

Building systems to monitor model health in production systems

Mohammed Ridwanul Islam ([🐦 @mohammedri_](#))

Product Manager @Dessa

Hello DataCouncil 🖐️

Focus on building SDK and infrastructure tools for
Machine Learning Engineers

Previously: Data Platform Engineering @Shopify

We build Continuous Delivery Tools for Machine Learning

Foundations suite



Foundations Atlas

For Machine Learning
Development across teams



Foundations Orbit

For post-production model
monitoring & maintenance

We also have fun with Machine Learning

GIZMODO

VIDEO REVIEWS SCIENCE IO9 FIELD GUIDE EARTHER DESIGN PALEOFUTURE

ARTIFICIAL INTELLIGENCE

This AI-Generated Joe Rogan Voice Sounds Eerily Like the Real Thing



Jennings Brown

5/17/19 12:10pm • Filed to: DEEPFAKES

33.1K 28 Save



AI built by Canadians for the world.



Join the Canada.ai Slack channel

Send

Search Posts, Events, a

MENU

BROWSE CATEGORIES



Canadian AI General AI News Mar 15, 2019 • Sera Wong

Dessa Engineers Build Machine Learning Supernova Identification System

The system is improving the accuracy of identifying supernovas by 10 percent



Agenda

- Why do we care about ML monitoring?
- Why current systems are not good enough to monitor ML?
- What do we monitor?
- How we do monitoring at Dessa?
- Our systems approach and design
- Q&A

Why care about Model monitoring?

Personal goal: see more and more companies *RELEASE* products that are powered by Machine Learning *SUCCESSFULLY*

Lots of companies are building models but few are able to move them out of POC.

Why?

- Teams are getting burnt after putting models in production
 - Real life-example: F100 Enterprise
 - Lack of automated ways of maintaining SLAs - reduces trust on models in production
- Models drop in performance rapidly & managing more than 3 models per team becomes quite difficult manually
- Monitoring, diagnosis, rollback and rebuilding strategies are currently quite ad-hoc and not scalable to many models

Data Infrastructure Engineers responsible for building strong monitoring and performance management systems that are easy to use by Data Scientists.

What usually goes wrong

Accuracy (or your metric of choice) declines without explanation 🙄

- Is this a bug in the upstream data that the model is consuming?
 - If so, which part of the pipeline is broken? Is it the feature engineering steps?
- Is this a bug in the training source code?
 - Was a new model deployed?
 - Was a new model trained with a drastically different dataset?
 - Are we overfitting / underfitting? Has the model gone stale?
- Has the wrong model been put into production? How do I rollback?
- Is there an adversarial user?
- Where do I start? Who is responsible? Who knows this model / dataset best?

THIS IS YOUR MACHINE LEARNING SYSTEM?

YUP! YOU POUR THE DATA INTO THIS BIG
PILE OF LINEAR ALGEBRA, THEN COLLECT
THE ANSWERS ON THE OTHER SIDE.

WHAT IF THE ANSWERS ARE WRONG?

JUST STIR THE PILE UNTIL
THEY START LOOKING RIGHT.



Solution: continuous testing

- Create tests & alerts on upstream datasets & models for quality
- Centralized model management with efficient rollback strategies
- Evaluate models using real outcomes and maintain pre-defined SLAs through re-training & recalibration
- Make it easy for your scientists!



Monitoring needs for ML

Traditional software vs. ML

- Source code changes (track using CI tools, build pipelines etc.)
- Infrastructure / Network Changes (track using DataDog, Splunk et. al.)
- Source code changes
- Infrastructure / Network Changes
- Bugs in the Data Pipeline
- Dataset shifts
 - Co-variate shifts
 - Prior probability shifts
 - Concept Drift
 - Uncertainty
 - Explainability
 - + more!

What needs to be monitored?

Bugs in the data pipeline

- Upstream changes impacting downstream data consumption e.g.
 - Schema changes:
 - Type changes
 - Column name changes
 - New features added/deleted
 - Systematic errors introduced

Bugs in the data pipeline - monitoring strategies

- Schema checks:
Compare schemas of incoming data to a reference dataset
- Special Value Checks:
E.g. % of nulls in a column, NaNs, -999999, -1, etc.
- Min-max ratio checks
- Uniqueness checks
- Special knowledge checks

Dataset Drift

Input Drift


Concept Drift

Input Drift

Input shift refers to the change in the distribution of the input variables present in the training and inference data

1. May or may not correspond to a drop in model performance
2. Indicates either a shift in the underlying data distribution or bugs in the data pipeline
3. Monitoring input drift can help you be proactive in choosing better features in the future

$$p_{train}(X) \neq p_{inference}(X)$$



On the Reduction of Biases in Big Data Sets for the Detection of Irregular Power Usage - Scientific Figure on ResearchGate. Available from: https://www.researchgate.net/figure/Example-of-covariate-shift-training-and-test-data-having-different-distributions_fig1_324168804 [accessed 12 Nov, 2019]

Input Drift - various monitoring strategies

1. Population Stability Index (PSI)
 - a. Based on Kullback–Leibler (KL) Divergence
 - b. Complex to interpret
2. L-Infiniti
 - a. Much more easily interpretable e.g. “allow changes of only up to 1% for each (feature)” [1]
3. Similarity score
 - a. Lots of other approaches - many papers have been written

[1] Eric Breck et. al., *Proceedings of SysML* (2019) (to appear) DATA VALIDATION FOR MACHINE LEARNING, 2019

Input Drift - various monitoring strategies

Choose the metric that is easiest to understand and use.

"A first-cut solution here might be to use typical distance metrics such as KL divergence or cosine similarity and fire an alert only if the metric crosses a threshold. The problem with this solution is that product teams have a hard time understanding the natural meaning of the metric and thus tuning the threshold to avoid false positives..."

...(distribution distance metrics instead) has a simple natural interpretation of the largest change in probability for a value in the two distributions, which makes it easier for teams to set a threshold (e.g., "allow changes of only up to 1% for each value")

Concept Drift

Occurs when the function mapping input to target has changed

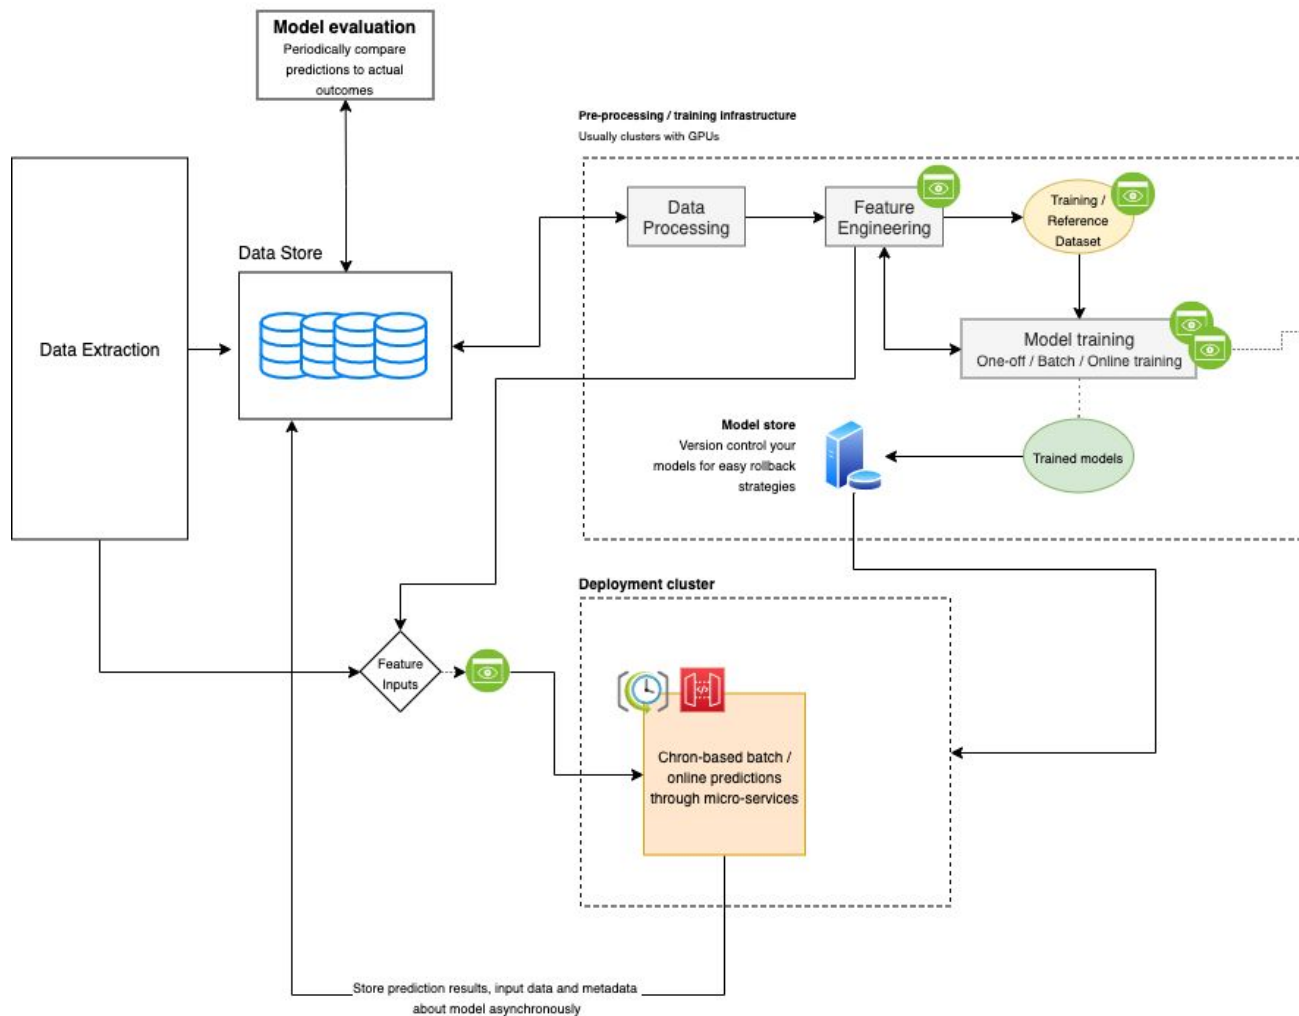
1. Corresponds to a drop in model performance
2. Indicates that a model may need to be trained
3. Change can take many forms[2]:
 - a. A gradual change over time.
 - b. A recurring or cyclical change.
 - c. A sudden or abrupt change.

$$p_{train}(Y|X) \neq p_{inference}(Y|X)$$

Source:

[2]: Jason Brownlee, <https://machinelearningmastery.com/gentle-introduction-concept-drift-machine-learning/>

**We need to be able to codify expectations
from our datasets**



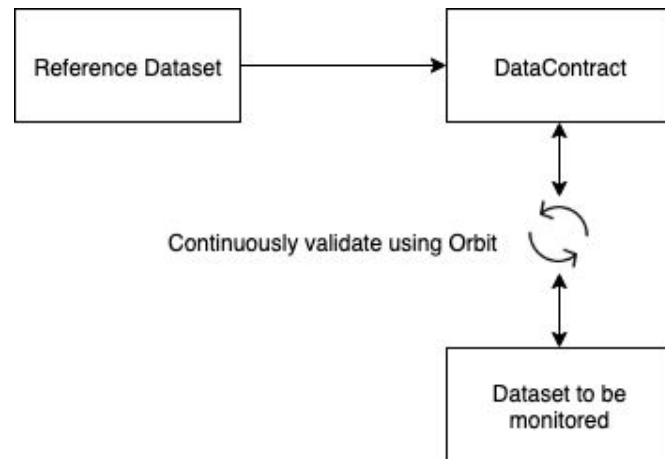
Orbit = Python SDK + Scheduler + GUI

DataContracts as expectations using the Orbit SDK

To monitor bugs in the pipeline and Dataset Drifts

Orbit Python SDK - DataContracts

- DataContracts = Data Expectations
- Every DataContract object ~= configurable metadata store, versioned & immutable
- Uses reference dataframe to compare incoming data and generate validation reports
- Configure and set expectations for any *reference* data set



```
from foundations_orbit import DataContract
```

```
data_contract = DataContract(data_contract_name,  
reference_dataframe)
```

```
data_contract.save(dir_path_to_save) # Save contracts in a  
registry
```

```
contract.distribution_test.configure(["col1", "col2"],  
threshold=0.2, method=my_custom_method())  
contract.distribution_test.configure(["col1", "col3"],  
threshold=0.1, method='l-infinity')  
contract.special_value_test.configure(["col1"],  
thresholds={np.nan: 0.1, -1: 0.2})
```

```
validation_report =  
data_contract.validate(dataframe_to_validate)
```

Product Recommendation Engine

Model Management

Data Validation Results

C

| Time | Monitor Name | Contract Name | |
|------------|--------------|------------------|---|
| 2019-08-22 | monitor_1 | input_contract_1 | 0 |
| 2019-08-22 | monitor_1 | input_contract_2 | 0 |
| 2019-08-22 | monitor_1 | input_contract_3 | 2 |
| 2019-08-22 | monitor_1 | output_contract | 0 |
| 2019-08-22 | monitor_2 | input_contract_a | 0 |
| 2019-08-22 | monitor_2 | input_contract_b | 0 |
| 2019-08-22 | monitor_2 | input_contract_c | 2 |
| 2019-07-22 | monitor_1 | input_contract_1 | 0 |
| 2019-07-22 | monitor_1 | input_contract_2 | 0 |
| 2019-07-22 | monitor_1 | input_contract_3 | 2 |
| 2019-07-22 | monitor_1 | output_contract | 0 |
| 2019-07-22 | monitor_2 | input_contract_a | 0 |
| 2019-07-22 | monitor_2 | input_contract_b | 0 |
| 2019-07-22 | monitor_2 | input_contract_c | 2 |
| 2019-06-22 | monitor_1 | input_contract_1 | 0 |
| 2019-06-22 | monitor_1 | input_contract_2 | 0 |
| 2019-06-22 | monitor_1 | input_contract_3 | 2 |
| 2019-06-22 | monitor_1 | output_contract | 0 |
| 2019-06-22 | monitor_2 | input_contract_a | 0 |
| 2019-06-22 | monitor_2 | input_contract_b | 0 |
| 2019-06-22 | monitor_2 | input_contract_c | 2 |
| 2019-05-22 | monitor_1 | input_contract_1 | 0 |
| 2019-05-22 | monitor_1 | input_contract_2 | 0 |
| 2019-05-22 | monitor_1 | input_contract_3 | 2 |

Overview

input_contract_3 ☐

Monitor Name: monitor_1
Job ID: 1dfjsk23k1as
Time: 2019-08-22 12:12 AM
User: User
Row count: 1,111,111 → 1,222,213 (+9%)

| | |
|--------------------|---------------------------------------|
| Schema Check | 0 Critical 68 Healthy 0 Warning |
| 1 Population Shift | 1 Critical 67 Healthy 0 Warning |
| Special Values | 0 Critical 68 Healthy 0 Warning |
| Min / Max | 0 Critical 68 Healthy 0 Warning |
| 1 Uniqueness | 1 Critical 67 Healthy 0 Warning |
| Domain | 0 Critical 68 Healthy 0 Warning |
| Custom Test | 0 Critical 68 Healthy 0 Warning |

Concept Drift - monitoring strategies

- Continuously monitor the metric you are optimizing for
- Continuously evaluate metric with actual outcome


Logic to calculate these metrics

```
foundations.track_production_metrics("accuracy",  
{str(eval_date): accuracy})  
foundations.track_production_metrics("roc_auc",  
{str(eval_date): roc_auc})  
foundations.track_production_metrics("revenue",  
{str(eval_date): revenue})  
foundations.track_production_metrics("n_active_custs",  
{str(eval_date): n_active_custs})
```

NUMBER OF DASHBOARD METRICS: 4


MSE OVER TIME

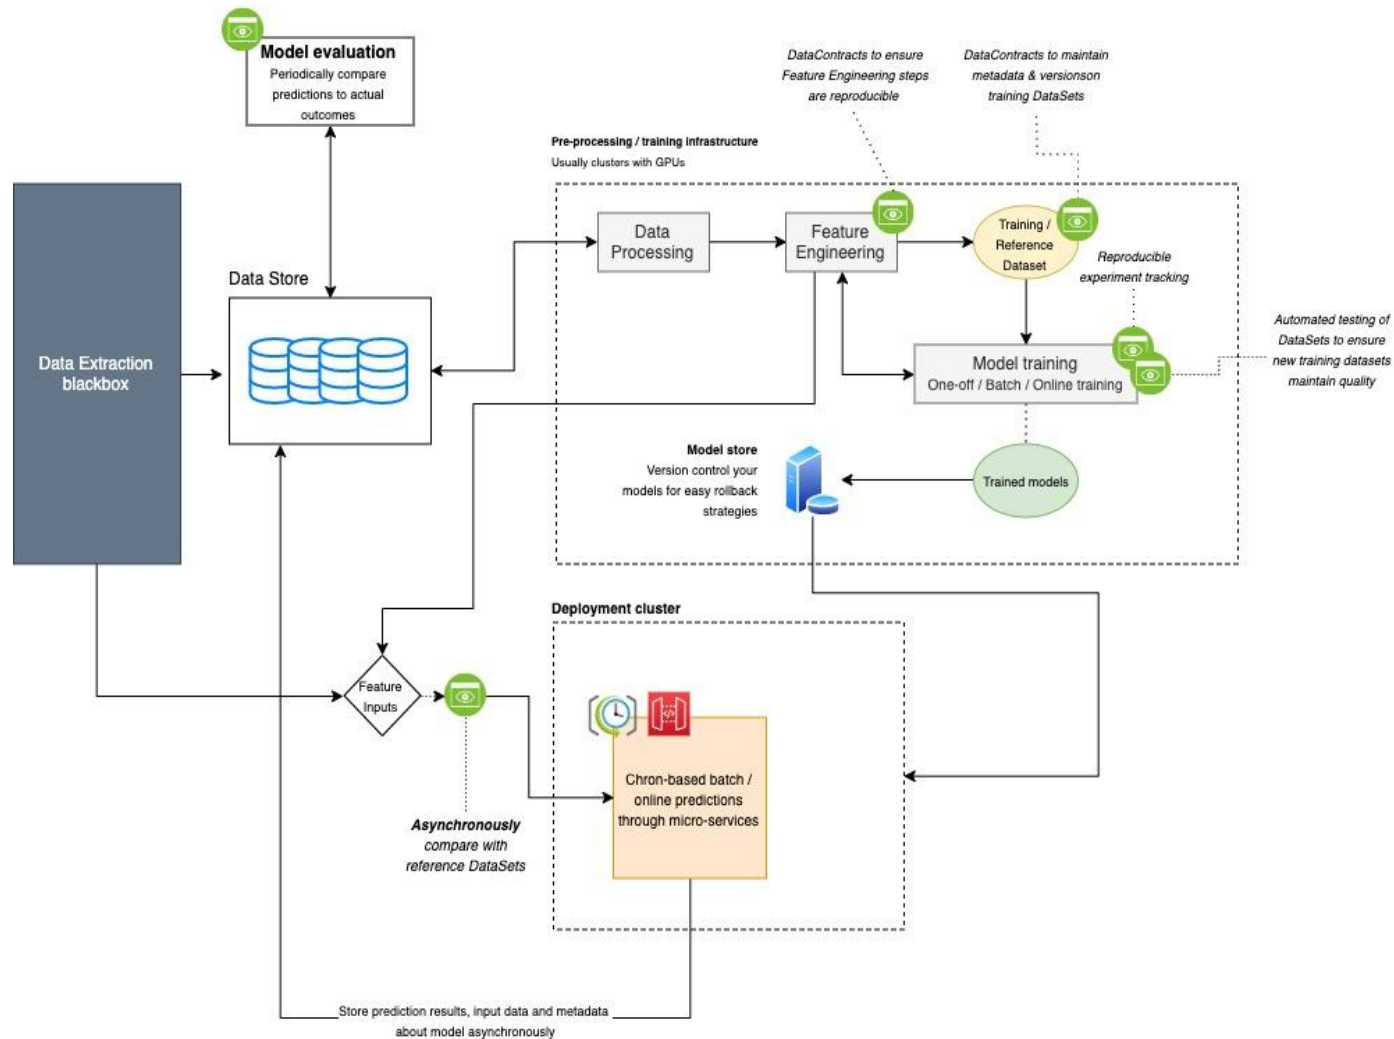
1 predictors



COST METRICS (\$) OVERTIME

1 predictors

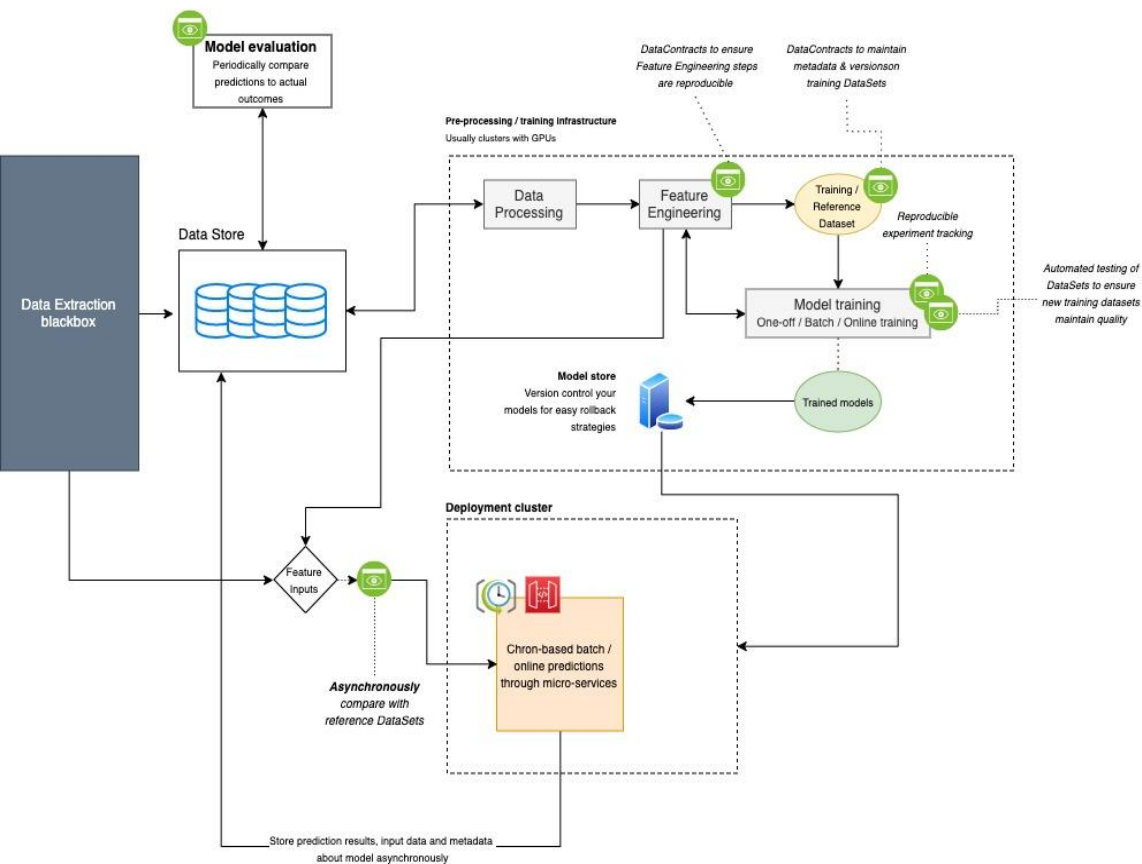




Other criterias for our monitoring system

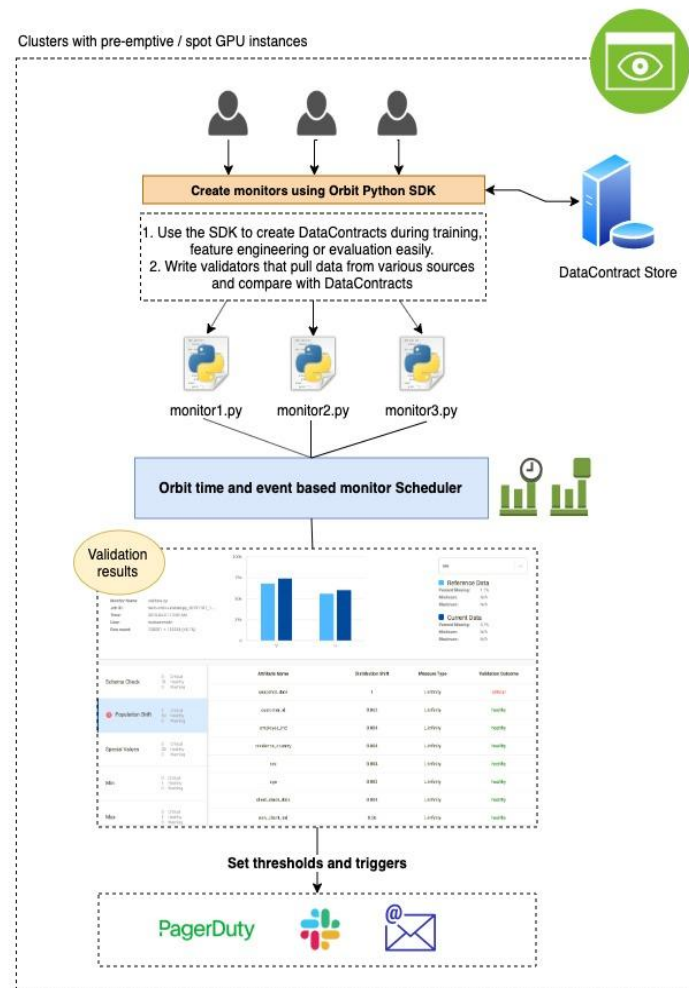
- Sidecar + Async architecture
- Built-in-standard monitoring such as schema checks, min-max's, L-infiniti etc.
- Ability to define custom monitors through lambda functions
- Chron based scheduler for monitoring
- Event based monitoring
- Dashboards for easy viewing and the ability to set thresholds and alerts

(Very) simplified machine learning workflow



Monitoring sidebar cluster

Clusters with pre-emptive / spot GPU instances



Data Scientists workflow


- Create a DataContract for a reference or training data set

```
21 # load all training data
22 train_dfs = []
23 for date in load_dates:
24     date_without_time = date.date()
25     # load one month of train data from a GCP bucket as a pandas dataframe
26     df = db.load(f'{data_key}-labelled-{date_without_time}')
27     # don't include customers who got promos in the training data
28     train_df = df[df['promo'] == 0]
29     train_df = train_df.drop(['promo'], axis=1)
30     train_dfs.append(train_df)
31 train_df = pd.concat(train_dfs, axis=0)
32
33 # split features from target
34 y_train = train_df['churn']
35 x_train = train_df.drop(['cust_id', 'date', 'churn', 'predicted_churn_probability'], axis=1)
36
37 # drop nans
38 x_train = x_train.fillna(0) # ensure the model can still train if NaN values are present
39
40 # train model
41 model = RandomForestClassifier()
42 model.fit(x_train, y_train)
43
44 # save model
45 pickle.dump(model, open("fitted_objects/model.pkl", "wb"))
```

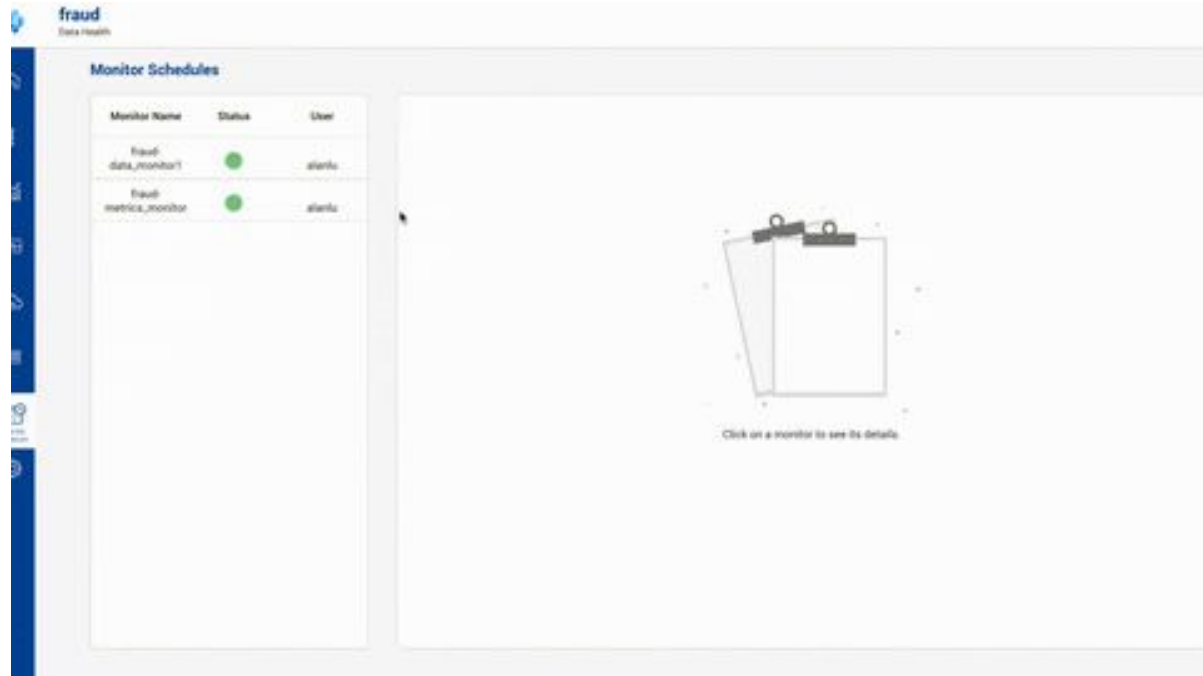

- Place DataContract in validation pipelines separate as monitoring code

```
train_driver.py train.py model_eval.py validate.py
1 import ...
12
13
14 def validate(inference_date):
15
16     inference_date = inference_date.split(' ')[0]
17     # load inference data from a GCP bucket as a pandas dataframe
18     inference_df = db.load(f'{data_key}-inference-{inference_date}')
19     # drop non-feature columns
20     x_inference = inference_df.drop(['cust_id', 'date'], axis=1)
21     x_inference = x_inference.fillna(0)
22
23     # load model from local directory
24     model = pickle.load(open("fitted_objects/model.pkl", "rb"))
25
26     # run inference, and create a dataframe to store the predicted probabilities
27     probs = model.predict_proba(x_inference)[: , 1]
28     probs_df = pd.DataFrame(probs, columns=["predicted_churn_probability"])
29
30     # save predictions to a GCP bucket
31     db.save(f'{data_key}-predictions-{inference_date}', probs_df)
32
```

- Deploy validation pipelines w/ DataContracts to Orbit via a REST backend



- Trigger monitoring using cron scheduler or event based triggers e.g. when a new dataset available




- Dig into reports or set triggers on data quality

fraud
(Data Health)

Data Validation Results

| Date | Monitor Name | Constraint Name | |
|------------|---------------|-----------------|--|
| 2019-03-01 | Data_Monitor1 | req_constraint | |
| 2019-03-01 | Data_Monitor1 | req_constraint | |
| 2019-03-01 | Data_Monitor1 | req_constraint | |
| 2019-03-01 | Data_Monitor1 | req_constraint | |



Click on a validation result to see its details.

Summary

Thank you :)

@mohammedri_