 Preparation

Understanding the Phases of Deployment and Steps of the Data Science Lifecycle can be valuable for companies in any phase of AI maturity. Whether a company is considering AI adoption, beginning a decided-on AI project, or in the middle of AI implementation and testing - understanding the steps and phases can help make AI deployment easier and less risky.

We recommend that leaders ask themselves a set of questions to determine their alignment in the use of the steps and phases - and these questions vary depending on what situation their company is in:

Assessing AI Opportunity

For leaders or consultants who haven't yet decided on an AI initiative, and are assessing their options:

- Do we have the requisite team components to build an AI team (leadership who understand AI conceptually, in-house data science talent with an understanding of our business processes, subject-matter experts with willingness to be part of a cross-functional AI team), or do we need time to establish these core team components?
- Do we have the necessary budgets - and high-value potential uses-cases - to justify the lengthy, hands-on process that AI deployment requires?
- Where do we stand with regards to the Prerequisites to AI Deployment (section 1 above)?
- When examining the top use-cases for AI that we are considering for our company - what might a PoC or Incubation phase of those use-cases look like? What kind of resources and timelines might each of them require?

Beginning an AI Project



For leaders or consultants who have already decided on an AI initiative, and want to begin a proof of concept:

- What will each of the Phases of Deployment look like for this project?
 - Proof of Concept:
 - What kind of data might we need in order to run a controlled, “sandbox” PoC project?
 - What permissions do we need to access that data - and what expertise do we need to clean the data and determine the value features within it?
 - What are the criteria of success from which we can move beyond the PoC to the incubation phase?
 - Incubation:
 - Without risking our user experience or a total change in our workflows and data infrastructure, how can we take our ML model “live” in a limited incubation phase?
 - What permissions would we need in order to move to this phase?
 - Which customer groups or internal team members would be impacted by this incubation phase, and how could we recruit their help?
 - Deployment:
 - What would be the desired scope of our deployment? (i.e. will we be using this chatbot for all customer service inquiries, or just for a small portion of them? Will we be leveraging predictive analytics on all our manufacturing equipment, or only certain types of machines in a single plant?)
- Are there any steps in the Data Science Lifecycle we have left out?
- Do we have the right mix of team members actively involved throughout the steps of the Data Science Lifecycle?
- Who will form the core of our cross-functional AI team?
- Do we have someone in leadership who can champion this project and help us obtain resources and team participation? Does this person have enough of a conceptual grasp of AI and AI opportunity to actively support this project?

In the Midst of an AI Initiative



For leaders of consultants in the middle of a proof-of-concept or incubation-phase AI deployment:

- Have we determined the measures of success for this AI project?
- Have we secured someone in leadership who can champion this project and help us obtain resources and team participation? Does this person have enough of a conceptual grasp of AI and AI opportunity to actively support this project?
- Where are we now across the 3 Phases of Deployment?
 - Have we determined our threshold goals for when we are able to move into the next Phase of Deployment?
- Where are we now in the 7 Steps of the Data Science Lifecycle?
 - (See the “Goals” listed for each step) Are there any Data Science Lifecycle steps that we have neglected or overlooked?
 - (See the “Likely Persons Involved” listed for each step) Are there any kinds of team members whose participation we need more of during these steps and phases?

These questions may lead to the decision to hold off on adopting artificial intelligence, or to working on developing Critical Capabilities in advance of moving forward to a PoC project. This outcome is far better than leaping into an AI project without the requisite skills, resources, or expectations.

The role of AI strategist (whether an employee or outside expert) is not merely in convincing leadership of the value of AI, but in preparing a business to understand the challenges of AI, and to encourage them to see progress not merely in financial ROI, but in building the Critical Capabilities that are so important to actually deploying AI in the first place.

5.2 Ongoing Consideration

While this guide should be able to help companies avoid some of the more blatant pitfalls of applying AI in an existing business - there will never be a perfect project.

A company's first dozen AI initiatives will be tremendous learning experiences - where different Phases and Steps will be refined, skills will be built, and team dynamics and talent needs more carefully calibrated over time.



Bear the following in mind throughout your initial set of AI initiatives:

Expect to Get it Wrong

Getting it wrong is okay. AI is hard. That's why selecting the best projects and having the right upfront expectations is critical. Getting it wrong is part of the process.

Companies that stumble with AI adoption and avoid AI altogether and setting themselves up for failure. Companies that embrace inevitable stumbles and encourage their teams (data scientists, leadership, subject-matter experts) to learn from them - will be to better prepared to take advantage of AI moving forward.

One of the Critical Capabilities is "Culture of innovation" and "Value of data", both of which can be improved by taking a tolerance and patient view of the challenges of AI adoption.

Learning is Progress

Companies with clear goals for their AI initiative are likely to measure the success of their AI efforts by asking the following questions:

- Has our ML model outperformed our previous approach in a PoC or incubation phase?
- Based on current ML model performance, what kind of cost savings or revenue improvement might we expect in the coming year?
- What have we spent thus far on this AI initiative, and does it seem promising to continue investing?

These questions are all appropriate to ask, and no AI project should be established without some quantitative standards or timelines established upfront - even if these evolve over time.

Companies with realistic expectations about AI deployment, however, will also periodically ask the following questions throughout the course of an AI initiative:

- What Critical Capabilities have our team developed through the course of this project thus far?



- Have our data assets or data accessibility improved over the course of working on this project?
- Have we discovered any new challenges to AI deployment, or any new AI deployment areas over the course of our learning thus far?

Initial AI projects are rarely a massive, measurable ROI win. Rather, they are opportunities to learn hard lessons, build Critical Capabilities, and become more and more prepared to take advantage of future AI applications and gains in the future.

Involve and Inform Leadership

It is not uncommon for business leadership to be the largest barrier to AI deployment within a company.

Projects are often approved or denied based on the judgement of leaders with very little conceptual understanding of AI, or knowledge of AI's strategic value. Consequently, resources are allocated to the wrong projects, with the wrong expectations.

AI strategists and data science leaders should focus not just on getting approval from leadership, but on educating and communicating with leadership - to not just get approval - but to get buy-in and connection to common vision.

Only with some degree of this shared vision can data science teams and AI project managers expect to have the freedom and resources they need to finish an AI project, or to persist through the inevitable setbacks in a project.

If your leadership have little context on AI's applications and strategic value, use initial AI projects as a way to continue ongoing dialogue with leadership, educating them, and helping to develop their own sense of AI's strategic value, and their own sense of realistic expectations around AI.

Business people often assume that the AI advantage of Google, Facebook, Amazon, and other tech giants is their profitability, ability to afford hordes of data scientists, and their gigantic data warehouses.

They often overlook the tremendous benefit of having leadership that - by and large - understands AI conceptually and can align company goals and resources to the most



promising AI ideas. The sooner that other businesses can obtain this same level of resonance between leadership and cross-functional AI teams, the better. Each initial project should be seen as an opportunity to get closer to that ideal.

5.3 Retaining Enterprise AI Competence

One could argue that deploying initial AI projects is more about the ROI of learning than it is about the ROI of any first AI application itself. In truth, that learning won't happen unless a cross-functional AI team strives for a real ROI from the application itself - but even if that application produces no financial return, a team can be better off than it was before the project started - and more able to make the most of future projects.

This "learning ROI" results from improving the Critical Capabilities mentioned in section 1.1 of this report.

The Critical Role of Retained Learning

From the outset of any project, a cross-functional AI team should have:

- a) A strong grasp of the problem to be solved or opportunity to be exploited
- b) A reasonable understanding of how success will be measured, and what kind of tangible return would make the project worthwhile (see our full report on [Generating AI ROI](#)).
- c) An outline and initial plan for how to approach the project, phase-by-phase (the guide you'd reading now is intended for that exact purpose)
- d) A place to record lessons learned from this particular AI project.

Each company manages knowledge differently, and the particular system or repository isn't as important as the habit of tracking and maintaining a categorized list of insights.

The person in charge of retaining these lessons learned will sometimes be the project manager themselves, or a subject-matter expert (working closely with the project or product manager) assigned specifically to the task of tracking lessons learned.

Categories of Retained Learning



Retained insights will generally be categorized according to the data science life cycle step or the phase of AI deployment that they pertain to. Alternatively (or additionally), insights can be sorted or categorized along the lines of the Critical Capabilities.

We'll explore a number of examples of the kinds of insights that a cross-functional AI team will want to retain.

Phases of AI Deployment

- **PoC** - An AI team may discover that in order to run a PoC successfully, they need to handle data access, legal, and compliance issues upfront - rather than waiting to ask those questions when heading into the incubation phase, when insurmountable issues may come up after months of work on a seemingly viable PoC. The team may devise a specific legal and compliance checklist for all PoC projects which can be used to screen future vendors and "sandbox" projects of any kind, preventing future waste.
- **Incubation** - An AI team in the Incubation phase may discover that an Incubation project cannot be fully deployed unless a maintenance team can be established ahead of time. The team may learn that team members are reluctant to become part of the maintenance of an AI application (say, a fraud detection system within an eCommerce company) if that isn't the job they signed up for with the company. Hence, future Incubation periods may involve determining hiring and recruiting needs for the maintenance team - and a robust process of recruiting subject-matter experts from within the company to flag themselves of interest in deploying the project, well before it can actually be deployed. A process or checklist might be developed for this - to be tweaked and used across future AI projects.
- **Deployment** - An AI team may learn that the deployment of some AI systems requires new types of interoperability with existing technology systems. A machine learning integration map might be created to determine the required integrations and potential risks and challenges of different kinds of deployments, allowing teams to better establish IT and data infrastructures that could sustain a live AI deployment.

Steps of the Data Science Lifecycle

- **Step 2 - Data Understanding** - An AI team that works continuously with its most important data may discover important patterns or errors within that data. For



example, a bank may discover key patterns in its bank transfer data, or common errors that cause some data to be stored differently or incorrectly. Recording these common issues (or repairing them at their source) would help future AI teams who will work with the same data later on.

- **Step 4 - Data Preparation** - An eCommerce company may discover critical insights about how to prepare its transaction and behavior data for recommendation-related AI applications. This data preparation routine could be recorded and documented for future teams - and could be used as a template for helping to improve analytics systems in the future.
- **Step 6 - Evaluation** - An AI team may determine that the Evaluation phase must require agreement on the next steps on the project from data scientists, subject-matter experts, and executive leadership. Developing a set of Evaluation meeting templates could help the company to have the right Evaluation “checkpoints” for its data science lifecycle interactions, ensuring that success or failure are agreed upon, and the project can continue on the best possible path forward.

Critical Capabilities

- **Contextual AI Understanding for Business Leaders** - A company may learn over time that without a specific amount of basic AI understanding, and use-case understanding - functional business leaders are unable to properly handle conversations with AI vendors, or properly allocate resources (with realistic expectations) for AI projects. For this reason, the firm may develop an internal AI curriculum for any functional business leader (a Head of Compliance, a VP of Marketing, etc) who will be involved with AI initiatives, or with vetting AI vendors.
- **A Culture That Values Data** - A company may discover, after a number of AI projects, that the data assets in some departments or silos are more uniform, more reliable, and better kept than the data in other departments or silos. Diagnosing those differences and building a set of protocols for handling data and ensuring data quality could help make future AI projects easier, and could help shift the corporate culture for the better.
- **Data science Skills and Talent** - After even just a few AI initiatives, a company will get a sense of the data science talent mix that is required for different kinds of projects (data science lead, ML engineers, data engineers, etc). By developing a set of norms and standards for AI teams, a company can better estimate its hiring and recruiting needs, and can onboard the right kind of staff to contribute meaningfully to existing projects.



These are a number of hypothetical retained lessons from AI deployments, but they should serve to showcase the kinds of specific frameworks and best practices that can be developed by a team that keeps their eyes open to areas of improvement, and specific insights about AI adoption within their business.

For any of the lessons above, a company may have to make a painful mistake. Many lessons will be learned “the hard way”, and in the nascent world of enterprise AI, it will be hard to avoid all such mistakes.

The ROI of initial AI projects is, in large part, a company and team’s ability to better leverage AI into the future. The company with the best Critical Capabilities (rich soil for new AI opportunities to grow) is more likely to win in the marketplace than a company with one high-ROI AI application, but no core foundational skills and resources to support it.

Deliberately turning AI deployments into learning experiences is an advantage that most companies will neglect. Those who bolster their core skills and retain the lessons of applying AI to their specific data assets and their specific business processes will be setting themselves up to win in the years ahead.

5.4. IT Services and Consulting Considerations

Telling clients that “learning” is among the most important kinds of ROI for early AI projects is not only a way to buffer against disappointment with early projects (which is common, and hard to avoid, for all the reasons we’ve laid out in this report) - but it’s absolutely true.

Clients who see a consulting or services firm as a conduit to enterprise learning are able to find an ROI in almost any well-intended AI project - potentially giving that services firm more opportunities to work with that client in the future.

It is somewhat unlikely that a strategy consultant or IT services firm will be the one actually recording the learns gained as the AI project progresses - but at the very least - services firms should emphasize the opportunity to learn (and build Critical Capabilities) as an important part of the engagement, and should encourage someone on the client’s team to be a point of contact to share and build on these lessons through the project.

While it’s important to always aim for a measurable ROI, and to hold any AI initiative accountable to measurable benchmarks - it is likely that initial proof of concept projects



will not be able to be integrated into the business successfully, and that progress on Critical Capabilities will be the most tangible ROI from a given project. If this progress can be seen as a “win” for the client (by helping to set their expectation ahead of time), it helps a client have more realistic expectations about AI adoption.



Emerj Artificial Intelligence Research

Emerj Artificial Intelligence Research is where executive leaders turn to understand how AI is impacting their organization or industry – and what to do about it. We're the industry source for authoritative market research and competitive intelligence for the business applications of artificial intelligence.

Our objective, jargon-free research and industry overviews are designed to give executives and decision-makers exactly what they need for competitive insight, informed AI technology procurement and strategic planning around AI.

With a finger on the pulse of academia, Fortune 500s, and the global artificial intelligence startup ecosystem, organizations call upon us for insight and research for their most important AI-related strategic decisions.



Through our [Research Services](#), [AI Capability Maps](#) and [AI Business Strategy Process](#), we help clients win market share and make more profitable decisions – with a firm grounding in the current realities of the AI landscape.

Contact Emerj

services@emerj.com

1-617-945-8567

