Why do most Machine Learning Projects Never Make it to Production?

Cameron Vetter

Microsoft®
Most Valuable
Professional

Cameron Vetter

- ✓ Technologist
- ✓ Fractional CTO
- ✓ Enterprise Architect
- ✓ ML Engineer
- ✓ MR Enthusiast
- ✓ Microsoft MVP - AI

OCTAVIAN
TECHNOLOGY GROUP

Microsoft Partner | Silver Data Analytics

- ✓ Machine Learning
- ✓ Cloud Architecture
- ✓ Fractional CIO/CTO
- ✓ Innovation Workshops

Octavian Technology Group

slido

Join at slido.com
#107213

ⓘ Start presenting to display the joining instructions on this slide.

slido

What is your Organizational Role?

ⓘ Start presenting to display the poll results on this slide.

slido

Why did you choose to come to this talk? What do you hope to gain?

ⓘ Start presenting to display the poll results on this slide.

slido

Have you taken part in a Machine Learning Project?

ⓘ Start presenting to display the poll results on this slide.

?

What goes wrong?

We discuss common mistakes in Machine Learning Projects.

Why does it go wrong?

My opinion on why these mistakes are made.

PROBLEM

SOLUTION

How to Avoid Becoming a Statistic

What not to do in your machine learning project.

My Advice

Some tips on how I setup my machine learning projects for success.

A Bad Foundation

LEADERSHIP THROWS **MONEY** AT **MACHINE LEARNING?**

"**Money is only a tool. It will take you wherever you wish, but it will not replace you as the driver.**"

@ Ayn Rand

LACK OF **LEADERSHIP SUPPORT**

Leadership Misalignment

Is only part of your leadership team onboard?

The Experiment

Is your machine learning project an experiment not expected to succeed?

!

POORLY DEFINED **ROI**

Why?
Why is this project critical to your organization?

What?
What will it bring to the company? Increased Revenue? Lower Cost? New Line of Business?

When?
When will this investment be realized?

The Ownership Problem

Data Scientists Focused On **Experimental Model Creation** Only

Data Science teams often focus on iterative experimentation and often consider their work done once the model works.

Not Scalable

Model is not computationally or financially possible to run at scale.

Too Slow

Model takes too long to execute to provide production level performance.

Not Refactorable

Model is very difficult to refactor into something production worthy.

Input Shape

The input shape required for the model is not possible in a real time or production environment.

Models are **tossed over the wall** to an engineering team

This is **not MY model** the other team owns.

When one team creates the model and tosses it over the wall to another team to iterate the model, often neither team feels ownership of the model. Treat a model like any other software unit in your organization!

Not Starting Simple

Chasing the Unicorn

Starting with the earth shattering, company redefining, machine learning project, instead of the simple and achievable.

Not Demonstrating Progress

The team is so focused on the end goal they fail to perform simple incremental progress that shows business value.

NeverEnding Projects

Projects so large or so infeasible they drag on forever and create organizational fatigue.

No MVP Plan

Every machine learning project needs an MVP plan, just like every software project.

The Wrong Team

Are they a team?

DATA EXPERT

I know the data inside and out, and I know where it all physically exists. I understand how the different data relates to each other and I understand the business domain well enough to provide context for the rest of the team.

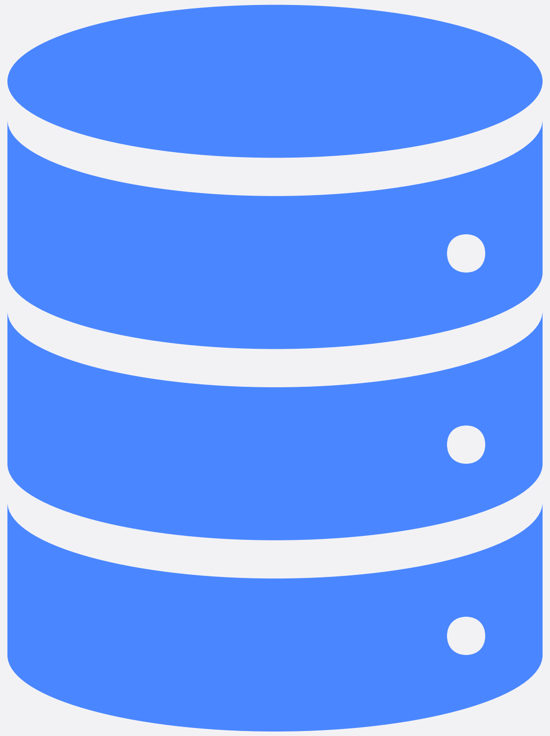

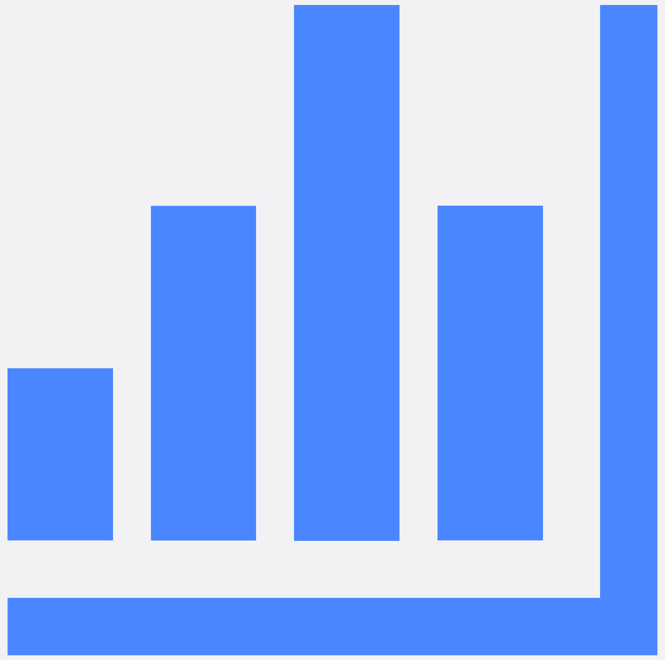

DATA SCIENTIST [optional]

I live and breath data. I know how to manipulate it and apply world class techniques. I may not be very good at coding, but I can do enough to get the result I want. I often have a PHD and the math and statistics skills to back it up.

BUSINESS SPONSOR

I’m strongly positioned within the business and well respected by the company's leadership team. I am the champion for this project and truly believe in it. I can easily speak to the business benefit and when it will be delivered.

ML ENGINEER

I’m a great engineer that understands design patterns, development process, and production quality software as well as any other engineer developing software within the company. What differentiates me is a thorough understanding of how to apply machine learning and the ability to build custom ML models.

MLOPS ENGINEER

I understand how to manage the code, data, and models associated with machine learning. I will create a place and versioning scheme for each of these. I also will create and manage a pipeline to automate the training, testing, and deployment of the machine learning models for this project. I bring the rigor and repeatability of a traditional development project to a machine learning project.

DATA LABELING TEAM

Our team is able to label each row of training data based on our business domain knowledge. We will be able to create massive data sets that are used by the rest of the team to train and validate the models. We are usually lower cost labor but are the most important part of the team, the quality of your machine learning model will never be any better than the quality of our labeling work.

Missing Process

SDLC

Although machine learning work is experimental you still need process.

Test Strategy

How will I validate this model works? How will I prove it works in the real world?

Release Strategy

Will I do Side by Side releases of different models? How will I roll back? Can I turn off the ML completely?

SDLC

Although machine learning work is experimental you still need process.

Test Strategy

How will I validate this model works? How will I prove it works in the real world?

Release Strategy

Will I do Side by Side releases of different models? How will I roll back? Can I turn off the ML completely?

SDLC

Although machine learning work is experimental you still need process.

Test Strategy

How will I validate this model works? How will I prove it works in the real world?

Release Strategy

Will I do Side by Side releases of different models? How will I roll back? Can I turn off the ML completely?

Data, what data?

ACCESS	Data is often siloed to business units. Do not start a project until full access to all data is secured for the entire team.
FORMAT	Data is found in different database formats, and different storage mediums. Often data is hiding in images and video.
PRIVACY	Security and privacy requirements within an enterprise or enforced by regulatory bodies.
QUANTITY	The amount of data needed for machine learning model development is almost always greater than available.
LABELING	A plan is needed to bring together the vast amount of manual labor and the domain knowledge to execute data labeling.

ACCESS	Data is often siloed to business units. Do not start a project until full access to all data is secured for the entire team.
FORMAT	**Data is found in different database formats, and different storage mediums. Often data is hiding in images and video.**
PRIVACY	Security and privacy requirements within an enterprise or enforced by regulatory bodies.
QUANTITY	The amount of data needed for machine learning model development is almost always greater than available.
LABELING	A plan is needed to bring together the vast amount of manual labor and the domain knowledge to execute data labeling.

ACCESS	Data is often siloed to business units. Do not start a project until full access to all data is secured for the entire team.
FORMAT	Data is found in different database formats, and different storage mediums. Often data is hiding in images and video.
PRIVACY	Security and privacy requirements within an enterprise or enforced by regulatory bodies.
QUANTITY	The amount of data needed for machine learning model development is almost always greater than available.
LABELING	A plan is needed to bring together the vast amount of manual labor and the domain knowledge to execute data labeling.

ACCESS

Data is often siloed to business units. Do not start a project until full access to all data is secured for the entire team.

FORMAT

Data is found in different database formats, and different storage mediums. Often data is hiding in images and video.

PRIVACY

Security and privacy requirements within an enterprise or enforced by regulatory bodies.

QUANTITY

The amount of data needed for machine learning model development is almost always greater than available.

LABELING

A plan is needed to bring together the vast amount of manual labor and the domain knowledge to execute data labeling.

ACCESS	Data is often siloed to business units. Do not start a project until full access to all data is secured for the entire team.
FORMAT	Data is found in different database formats, and different storage mediums. Often data is hiding in images and video.
PRIVACY	Security and privacy requirements within an enterprise or enforced by regulatory bodies.
QUANTITY	The amount of data needed for machine learning model development is almost always greater than available.
LABELING	A plan is needed to bring together the vast amount of manual labor and the domain knowledge to execute data labeling.

Lack of Customer Focus

THE Perfect Model?

Good Enough?

Customer Wants?

or Customer Needs?

Not Invented Here Syndrome

THE PURSUIT OF **HAND-CRAFTED PERFECTION** MISSES **EASY SUCCESS**

RESISTANCE TO **OFF THE SHELF** SOLUTIONS

ML as a Service

Offerings by many cloud providers and many other companies.

Open-Source Model

Start with an existing model known to solve a similar problem and build from there.

Auto ML

Machine learning to create machine learning models.

Lack of Explainability

Can your team explain your model?

If not, will your industry accept it?

Will regulatory bodies allow it?

Understand Business Needs

BUSINESS NEEDS

USER NEEDS	The business needs should be directly related to your users needs, if these are misaligned you have a problem.
BOTTOM LINE	What is the return on investment? Translate your ML project into an impact on the business bottom line.
CHAMPION	Find your business champion, if you are not able to find one you likely don’t have a viable project.
EXPECTATIONS	Manage expectations, make sure your promises are realistic in a short timeline.

BUSINESS NEEDS

USER NEEDS	The business needs should be directly related to your users needs, if these are misaligned you have a problem.
BOTTOM LINE	What is the return on investment? Translate your ML project into an impact on the business bottom line.
CHAMPION	Find your business champion, if you are not able to find one you likely don't have a viable project.
EXPECTATIONS	Manage expectations, make sure your promises are realistic in a short timeline.

BUSINESS NEEDS

USER NEEDS

The business needs should be directly related to your users needs, if these are misaligned you have a problem.

BOTTOM LINE

What is the return on investment? Translate your ML project into an impact on the business bottom line.

CHAMPION

Find your business champion, if you are not able to find one you likely don't have a viable project.

EXPECTATIONS

Manage expectations, make sure your promises are realistic in a short timeline.

BUSINESS NEEDS

USER NEEDS	The business needs should be directly related to your users needs, if these are misaligned you have a problem.
BOTTOM LINE	What is the return on investment? Translate your ML project into an impact on the business bottom line.
CHAMPION	Find your business champion, if you are not able to find one you likely don't have a viable project.
EXPECTATIONS	Manage expectations, make sure your promises are realistic in a short timeline.

Multiphase Project

POC

Start with a proof of concept, fail fast, pivot fast, and repeat until you have a feasible project.

MVP

Target a true minimal viable product and get it in front of users fast.

ITERATION

Plan on iterating quickly and moving fast. It is not uncommon to ship multiple models in a week.

POC

Start with a proof of concept, fail fast, pivot fast, and repeat until you have a feasible project.

MVP

Target a true minimal viable product and get it in front of users fast.

ITERATION

Plan on iterating quickly and moving fast. It is not uncommon to ship multiple models in a week.

POC

Start with a proof of concept, fail fast, pivot fast, and repeat until you have a feasible project.

MVP

Target a true minimal viable product and get it in front of users fast.

ITERATION

Plan on iterating quickly and moving fast. It is not uncommon to ship multiple models in a week.

Evaluate Existing Solutions

off the shelf

Evaluate all off the shelf solutions that are applicable to this project

don't hand craft

Initial machine learning projects should avoid hand crafting

pivot

If something simple and off the shelf won't work you have the wrong starting project

off the shelf

Evaluate all off the shelf solutions that are applicable to this project

don't hand craft

Initial machine learning projects should avoid hand crafting

pivot

If something simple and off the shelf won't work you have the wrong starting project

off the shelf

Evaluate all off the shelf solutions that are applicable to this project

don't hand craft

Initial machine learning projects should avoid hand crafting

pivot

If something simple and off the shelf won't work you have the wrong starting project

Create a Data Plan

EXPLORE

Start with exploratory data analysis

EVALUATE

COLLECT

PAUSE

FEATURE

EXPLORE

EVALUATE

Is this the correct data and is there enough data?

COLLECT

PAUSE

FEATURE

EXPLORE

EVALUATE

COLLECT

Collect any additional data needed to enable success before starting

PAUSE

FEATURE

EXPLORE

EVALUATE

COLLECT

PAUSE

Shift to another project and wait until enough data is collected

FEATURE

EXPLORE

EVALUATE

COLLECT

PAUSE

FEATURE

Feature engineering will be a significant portion of your project

The Right Team and Process

THE RIGHT TEAM

identify:

- ✓ Data Expert
- ✓ Data Scientist
- ✓ Business Sponsor
- ✓ ML Engineer
- ✓ MLOps Engineer
- ✓ Data Labeling Team

OWNERSHIP

The Team

Clearly establish the teams long term ownership and maintenance plan for the models.

PROCESS

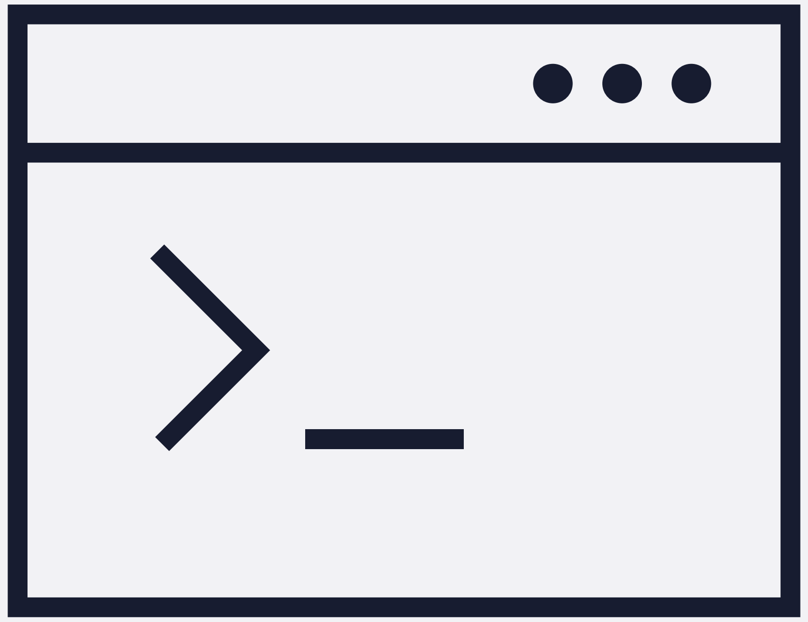

MLOps

Initial MLOps pipelines are in place and can deploy to a DEV environment.

SDLC

A team development process is in place and members of the team understand and participate in it.

Test and Release Planning

Model Testing

Have a plan for validating the model with labeled data not used during the training process.

Real World Validation

Have a small test group of real end users that can put your model through real world usage.

Have a Documented Test Plan

Have a RELEASE PLAN

Go / No Go

Establish metrics that will determine if a release to production has been successful.

A / B Transitioning

When updating models, move a small part of your user base at a time.

Rollback Plan

Have a plan to enable rolling back to a previous model or disabling its use.

Microsoft® Most Valuable Professional

OCTAVIAN TECHNOLOGY GROUP

Thank You for Watching!

- ✓ cameron@cameronvetter.com
- ✓ linkedin.com/in/cameronvetter
- ✓ twitter.com/poshporcupine

Microsoft Partner | Silver Data Analytics