

Defining A Successful AI Project

A Framework For Choosing The Right Use Cases



A WHITE PAPER BY DATAIKU


www.dataiku.com

INTRODUCTION

As hype around data science, machine learning, and AI continues to grow, more and more organizations are feeling the pressure to modernize their business by implementing advanced data solutions or else risk falling behind the technology curve.

But how to choose the right project? Should it be a trendy, consumer-facing solution such as a chatbot, virtual assistant, or computer vision application? Or is it better to focus on embedded, back-of-the-house projects that quietly help to improve internal operations, employee efficiency, or decision-making?

Generally, AI projects fit into one of the following buckets:



With dozens of potential use cases but limited resources, it is important to prioritize projects that have both high business value and a high likelihood of success. That's where a time-honored framework - the five Ws and an H - comes into play for organizations starting their path to AI.

An ideal AI project will have clear and compelling answers to each of these questions:

- **WHO** will this project benefit?
- **HOW** will it specifically improve experience or outcomes, and **HOW** can this be measured?
- **WHY** is using AI for this purpose better than existing processes?
- **WHAT** is the upside if it succeeds, and **WHAT** are the consequences if it fails?
- **WHERE** will the data come from, and does it already exist?
- **WHEN** should an initial working prototype and, subsequently, a final solution in production be delivered?

This white paper will go through each of these considerations in detail and provide examples, strategies, and follow-up questions to help guide the selection of AI use cases that will ultimately bring business value.

WHO WILL THIS PROJECT BENEFIT?

Identifying the intended target audience is a critical step in defining an AI project. One important consideration is recognizing that “audience” doesn’t necessarily have to be external - the people benefiting from an AI project can certainly be employees inside the organization, especially if the use case is one aimed at bringing internal speed-to-value or team efficiency.

For example:

A financial services organization is considering whether to build an external-facing application -- specifically, a virtual assistant that integrates both with a mobile banking app and Amazon Echo. Before diving in to the project, it's critical to determine whether there are specific segments of customers for whom the product would be most relevant. One important question might be: How important to the overall business strategy are these particular segments?

The WHO question is therefore directly related to an AI project’s ultimate success or failure. Data science teams can certainly develop innovative models and systems, but if they don’t design the project with the target audience’s application of the solution in mind, they risk their work product being neither useful nor adopted.

Takeaways

- Define a specific audience for an AI project before building the solution.
- Collaboration is key - those developing AI projects need to work closely with business stakeholders or other subject matter experts to not just choose, but truly understand the audience. Who are they? What are their needs? What are their current processes and habits?



HOW WILL IT SPECIFICALLY IMPROVE THE INTENDED AUDIENCE'S EXPERIENCE OR OUTCOMES? CAN WE MEASURE THIS?

With a target audience identified, it is then important to consider exactly how the AI application or system will improve that audience's experience. This HOW should be as quantitative as possible in order to show the return on investment (ROI) of the AI system once implemented.

Consider the use case of building a back-office AI application at an insurance company that takes freeform text comments (such as those found in insurance and warranty claims, call center notes, or medical health records) and automatically tags text with all relevant categories, extracting key entities such as dates, locations, and names.

- The audience (the WHO) here is internal employees at the company, specifically those dealing with claims or managers analyzing customer service interactions.
- One HOW question might be: How much time would this system save a subject matter expert who previously had to do these tasks manually?

Not properly quantifying potential AI use cases with the right HOW questions up front puts executive sponsorship for future endeavors at risk - it is critical to be able to measure expected lift. In the example above, measuring man-hours saved is one way to quantify ROI, but another could be estimating the incremental revenue earned by automatically detecting insights (or the money saved by predicting fraudulent claims or future risk events) which previously would have gone unnoticed, whether due to human error, lack of ability to scale manual review to all data, or data complexity.

Real-Life ROI: By using the approach outlined above to automatically scan patient records at scale, a major hospital in Europe was able to both improve patient outcomes and gain an ROI in the millions of euros. They did so by using text analytics to detect secondary conditions and treatments mentioned in medical notes that had previously gone undocumented, and therefore also were unbilled to insurance; furthermore, automating this process saved hundreds of costly doctor-hours in manual review.

Go Further:



Takeaways

- Always ask: HOW will this AI use case help me make money, save money, or do something I can't do today?
- If the answer is not clear, then scratch it off the list. AI use cases (especially early ones when garnering organizational-wide support for AI efforts) should focus on opportunities with real and measurable business results.

WHY IS USING AI FOR THIS PURPOSE BETTER THAN EXISTING PROCESS?

It is easy to be seduced by the idea of AI becoming a new “easy” button. But new doesn't always mean more valuable - when weighing the pros and cons of various AI projects, consider whether the anticipated benefits justify the time and investment required to replace an existing process.

For example, Let's say:

- A marketing department at a large retailer manually tags newly licensed images every week so they may be properly indexed in the digital asset management system.
- This way, employees can find corporate-compliant images for their presentations that contain the specific elements they want (e.g., a business person, a tablet, an airplane setting, etc.).
- **Would it be useful to build an object recognition model to automatically detect and tag the items in each photo?**

In this case:

- The low volume of images and the business value of automating this process probably are not sufficient to overcome the significant effort it would take to train and maintain a robust deep learning model for object detection from scratch.
- However, perhaps using a pretrained model (e.g., YOLO, COCO, Mask R-CNN, etc.) could provide the shortcut needed to make this wish a reality without expending too many resources.

Takeaways

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



WHAT IS THE UPSIDE IF THE PROJECT SUCCEEDS? WHAT ARE THE CONSEQUENCES IF THE PROJECT FAILS?

By now, the importance of quantifying the potential benefits of a successful AI project - whether they be increased revenue, operational cost-savings, etc. - should be clear. But risk is also something that should be examined up front.

No one likes to think about AI projects failing before they begin, but the reality is that VentureBeat AI reports¹ 87 percent of data science projects never make it into production. So it is important to evaluate the consequences should the worst happen.

- Will there be significant damage to brand reputation or customer trust?
- Will there be damage to the level of organizational trust in AI projects and the ability of data teams to execute?
- Are there risks regarding regulatory compliance or customer privacy?

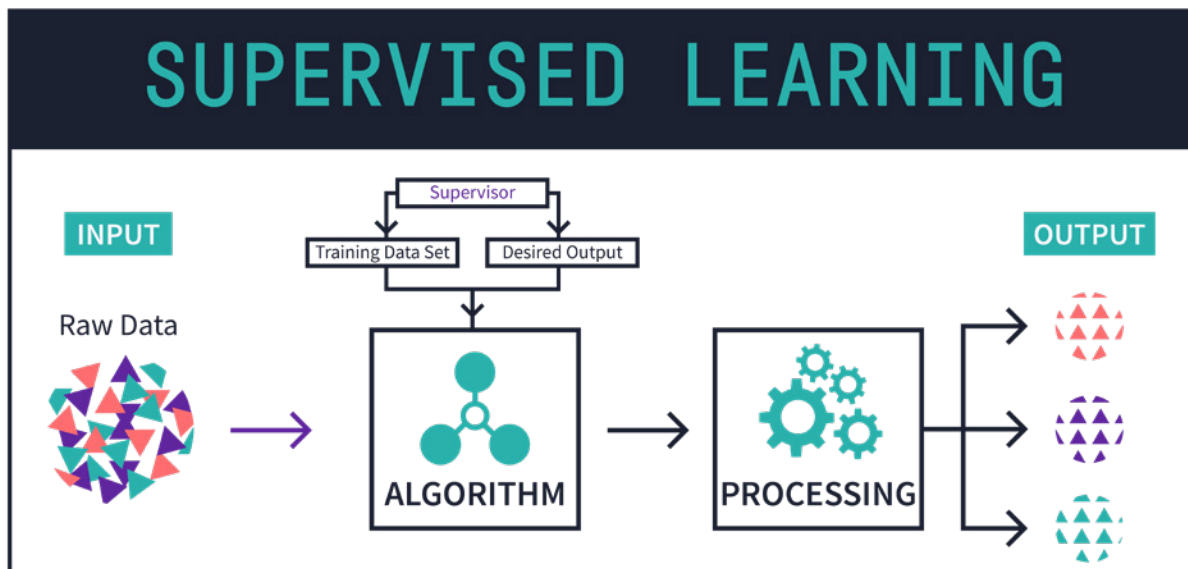
Celent analyst Dan Latimore put it best in his 2018 publication Artificial Intelligence in Banking² when he said, "Pick a well-defined and manageable project whose failure won't be catastrophic and whose success will be inspirational."

Takeaways

- *Choose an AI project that won't have devastating consequences in case of failure.*
- *Evaluating the following are critical to identifying a project whose failure would not be catastrophic to either the company or AI programs themselves.*
 - *Project scope*
 - *Access to subject matter experts*
 - *Availability of analytical resources (both human and infrastructure)*


WHERE DOES THE DATA COME FROM, AND IS IT APPROPRIATE FOR THE OBJECTIVE?

It should come as no surprise that data is the most important element of an AI system. The majority of modern AI solutions today leverage supervised machine learning algorithms, which require a labeled set of data from which data scientists train the machine. In other words, the computer learns from the examples (which include the answer key) that it is fed, and typically, these algorithms require quite a lot of labeled training data in order to produce highly accurate results.



The disconnect occurs when organizations want to implement cool AI solutions but don't have the data readily at hand for the specific task. Those businesses tackling first AI projects might consider selecting initial projects that only rely on internal data rather than ones that heavily rely on third-party sources or custom-manufacturing large quantities of training data (think Stanford's Snorkel, Amazon's Mechanical Turk, or Google's AI Platform Data Labeling Service).

Go Further:



Takeaways

- Minimize external dependencies for a first AI project. This not only reduces complexity, but also steers teams toward early projects with a high speed to value, which is important for gaining organizational traction.
- Collaboration between the business, IT, and data science groups will be critical for this consideration as well, so involve relevant stakeholders early and often in project design.



WHEN DOES AN INITIAL WORKING PROTOTYPE NEED TO BE COMPLETED? WHEN WILL A FINAL SOLUTION IN PRODUCTION BE IMPLEMENTED?

Finally: the all-important timeline. Like a do-it-yourself home construction project, it can be easy for an AI application build to stretch on and on; there always seems to be a bit more fine tuning to be done, another small feature to be added.

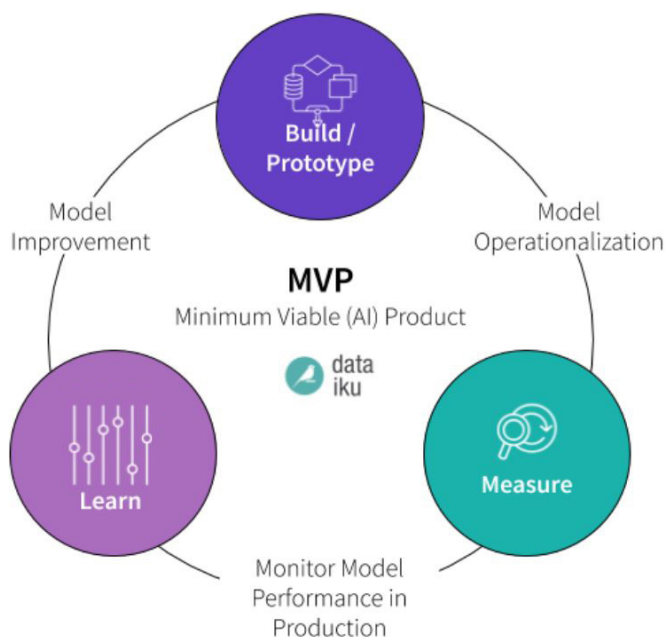
But in order to build credibility with internal stakeholders, it's best practice to have a limited slice of the solution working from end-to-end in a short period of time rather than aiming for a fully baked solution in the first pass.

Real-life example: Take the example of a customer-facing chatbot; the first prototype can either focus on breadth (i.e., it can answer a wide variety of simple inquiries across diverse topics), or depth (it can answer very detailed questions asked in many different ways, but only about a few topics). However, it's probably not realistic to strive for both of these goals in a first iteration. Once the chatbot preview earns internal user acceptance -- that is, stakeholders confirm that once built out, it will in fact deliver the desired business outcomes--only then is it appropriate to fully flesh out the application.



The WHEN question also touches on a - or maybe the - key step in an AI project: operationalization. Operationalization is the process of converting data insights into actual large-scale business and operational impact. This means bridging the huge gap between the exploratory work of designing machine learning models and the industrial effort (not to mention precision) required for deployment within actual production systems and processes.

AI project timelines should not only include time for development and a deadline for a working prototype, but also deadlines for deployment into production and for second (or third) iterations. AI projects fit well into the [MVP methodology](#) because the best way to see how a machine learning model performs is in production, so the key to successful projects is operationalizing early, adjusting, and pushing new versions often.



Takeaways

- A “progress before perfection” and MVP mindset enables stakeholders to get excited about the long term vision of an AI project and give feedback while it's still early enough for major course corrections.
- AI project timelines should include not only development to prototype, but development to operationalization.
- It's critical to set these timelines upfront, working through any potential impending obstacles to operationalize before getting there and discovering roadblocks, if possible.

Go Further:



CONCLUSION

Today's most successful organizations have embraced the idea that effectively leveraging data and technology can not only drive competitive advantage, but also improve the experiences of both their employees and customers. But it's important to be strategic and deliberate when planning an organization's AI strategy.

Following the framework outlined in this white paper will help avoid false starts on AI projects that are ill-defined and create an environment for success. This final worksheet is designed to be printed and used as a guide for AI project planning:

- ☐ **WHO** will this project benefit?
 - ☐ Who is the audience?
 - ☐ What are the needs of the audience?
 - ☐ What are the audience's current habits and processes?
 - ☐ If considering an AI project for an external audience, who within the organization knows this audience best and can serve as a subject matter expert (SME) on the project?
- ☐ **HOW** will it specifically improve experience or outcomes, and HOW can this be measured?
 - ☐ Will the AI project help:
 - ☐ Reduce risk
 - ☐ Increase revenue
 - ☐ Decrease costs
 - ☐ Give a competitive edge
 - ☐ Increase speed-to-value or internal team efficiency?
 - ☐ What are the key metrics that will be tracked to show the project's ROI?
 - ☐ Who will own the task of calculating ROI for the project?
 - ☐ How much time after the project is implemented will ROI be measured?
- ☐ **WHY** is using AI for this purpose better than existing processes?
 - ☐ What are the existing processes?
 - ☐ What are the key metrics that define these existing processes?
- ☐ **WHAT** is the upside if it succeeds, and WHAT are the consequences if it fails?
 - ☐ What is the project scope?
 - ☐ Does the team have access to the necessary subject matter experts?
 - ☐ Are all required resources (both human and infrastructure) accessible to not only start, but complete the project?
- ☐ **WHERE** will the data come from, and does it already exist?
 - ☐ What data will be used for the project?
 - ☐ Who owns the data?
 - ☐ How will the team gain access to the necessary data?
- ☐ **WHEN** should an initial working prototype and, subsequently, a final solution in production be delivered?
 - ☐ Define a project timeline, including deadlines and dates for:
 - ☐ Getting the right data
 - ☐ Cleaning and preparing data
 - ☐ Building model(s)
 - ☐ First working prototype
 - ☐ Feedback cycle
 - ☐ Deployment to production



300+
CUSTOMERS

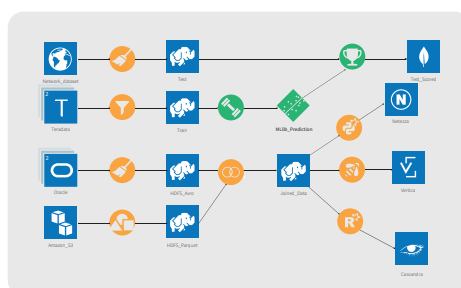
30,000+
ACTIVE USERS




KUKA




Name	Sex	Age
Michael, J	Gender	32
Braund, Mr. Owen Harris	19.00	21
Allen, Mr. James	19.00	38
Harbeck, Remove rows containing Mr		26
Patella, Keep only rows containing Mr		35
Allen, Mr.	Split column on Mr	35
McCarty, Explosion Mr Sp...		29
Harbeck, Remove rows equal to Mr. James		
	Keep only rows equal to Monan, Mr. James	
	Clear cells equal to Monan, Mr. James	
<input checked="" type="checkbox"/> Filter on Monan, Mr. James		
<input type="checkbox"/> Filter on Mr		
<input type="checkbox"/> Toggle row highlight		
<input type="checkbox"/> Show complete table		




5. Monitor & Adjust



4. Deploy to production



3. Mining & Visualization



WHITE PAPER

www.dataiku.com