Agile Machine Learning, in Production

Kishau Rogers

@kishau #deliverAgile2019 #AgileMachineLearning

TIME : STUDY

Kishau Rogers, Founder at Time Study Inc.

- Background: Computer Science, Entrepreneur, 24yrs delivering enterprise software solutions
- Sector: Enterprise SaaS, Health & Research
- Current Focus: Machine Learning at Time Study

TIME : STUDY

- We provide Enterprise tools for simplifying complex time sheets
- Our clients: U.S. Healthcare systems, Research enterprises
- ML: A computer is learning if its performance of a certain task, as measured by computable score, improves with experience
- How we're using Machine Learning to deliver value to our clients
  - Predicting activities
  - Predicting hours worked
  - Classification (departments, groups, roles)
  - Anomaly detection (validation of time)

Stanford HEALTH CARE
YALE-NEW HAVEN HOSPITAL
Lifespan Delivering health with care.
uchealth
MedStar Health
NewYork-Presbyterian
Stony Brook Medicine
PRESBYTERIAN

Agenda – Barriers for Agile Machine Learning

- 1: People
- 2: Tools
- 3: Data
- 4: Development
- 5: Production-Readiness
- Q&A

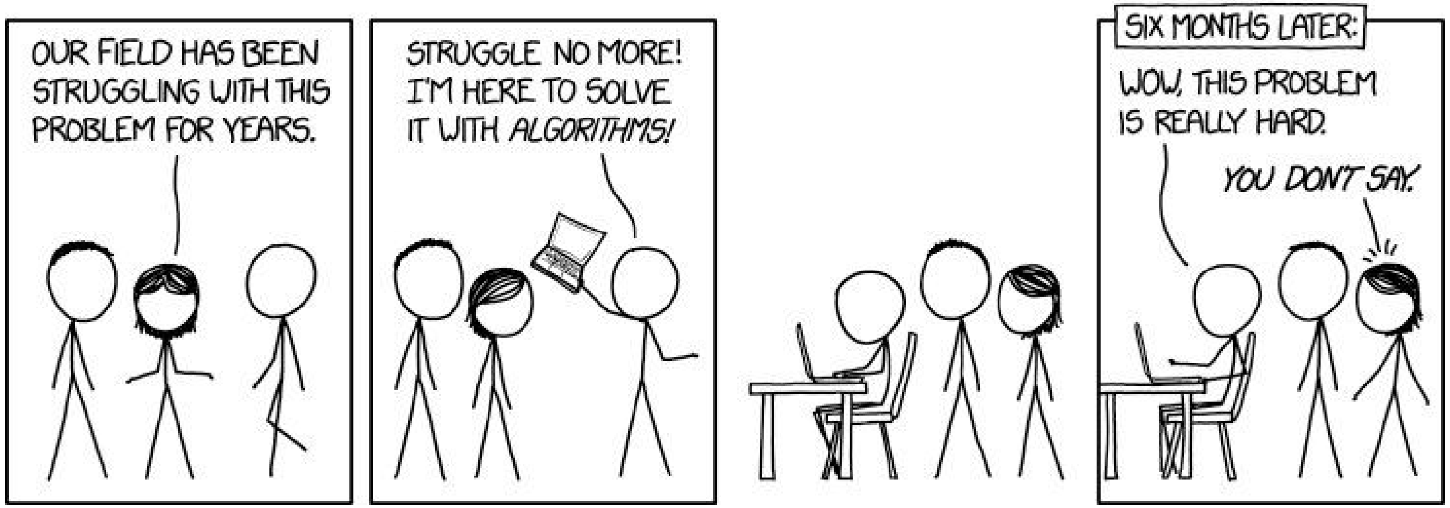

Lesson : Data Science fails to reach production when performed in silos

1: People – Collaborative ML Team

- Data Science fails to reach production when performed in silos
- Collaboration supports agility
  - Key Stakeholders (KS) – sponsor, establishing KPIs
  - Product manager (PM) – communicating vision, road-mapping
  - SME/domain expert (SME) – provide customer insights
  - Data Engineer (DE) – create & maintain the data marketplace used for model development
  - Data Scientist (DS) – researching/exploring, experimenting & developing models
  - IT/DevOps (IT) – create the environments to support self service data science, data governance
  - Developers (D) - end-to-end application dev., leveraging models in production

ML ROADMAP	CLASSIFY	ACQUIRE	PREPARE	BUILD	VALIDATE	DEPLOY	MONITOR
GOAL	Identify hypothesis	Acquire data assets & establishing context	Improve data quality	Develop an appropriate learning system	Identify & Reduce error	Present results	Monitor change
TASKS	Establishing KPI, baseline(s)	Discover, Gather Data	Ingest, Understand, Transform	Featuring Engineering, Model Dev	Training, Testing, Model Selection	Serving Results	Logging, Monitoring
TOOLS	Kanban Roadmap	Data Lake, Analytics	Warehouse, Azure ML Data Prep, Python	Python, R, ML Studio (or Service w/ Python),	ML Studio, Python, R	ML Services	Power BI
DELIVERY	Questions That Data Can Answer	Centralized Data Market	Useable Datasets for training	Model MVP & Algorithms	Model Scores & Results	Model API, Pipeline for New Data	Performance
TEAM	KS, PM, DS	DE, DS, IT	DS	DS	DS, KS, SME	DS, IT/D	KS, DS, IT/D
NEXT	Go/No Go	Research / Implement	Research / Implement	Experiments	Refine / Deploy	Apps, Integration	Retrain / Remodel

Lesson : Cross-functional product team. One product roadmap to assess business/customer value and (if useful) establish the appropriate use for ML tools now vs. later.

2: Tools vs. Value

- First, What's the problem? What's the question to be answered with data?
- Next, Assess current state.
- Next, Review impact and business/customer value.
- Finally, Determine if ML useful? (customers don't care about your tools) ML is a tool, not the product. Carefully define the problem & desired impact, then iterate toward that outcome
- Start small, scope down before scaling up

2: Why ML? When ML?

	PROBLEM CHARACTERISTICS	ML
INSIGHTS	How well do you understand this problem?	Absence of insights, ML becomes a Black box
COMPLEXITY	Can you code the rules (using a simple deterministic rule-based solution)? Is this a simple problem to solve? How many factors are involved?	Solution not easily coded & Involves Many Variables
ACCURACY	What accuracy rate is required? Mission critical? How quickly does your process need to adjust & learn from mistakes?	Enables Automated learning Accuracy is based on quality of data & model
SCALABILITY	Can a human perform this in a series of repeatable steps? Are you able to scale their efforts?	Offers speed & convenience Limited transparency
DATA ASSETS	Do you have the “right” data to “learn from”? Is the data clean (i.e. not “noisy”)	The “Right” Data is Required Requires data transformation & noise reduction
RESOURCES	Do you have resources to maintain your ML solution? Rhythm for ML?	Data Science, IT, SME
BOTTOM LINE	What is the business impact for solving this problem? What is the current state of the product? Product-market fit?	Works well for enhancing stable products, solving problems with significant bottom-line impact.

2: Asking Well-Formed Questions

QUESTION TYPE	SAMPLE QUESTIONS
Categorizing or Classifying Things	Will this person provide patient care or administrative time?
Predicting Values	How many hours will this person work next week/quarter/year?
Detecting Anomalies	Is this work week typical?
Grouping Related Data	What is the best way to group employees in 25 departments?
Determining the next best step	Where should this information be placed on the page so that the employee is most likely to click it?

ML ROADMAP	CLASSIFY	ACQUIRE	PREPARE	BUILD	VALIDATE	DEPLOY	MONITOR
GOAL	Identify hypothesis	Acquire data assets & establishing context	Improve data quality	Develop an appropriate learning system	Identify & Reduce error	Present results	Monitor change
TASKS	Establishing KPI, baseline(s)	Discover, Gather Data	Ingest, Understand, Transform	Featuring Engineering, Model Dev	Training, Testing, Model Selection	Serving Results	Logging, Monitoring
TOOLS	Kanban Roadmap	Data Lake, Analytics	Warehouse, Azure ML Data Prep, Python	Python, R, ML Studio (or Service w/ Python),	ML Studio, Python, R	ML Services	Power BI
DELIVERY	Questions That Data Can Answer	Centralized Data Market	Useable Datasets for training	Model MVP & Algorithms	Model Scores & Results	Model API, Pipeline for New Data	Performance
TEAM	KS, PM, DS	DE, DS, IT	DS	DS	DS, KS, SME	DS, IT/D	KS, DS, IT/D
NEXT	Go/No Go	Research / Implement	Research / Implement	Experiments	Refine / Deploy	Apps, Integration	Retrain / Remodel

Lesson : Data acquisition, discovery & preparation is half the battle..

3: Data Discovery & Prep is half the battle

- 70–80% of the effort in building an ML solution involves discovering, acquiring, cleaning and representing the data in a format specific to the use case
- Wrangling data from siloed sources
- High IT dependence
- Tech. debt for data work

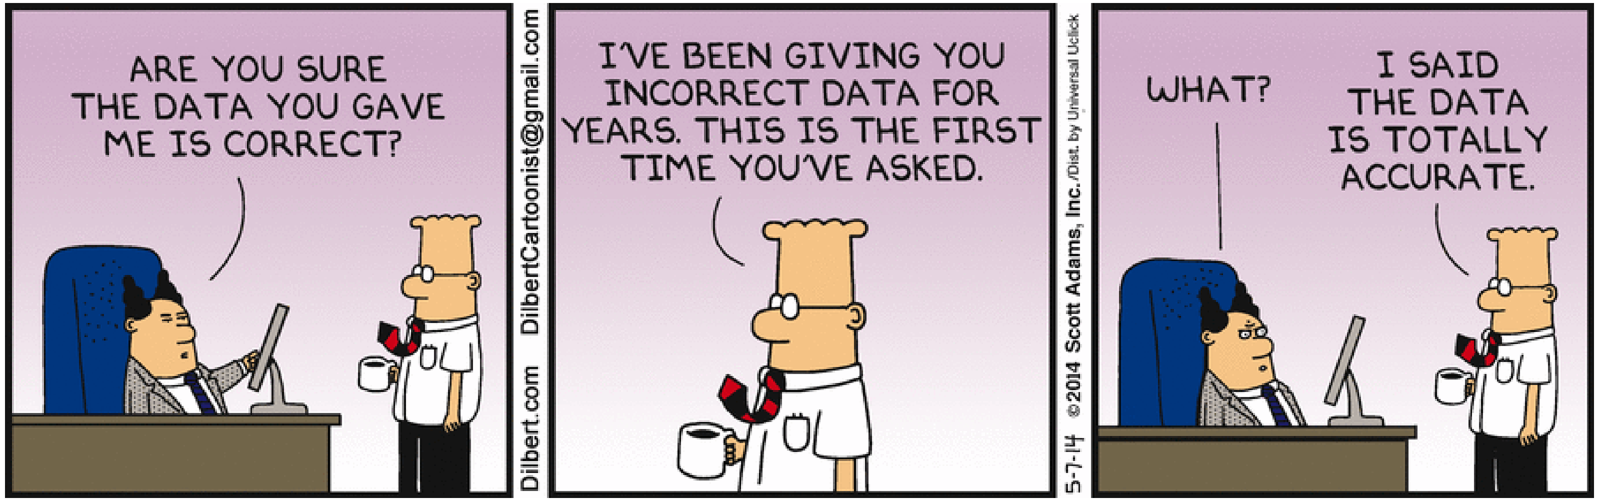

3: DataOps

- Data Marketplace - Having enough good data is fundamental to agile ML
- Data Discipline – Meta data maintenance supports understanding and visibility, enables governance, pipeline for new data, cleaning
- Building Data Assets Pipelines for ML Training & Testing – Data Marts for specific use cases. Usage drives architecture.
- Modular datasets (can be used across experiments)
- Time Box data exploration

ML ROADMAP	CLASSIFY	ACQUIRE	PREPARE	BUILD	VALIDATE	DEPLOY	MONITOR
GOAL	Identify hypothesis	Acquire, govern data, establish context	Improve data quality	Develop an appropriate learning system	Identify & Reduce error	Present results	Monitor change
TASKS	Establishing KPI, baseline(s)	Discover, Gather Data	Ingest, Understand, Transform	Featuring Engineering, Model Dev	Training, Testing, Model Selection	Serving Results	Logging, Monitoring
TOOLS	Kanban Roadmap	Data Lake, Analytics	Warehouse, Azure ML Data Prep, Python	Python, R, ML Studio (or Service w/ Python),	ML Studio, Python, R	ML Services	Power BI
DELIVERY	Questions That Data Can Answer	Centralized Data Market	Useable Datasets for training	Model MVP & Algorithms	Model Scores & Results	Model API, Pipeline for New Data	Performance
TEAM	KS, PM, DS	DE, DS, IT	DE, DS	DS	DS, KS, SME	DS, IT/D	KS, DS, IT/D
NEXT	Go/No Go	Research / Implement	Research / Implement	Experiments	Refine / Deploy	Apps, Integration	Retrain / Remodel

Lesson : Self service tools can streamline the model development & deployment process, enabling teams to work together in their core area of expertise.

4: Iterative Model Development

- Models can take 8-12 week to develop
- Defining the ML MVP..
- Build simply, quickly + Iterate toward the outcome. Analyze->Define approach->Refine->Measure. Define approach to make course correction from failed experiments and assumptions.
- Self service tools streamline model development process & allows DS to function in their core area of expertise.
- Share results of experiments to support model selection.

4: Self-Service Tools

- Tools based on:
  - Existing tech stack
  - Available team knowledge & skillset
  - Hiring pool
- Self service options to streamline data science process
  - Central Data "Marketplace"
  - Model building environment
  - Production sandbox
  - Models served via API

ML ROADMAP	CLASSIFY	ACQUIRE	PREPARE	BUILD	VALIDATE	DEPLOY	MONITOR
GOAL	Identify hypothesis	Acquire data assets & establishing context	Improve data quality	Develop an appropriate learning system	Identify & Reduce error	Present results	Monitor change
TASKS	Establishing KPI, baseline(s)	Discover, Gather Data	Ingest, Understand, Transform	Featuring Engineering, Model Dev	Training, Testing, Model Selection	Serving Results	Logging, Monitoring
TOOLS	Kanban Roadmap	Data Lake, Analytics	Warehouse, Azure ML Data Prep, Python	Python, R, ML Studio (or Service w/ Python),	ML Studio, Python, R	ML Services	Power BI
DELIVERY	Questions That Data Can Answer	Centralized Data Market	Useable Datasets for training	Model MVP & Algorithms	Model Scores & Results	Model API, Pipeline for New Data	Performance
TEAM	KS, PM, DS	DE, DS, IT	DS	DS	DS, KS, SME	DS, IT/D	KS, DS, IT/D
NEXT	Go/No Go	Research / Implement	Research / Implement	Experiments	Refine / Deploy	Apps, Integration	Retrain / Remodel

Lesson : Production readiness procedures enable deployed models to be leveraged by production apps.

5: Production Readiness Bottlenecks

- Re-implementation time - Experiments & models trapped on the data scientists' local machine
- Developing data flow pipelines for processing new data
- Monitoring for drift
- Refining models, version control
- Technical & configuration debt

"First thing Monday, we're gonna scale back the Machine Learning budget."

brianmooredraws.com

5: Production Readiness

STANDARD	DESCRIPTION
Scope	An clear indication of the well-formed question to be answered & stakeholder impact
Data Quality	Real-time access to the “right” data; automated process for cleaning data.
Model Quality	Model produces timely feedback, results are repeatable and explainable
Performance	Performance standards defined (accuracy, speed etc.). Tested for scalability. Automated data ingestion.
Maintainability	Well documented model versioning, deployment & configuration procedures.
Monitoring	Real-time alerts on performance and adaptability.
Security	Meets data governance & info security standards.

ML ROADMAP	CLASSIFY	ACQUIRE	PREPARE	BUILD	VALIDATE	DEPLOY	MONITOR
GOAL	Identify hypothesis	Acquire data assets & establishing context	Improve data quality	Develop an appropriate learning system	Identify & Reduce error	Present results	Monitor change
TASKS	Establishing KPI, baseline(s)	Discover, Gather Data	Ingest, Understand, Transform	Featuring Engineering, Model Dev	Training, Testing, Model Selection	Serving Results	Logging, Monitoring
TOOLS	Kanban Roadmap	Data Lake, Analytics	Warehouse, Azure ML Data Prep, Python	Python, R, ML Studio (or Service w/ Python),	ML Studio, Python, R	ML Services	Power BI
DELIVERY	Questions That Data Can Answer	Centralized Data Market	Useable Datasets for training	Model MVP & Algorithms	Model Scores & Results	Model API, Pipeline for New Data	Performance
TEAM	KS, PM, DS	DE, DS, IT	DS	DS	DS, KS, SME	DS, IT/D	KS, DS, IT/D
NEXT	Go/No Go	Research / Implement	Research / Implement	Experiments	Refine / Deploy	Apps, Integration	Retrain / Remodel

Summary

- 1: People – Collaborative Teams
- 2: ML Tools vs. Value – Determine when and where ML is appropriate
- 3: DataOps – Streamline data acquisition & data processing ops
- 4: Self Service Model Development – teams working together, in their core area of expertise
- 5: Production-Readiness – addressing performance & refinement constraints
- Questions?

Contact

- Kishau Rogers
- (e) kishau@timestudy.co
- @kishau