Welcome

BEAM SUMMIT Austin, 2022

Welcome to Beam Summit

Pablo Estrada & Danielle Syse

BEAM SUMMIT Austin, 2022

Going since 2018...

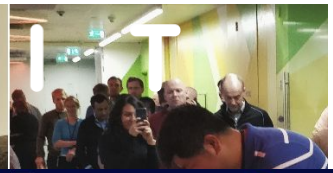

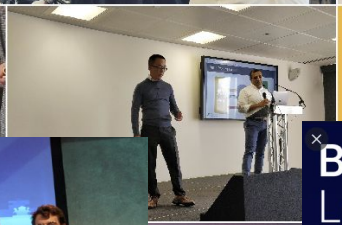

Beam Summit 2018

Photos from last year's talks, hands-on learning sessions.
#beamsummit

Thanks to our speakers!

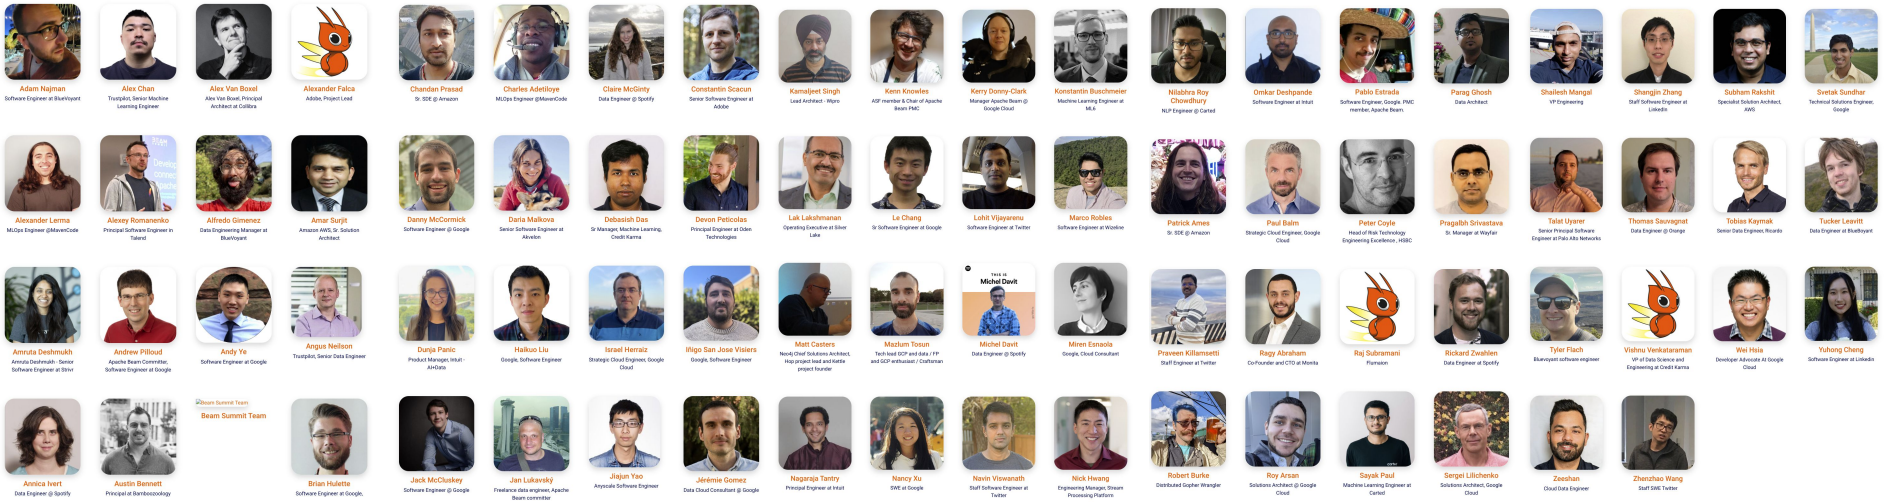

Thanks to our Sponsors

Google Cloud

GOLD

SILVER

Thanks to our partners

PARTNERS

THE APACHE SOFTWARE FOUNDATION

AKVELON
We make BIG software happen

Packt

MANNING

monita

AICamp

Data Council

WOMEN WHO CODE

ChickTECH

KEYNOTE SESSIONS

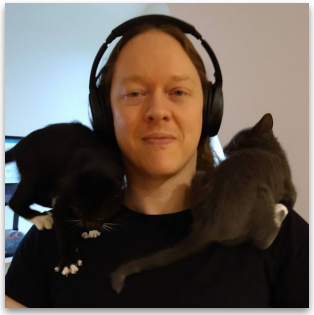

Kerry Donny-Clark

Manager Google Cloud
Dataflow

GOOGLE'S INVESTMENT ON BEAM AND ITS INTERNAL USE

10:00 - 10:25 AM

Lak Lakshmanan

Operating Executive
Silver Lake

MACHINE LEARNING DESIGN PATTERNS: BETWEEN BEAM AND A HARD PLACE

10:25 - 10:50 AM

Rickard Zwahlen

Data Engineer
Spotify

TAILORING PIPELINES AT SPOTIFY

10:50 - 11:15 AM

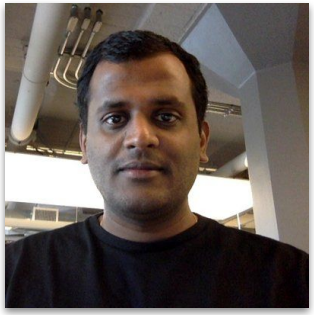

Lohit Vijayarenu

Principal Software Engineer
Twitter

THE ADOPTION, CURRENT STATE, AND FUTURE OF APACHE BEAM

11:15 - 11:40 AM

Before anything..!

Please, PLEASE fill our survey:

Monday Schedule

11:40 — Break

12:00
- 12:00 - 12:50 — Vega: Scaling MLOps Pipelines at Credit Karma using Apache Beam and Dataflow — by Debasish Das & Vishnu Venkataraman
- 12:00 - 12:50 — Houston, we've got a problem: 6 principles for pipelines design taken from the Apollo missions — by Israel Herraiz & Paul Balm
- 12:00 - 12:25 — RunInference: Machine Learning Inferences in Beam — by Andy Ye
- 12:30 - 12:55 — Speeding up development with Apache Beam (Adobe Experience Platform) — by Constantin Scacun & Alexander Falca

13:00 — Lunch

14:00
- 14:00 - 14:50 — Powering Real-time Data at Intuit: A Look at Golden Signals powered by Beam — by Omkar Deshpande, Dunja Panic, Nick Hwang & Nagaraja Tantry
- 14:00 - 14:50 — How the sausage gets made: Dataflow under the covers — by Pablo Estrada
- 14:00 - 14:25 — State of the Go SDK 2022 — by Robert Burke
- 14:30 - 14:55 — How to break Wordle with Beam and BigQuery — by Inigo-san-jose

15:00
- 15:00 - 15:50 — BlueVoyant: Detecting Security Dumpster Fires on the Internet — by Alfredo Gimenez, Adam Najman, Tucker Leavitt & Tyler Flach
- 15:00 - 15:50 — Migration Spark to Apache Beam/Dataflow and hexagonal architecture + DDD — by Mazlum Tosun
- 15:00 - 15:25 — Introduction to performance testing in Apache Beam — by Alexey Romanenko
- 15:30 - 15:55 — From script slums to beam skyscrapers — by Shailesh Mangal

16:00 — Break

16:15
- 16:15 - 16:40 — Data Integration on cloud made easy using Apache Beam — by Parag Ghosh
- 16:15 - 16:40 — How to benchmark your Beam pipelines for cost optimization and capacity planning — by Roy Arsan
- 16:15 - 17:15 — Cloud Spanner change streams and Apache Beam — by Haikuo Liu, Nancy Xu & Le Chang
- 16:45 - 17:10 — Collibra's Telemetry Backbone - OpenTelemetry and Apache Beam — by Alex Van Boxel
- 16:45 - 17:10 — Strategies for caching data in Dataflow using Beam SDK — by Zeeshan

17:15
- 17:15 - 18:05 — New Avro serialization and deserialization in Beam SQL — by Talat Uyarer
- 17:15 - 18:05 — Implementing Cloud Agnostic Machine Learning Workflows with Apache Beam on Kubernetes — by Charles Adetiloye & Alexander Lerma
- 17:15 - 18:00 — Cloud Spanner change streams and Apache Beam (continued) — by Haikuo Liu, Nancy Xu & Le Chang

18:05 — Reception 18:05 - 20:00 hrs

Where to Go Next?

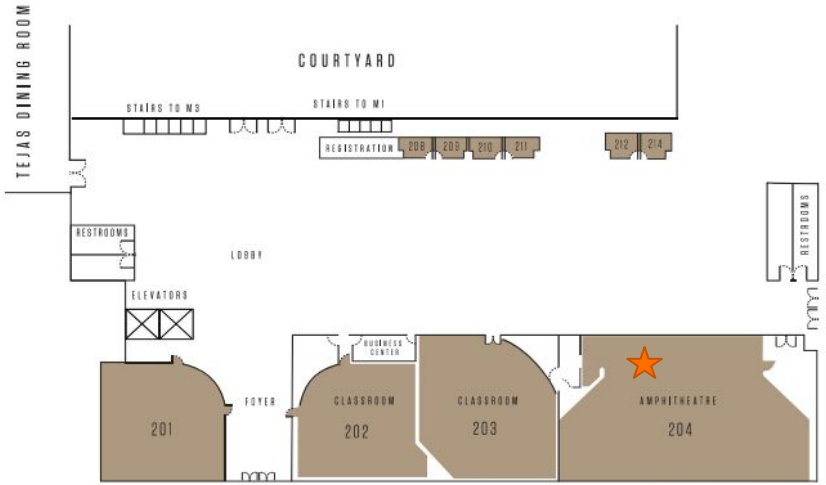

- All sessions will take place on this floor across 202-204
  - Keynotes will be held in the Amphitheatre only
- Lunch will take place from 1-2 PM in the Tejas Dining Room
  - Lunch box options include roast beef on ciabatta, chicken salad croissant and falafel fritter wraps
- Session rooms will be noted outside each door as well on each calendar invite/Beam Summit page
- Restrooms located at each end of the hall with elevators to our left next to the Dining Room

About the space...

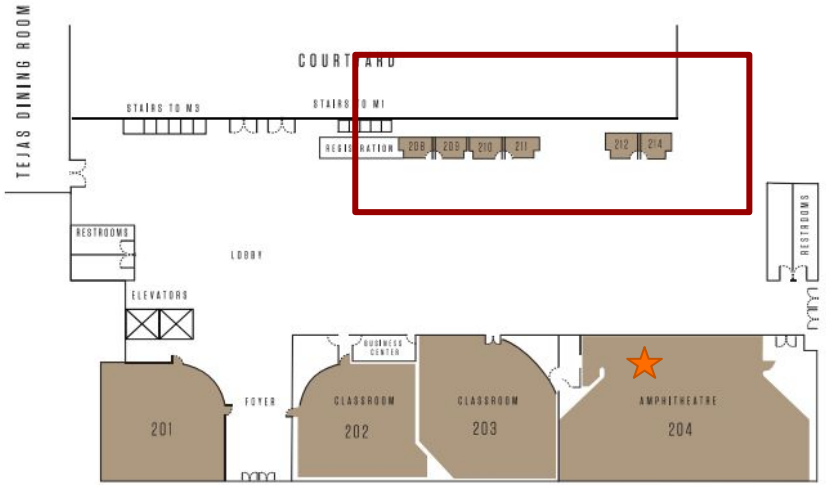

- We have rooms with whiteboards across the hall. Feel free to use the whiteboards.
  - We also have easel pads in the presentation rooms. Feel free to use in technical convos.

Networking Opportunities

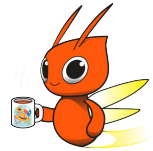

Please join us for networking opportunities while you're with us:

Reception tonight!

Join us for drinks after the event from **6:00 - 8:00 pm** at the **AT&T Conference Center Courtyard.**

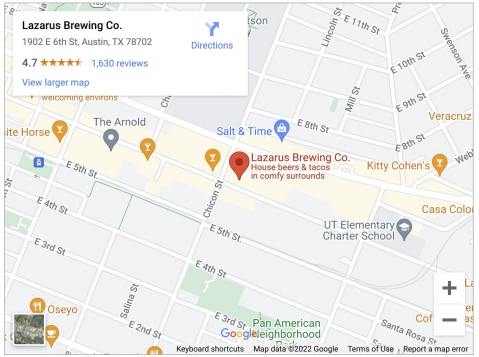

After Party Tuesday

Tuesday at 6:30pm at Lazarus Brewing Co, where beer on the house will be waiting for you! Send the directions to your phone by scanning the following QR code.

Networking Opportunities

Job Openings

Reminder to take a look at the current job openings gathered by our sponsors:

Speakers!

- Please arrive **a little early** to your room for setting up

Thank You

BEAM SUMMIT Austin, 2022

BEAM SUMMIT Austin, 2022

Google's investments in Beam

By Dr. Kerry Donny-Clark,
Google Engineering Manager for Beam

BEAM SUMMIT Austin, 2022

Hello!
I'm Kerry

Me at work

Me at home

My old job

Apache Beam

- Unified Model:
  - Batch and Streaming
- Many SDKs
  - Java, Python, Go, TS*
- Portability
  - Dataflow, Flink, Spark, Hazelcast, Ray*, Dask*, etc

*Experimental or in progress

Apache Beam used in Google

DeepMind

vertex.ai

Dataplex

YouTube

waze

The Beam Team at Google

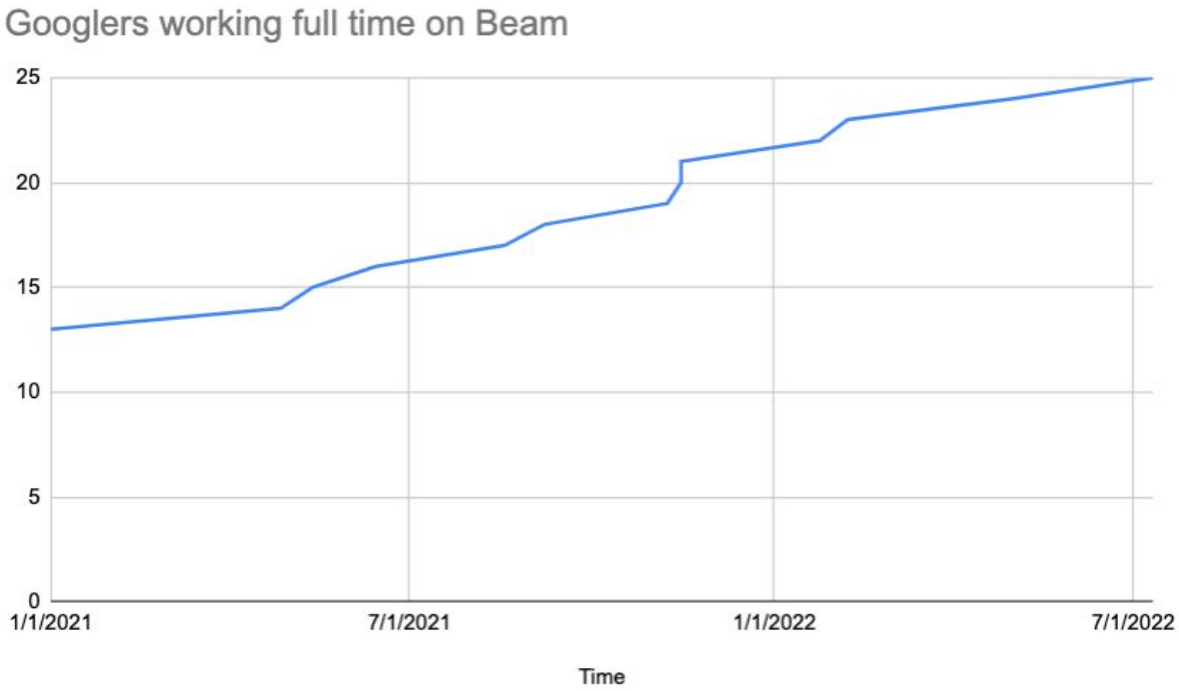

The Beam Team at Google


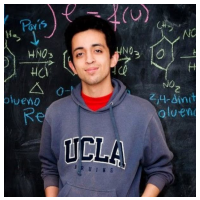










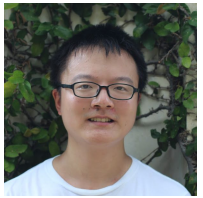






And more!

Beam going in new directions: Multi Language pipelines

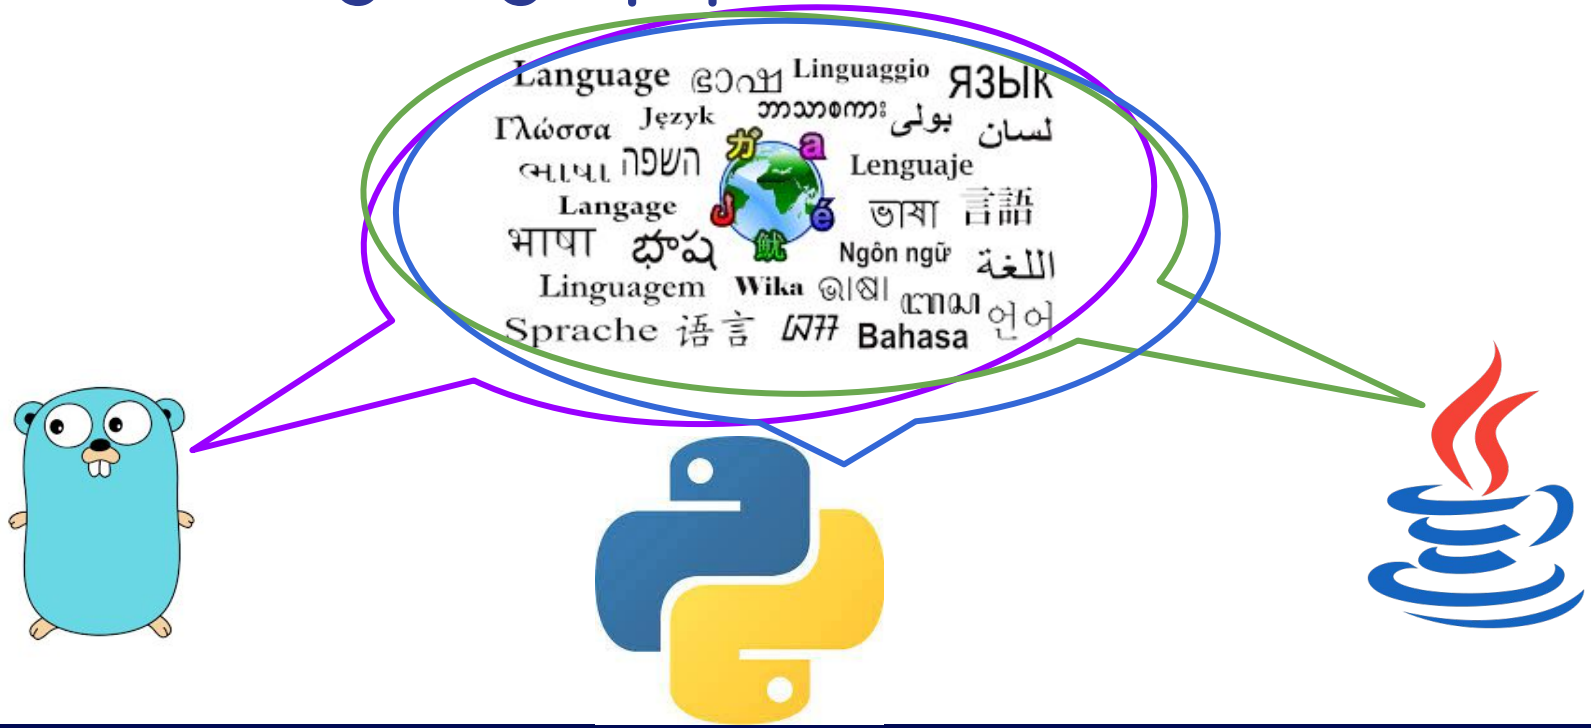

A Rebus Riddle

Beam 2.40, Dataflow GA 7/20

Beam Going in new directions

14:00 - 14:25. Today
State of the Go SDK 2022
by Robert Burke
Room: 202

16:15 - 16:40. Tuesday
Writing a native Go streaming pipeline
by Danny McCormick & Jack McCluskey
Room: 203

Beam going in new directions: RunInference in Beam Python

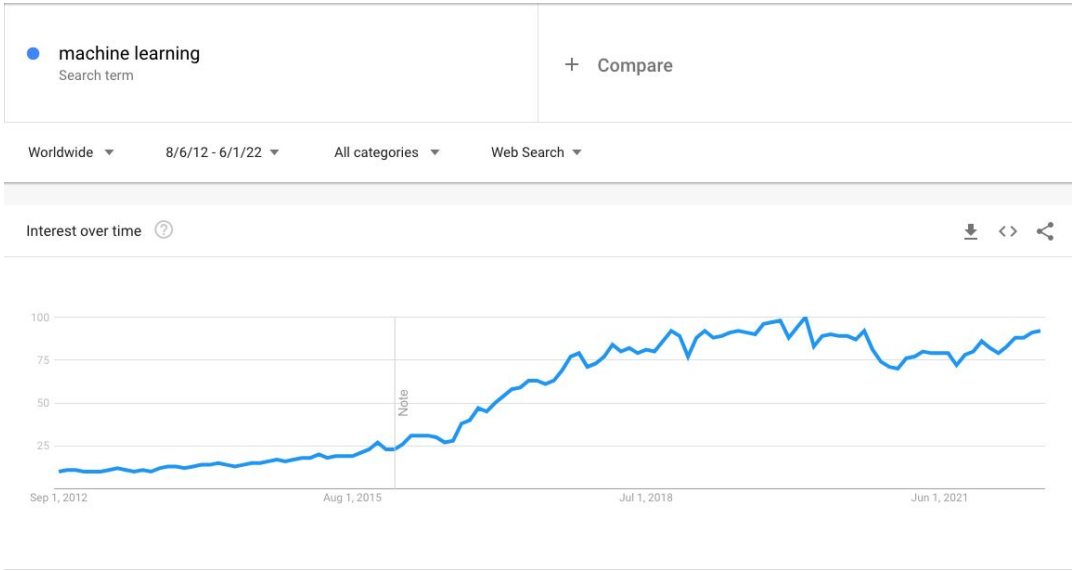

Beam going in new directions: RunInference in Beam Python

PYTORCH

TensorRT

TensorFlow

scikit learn

Beam going in new directions: RunInference in Beam Python

Today

12:00 - 12:25.

RunInference: Machine Learning Inferences in Beam

by Andy Ye

Room: 202

RunInference in Beam 2.40, GA on Dataflow 7/20

https://beam.apache.org/documentation/sdks/python-machine-learning/

Beam going in new directions: Typescript SDK

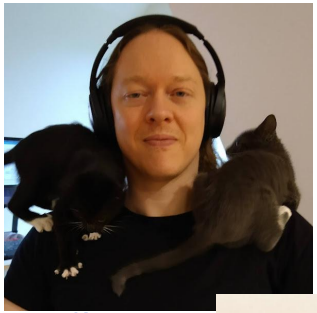

Kerry Donny-Clark

Jack McCluskey

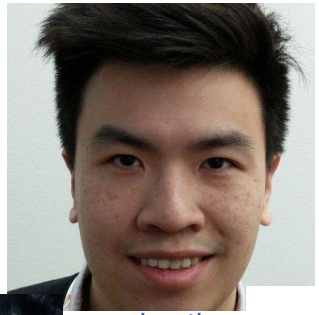

Jonathan Lui

Kenneth Knowles

Kevin Puthusseri

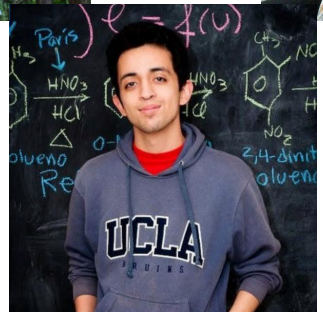

Pablo Estrada

Robert Bradshaw

Beam going in new directions: Typescript SDK

Beam going in new directions: Typescript SDK

... to contribute!

https://github.com/apache/beam/tree/master/sdks/typescript

A better way to learn Beam: Beam Playground

https://play.beam.apache.org/

A better way to learn Beam: Beam Playground

11:00 - 11:10. Wednesday

Beam Playground: discover, learn and prototype with Apache Beam

by Daria Malkova

Room: 201

A better way to learn Beam: Cloud notebooks

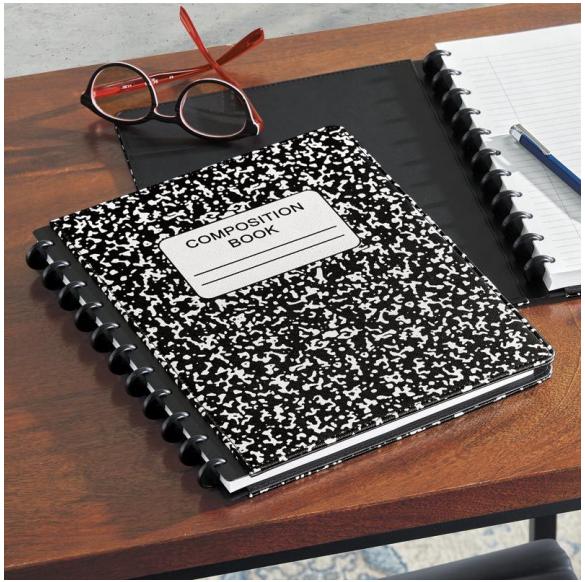

https://cloud.google.com/dataflow/docs/guides/interactive-pipeline-development

A better way to learn Beam:
A Tour of Beam

Coming in late 2022!

Contributing to Beam has never been easier: Github Issues

Inbox Apache Contributor permission for JIRA - Beam connecto
Inbox Apache Contributor permission for JIRA - org/jira/browse
Inbox Apache Contributor permission for Beam Jira tickets - ad
Inbox Apache Contributor permission for Beam Jira tickets - org
Inbox Apache Jira contributor permission - request > contributo
Inbox Apache Jira contributor permission - you to Jira. Thanks,
Inbox Apache Jira - contributor permission - you to Jira. Thanks
Inbox Apache RE: Re: Contributor permission for Jira tickets - is
Inbox Apache Contributor permission for Jira tickets - is your jir
Inbox Apache Jira Contributor Permission Request - Please gran
Inbox Apache Contributor permission for Beam Jira tickets - as
Inbox Apache Contributor permission for Beam Jira tickets - My
Apache Contributor permission for Beam Jira tickets - as a contr
Inbox Apache Re: Contributor permission for Beam Jira Tickets
Inbox Apache Re: Contributor permission for Beam Jira Tickets

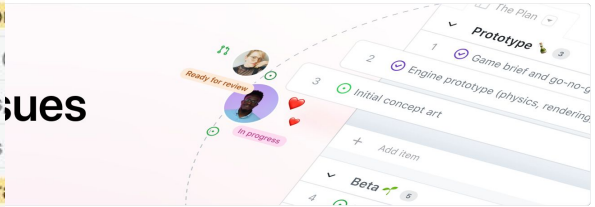

Contributing to Beam has never been easier: PR-bot

Conclusion

Beam is growing

- Multi Language
- Beam Go SDK
- RunInference in Python
- TypeScript SDK

Learn Beam

- Beam Playground
- Beam Notebooks
- A Tour of Beam

Contribute to Beam

- Github Issues
- PR-bot

BEAM SUMMIT Austin, 2022

Machine Learning Design Patterns: Between Beam and a Hard Place

Lak Lakshmanan

@lak_luster

BEAM SUMMIT Austin, 2022

Formalized best practices to solve common problems

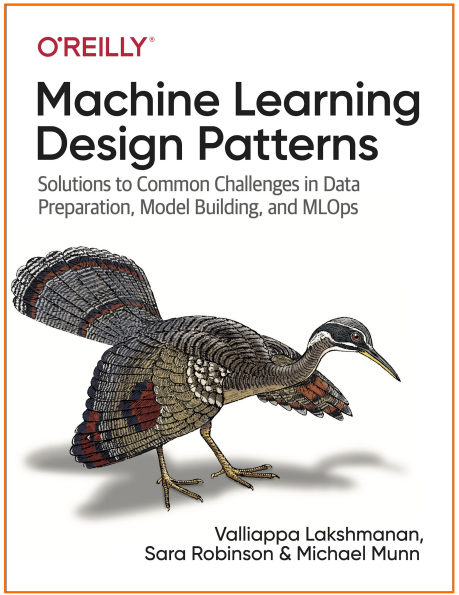

ML flavors of the same problems that arise in all software

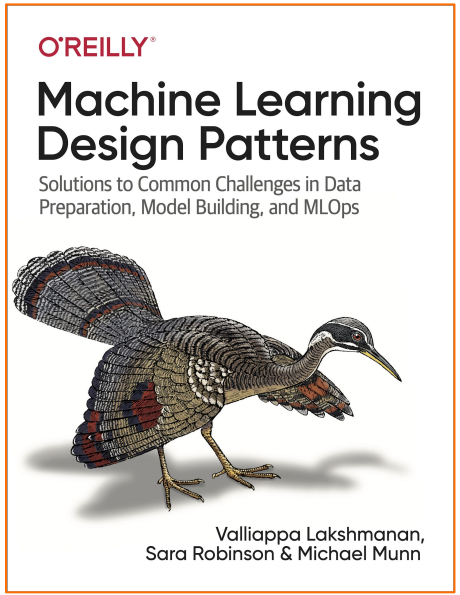

Maintainability
How do you represent categorical data when the vocabulary increases over time?

Reusability
How do you avoid having to relearn relationships between categorical variables used in related ML problems?

Beam is widely used in a few design patterns

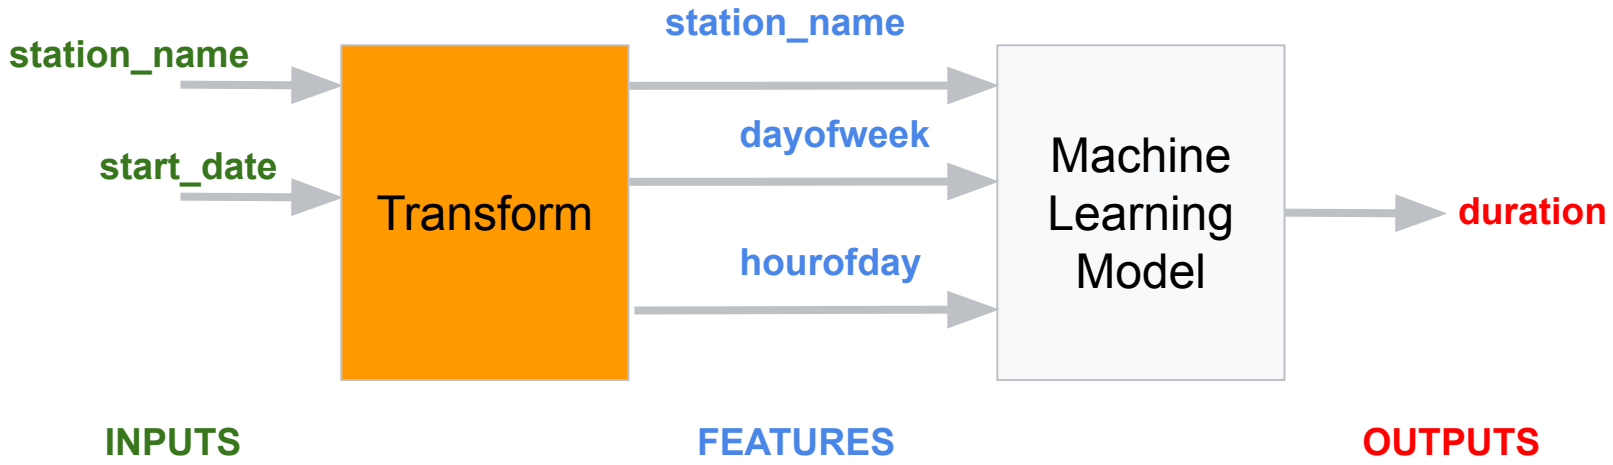

Ideally, client code does not have to know about all the transformations that were carried out

```
CREATE OR REPLACE MODEL ch09edu.bicycle_model
OPTIONS(input_label_cols=['duration'],
        model_type='linear_reg')
AS

SELECT
  duration
  , start_station_name
  , CAST(EXTRACT(dayofweek from start_date) AS STRING)
         as dayofweek
  , CAST(EXTRACT(hour from start_date) AS STRING)
         as hourofday
FROM
  `bigquery-public-data.london_bicycles.cycle_hire`
```

```
SELECT * FROM ML.PREDICT(MODEL ch09edu.bicycle_model,(
  350 AS duration
  , 'Kings Cross' AS start_station_name
  , '3' as dayofweek
  , '18' as hourofday
))
```

```
SELECT * FROM ML.PREDICT(MODEL ch09edu.bicycle_model,(
  350 AS duration
  , 'Kings Cross' AS start_station_name
  , CURRENT_TIMESTAMP() as start_date
))
```

Leading cause of training-serving skew

The Transform pattern: the model graph should include transformations

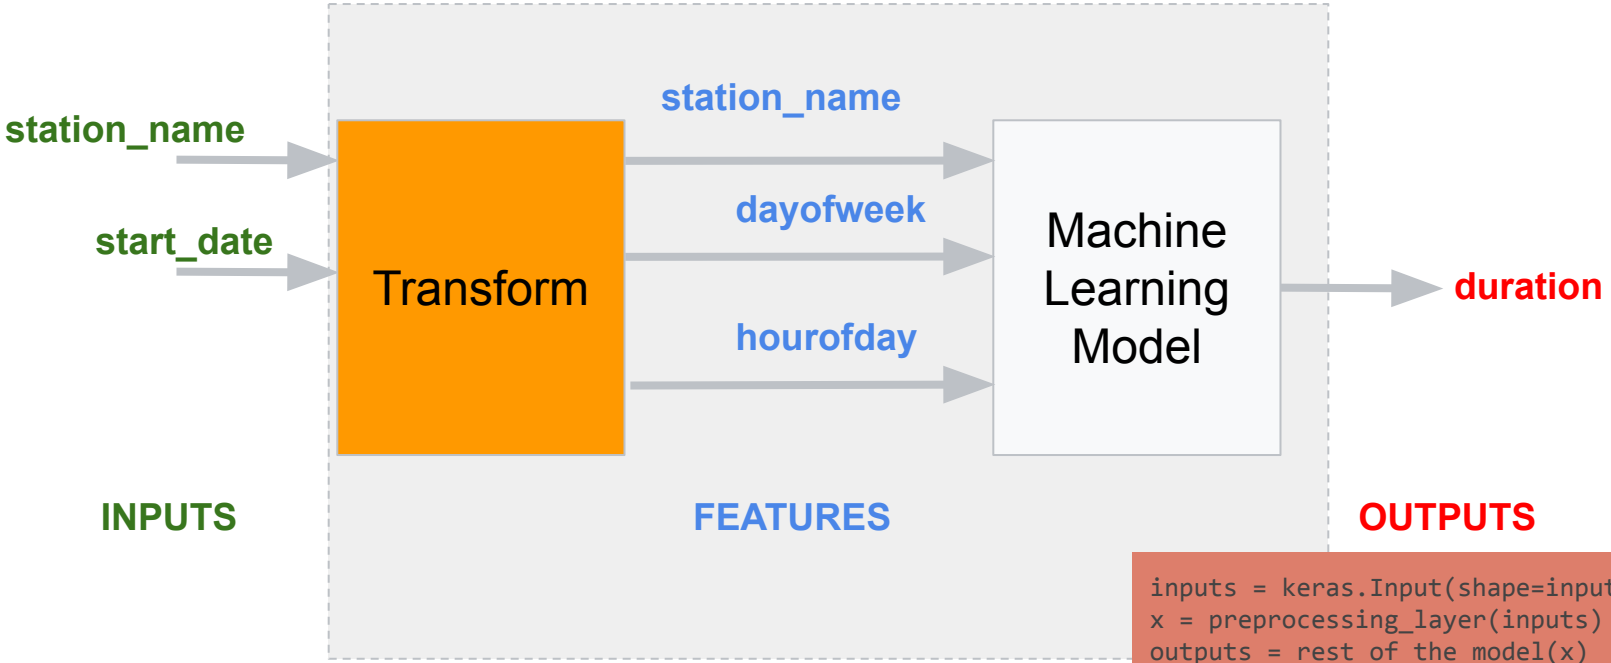


```
inputs = keras.Input(shape=input_shape)
x = preprocessing_layer(inputs)
outputs = rest_of_the_model(x)
model = keras.Model(inputs, outputs)
```

tf.transform provides reuse *and* efficiency

```
def main(output_dir):
  with tft_beam.Context(temp_dir=tempfile.mkdtemp()):
    transformed_dataset, transform_fn = (
        (raw_data, raw_data_metadata) | tft_beam.AnalyzeAndTransformDataset(
            preprocessing_fn))
  transformed_data, transformed_metadata = transformed_dataset

  # Save the transform_fn to the output_dir
  _ = (
      transform_fn
      | 'WriteTransformFn' >> tft_beam.WriteTransformFn(output_dir)

  return transformed_data, transformed_metadata
```

```
class ExportModel(tf.Module):
  def __init__(self, trained_model, input_transform):
    self.trained_model = trained_model
    self.input_transform = input_transform

  @tf.function
  def __call__(self, inputs, training=None):
    x = self.input_transform(inputs)
    return self.trained_model(x)
```

Other patterns that Beam supports well

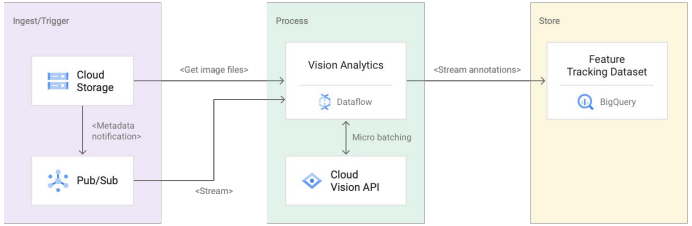

Batch Serving

Windowed Inference

```
## Time window into 2 hour windows, triggered every minute
WINDOW_INTERVAL = 2 * 60 * 60. # 2 hours, in seconds
PANE_INTERVAL = 10*60 # 10 minutes, in seconds
windowed = (data
        | 'window' >> beam.WindowInto(
                beam.window.SlidingWindows(WINDOW_INTERVAL, PANE_INTERVAL),
                accumulation_mode=beam.trigger.AccumulationMode.DISCARDING))
model_state = (windowed
        | 'model' >> beam.transforms.CombineGlobally(ModelFn()).without_defaults())
anomalies = (windowed
        | 'latest_slice' >> beam.FlatMap(is_latest_slice)
        | 'find_anomaly' >> beam.Map(is_anomaly, beam.pvalue.AsSingleton(model_state)))
```

Feature Store

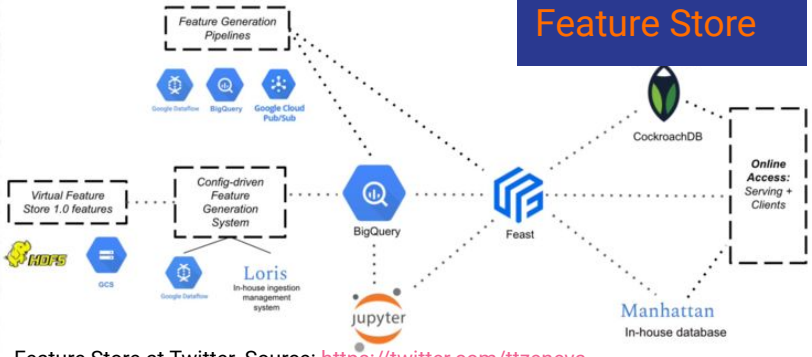

Feature Store at Twitter. Source: https://twitter.com/ttzeneva

Why?

Other patterns that Beam supports well

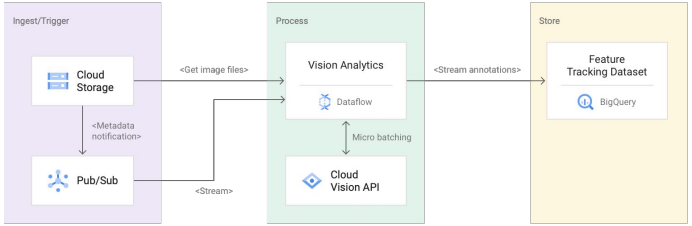

Batch Serving

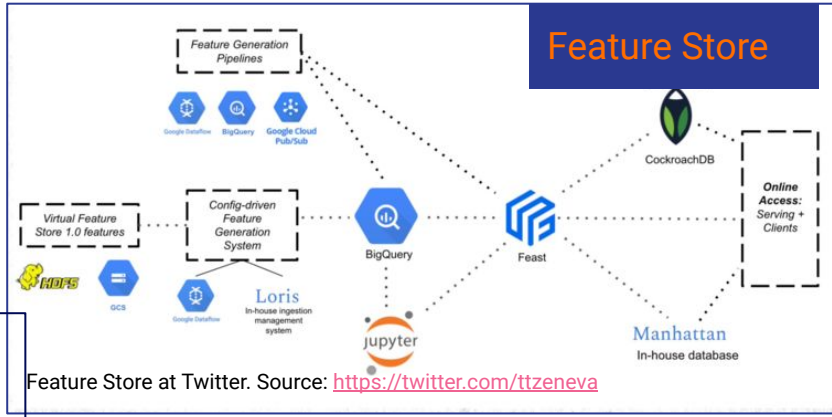

Feature Store at Twitter. Source: https://twitter.com/ttzeneva

Windowed Inference

```
## Time window into 2 hour windows, triggered every minute
WINDOW_INTERVAL = 2 * 60 * 60. # 2 hours, in seconds
PANE_INTERVAL = 10*60 # 10 minutes, in seconds
windowed = (data
        | 'window' >> beam.WindowInto(
                beam.window.SlidingWindows(WINDOW_INTERVAL, PANE_INTERVAL),
                accumulation_mode=beam.trigger.AccumulationMode.DISCARDING))
model_state = (windowed
        | 'model' >> beam.transforms.CombineGlobally(ModelFn()).without_defaults())
anomalies = (windowed
        | 'latest_slice' >> beam.FlatMap(is_latest_slice)
        | 'find_anomaly' >> beam.Map(is_anomaly, beam.pvalue.AsSingleton(model_state)))
```

- *Easy way to run in parallel*
- *Side inputs and windowing*
- *Unified Batch and Stream*
- *Much harder without Beam*

There are other patterns where Beam could be used, but isn't

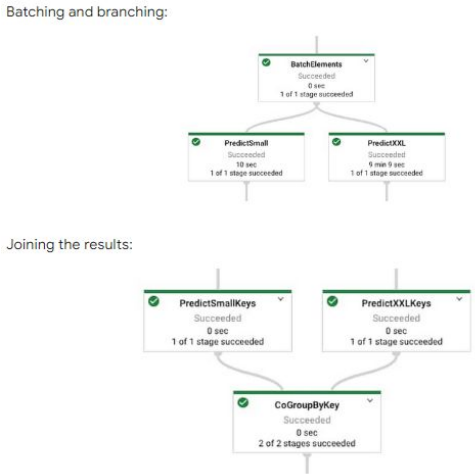

Cascade
Transfer Learning
Continuous Evaluation
Two Phase Predictions
Multimodal Input
Workflow Pipeline

https://cloud.google.com/blog/products/data-analytics/ml-inference-in-dataflow-pipelines

What's common to these?

Cascade
Transfer Learning
Continuous Evaluation
Two Phase Predictions
Multimodal Input
Workflow Pipeline

Why?

- *Training, evaluation: One-off, rare task*
- *Online serving: On-demand to millions*
- *Artifact Management among multiple ML models: Orchestration*

What if Beam could:

- scale from zero to millions of QPS
- consume/produce HTTP, cloud events
- be GPU-accelerated
- be run on-demand (start instantaneously)?

Imagine ...

A Beam Runner that runs on Cloud Run

Portable way to run Java/Go/Python across serverless container options on AWS, GCP, Azure

Scales to zero, suitable for rare ETL, scales on-demand code

Portable ML code across training, inference, evaluation

BEAM SUMMIT

Thank you!

BEAM SUMMIT

Austin, 2022

Tailoring pipelines at Spotify

By Rickard Zwahlen

BEAM SUMMIT Austin, 2022

Rickard Zwahlen

Data engineer @ Spotify

@rickardzwahlen

BEAM SUMMIT

(Mel: ABBA - Super Trouper)
Super Deduper runs as fast as lightning
Handles massive skew
Gets events to you
But only one of each, not two

Rickard Zwahlen

Data engineer @ Spotify

@rickardzwahlen

Smörgåsbord of data

BEAM SUMMIT Austin, 2022

BEAM SUMMIT Austin, 2022

How it started

How it's going

Why Scala?

- Productivity + performance
- Functional & type safe
- Large software ecosystem for data

The love triangle

Word count

```
val sc = ScioContext()
sc
  .textFile("shakespeare.txt")
  .flatMap { _ .split("[^a-zA-Z']+")
  .filter(_.nonEmpty) }
  .countByValue
  .saveAsTextFile("wordcount.txt")

sc.run()
```

Joins

```
val sc = ScioContext()
sc
   .avroFile[Artist](args("artists"))
   .keyBy(_.getArtistId)
   .hashJoin(musicLabels)
   .map { case (artistId, (artist, label)) =>
     (
       artistId,
       businessLogic(artist, label)
     )
    }

sc.run()
```

Bread & butter pipelines

A large majority of pipelines are written in Scio

Cake mix pipelines

Just add water

Data profiling

```
schedule: hourly
docker_image: grc.io/data-profiling/1.2.3@sha256:foo
docker_args:
        - wrap-luigi
        - --module
        - luigi_tasks
        - ProfileRunner
        - --input-dataset
        - Impressions.gcs
        - --partitioning
        - hours
        - --project
        - my-cloud-project
```

Data profiling (pt 2)

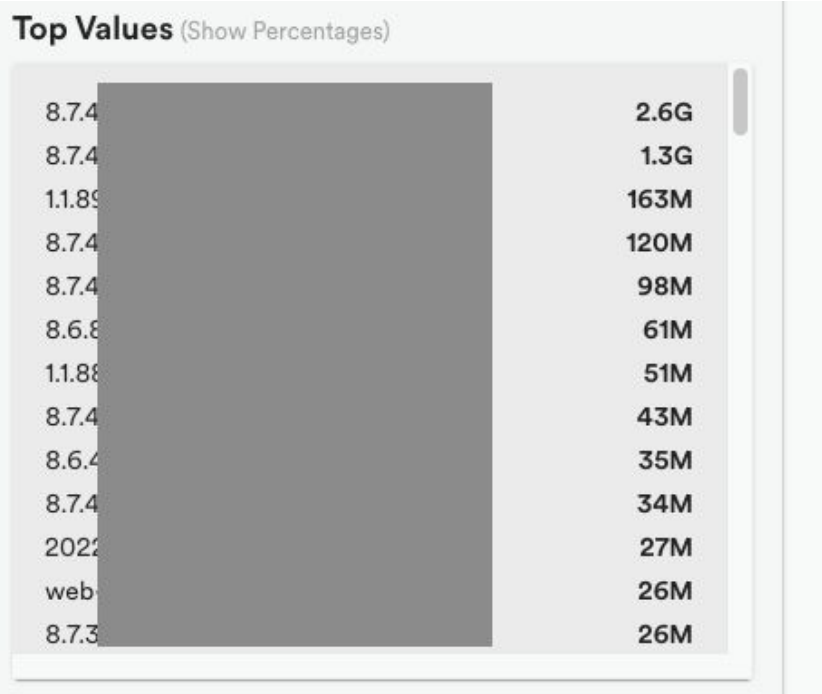

Data profiling (pt 3)

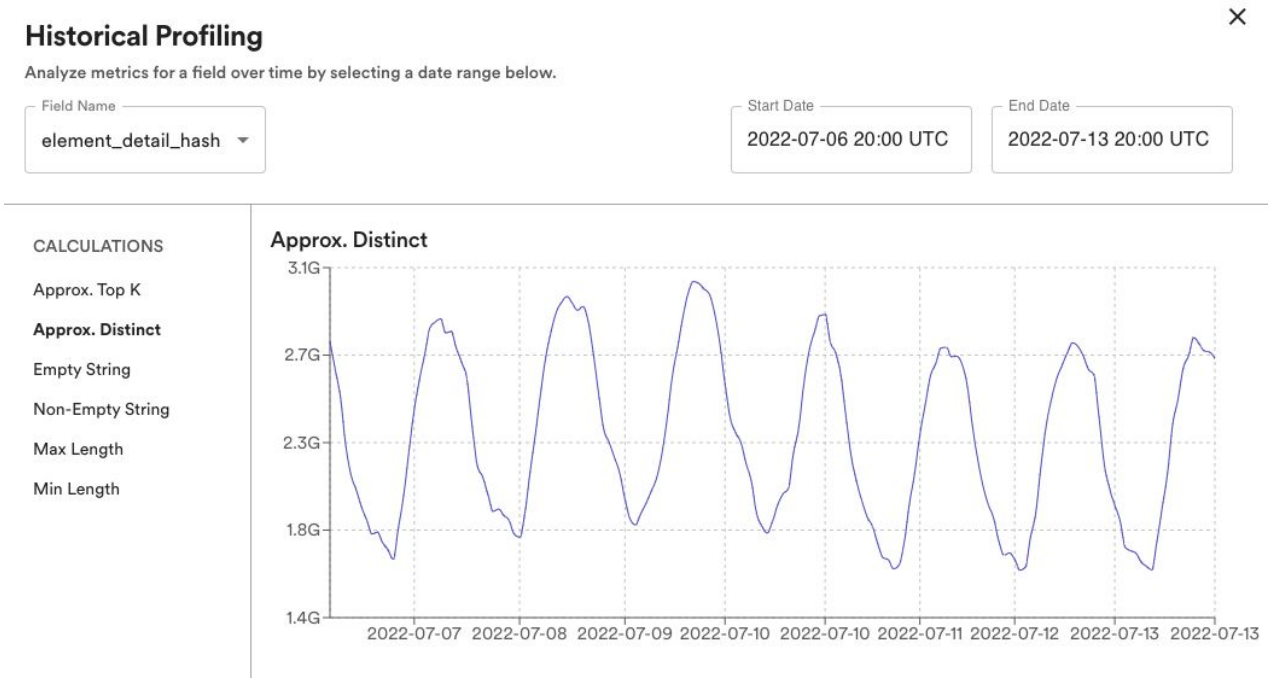

Anomaly detection

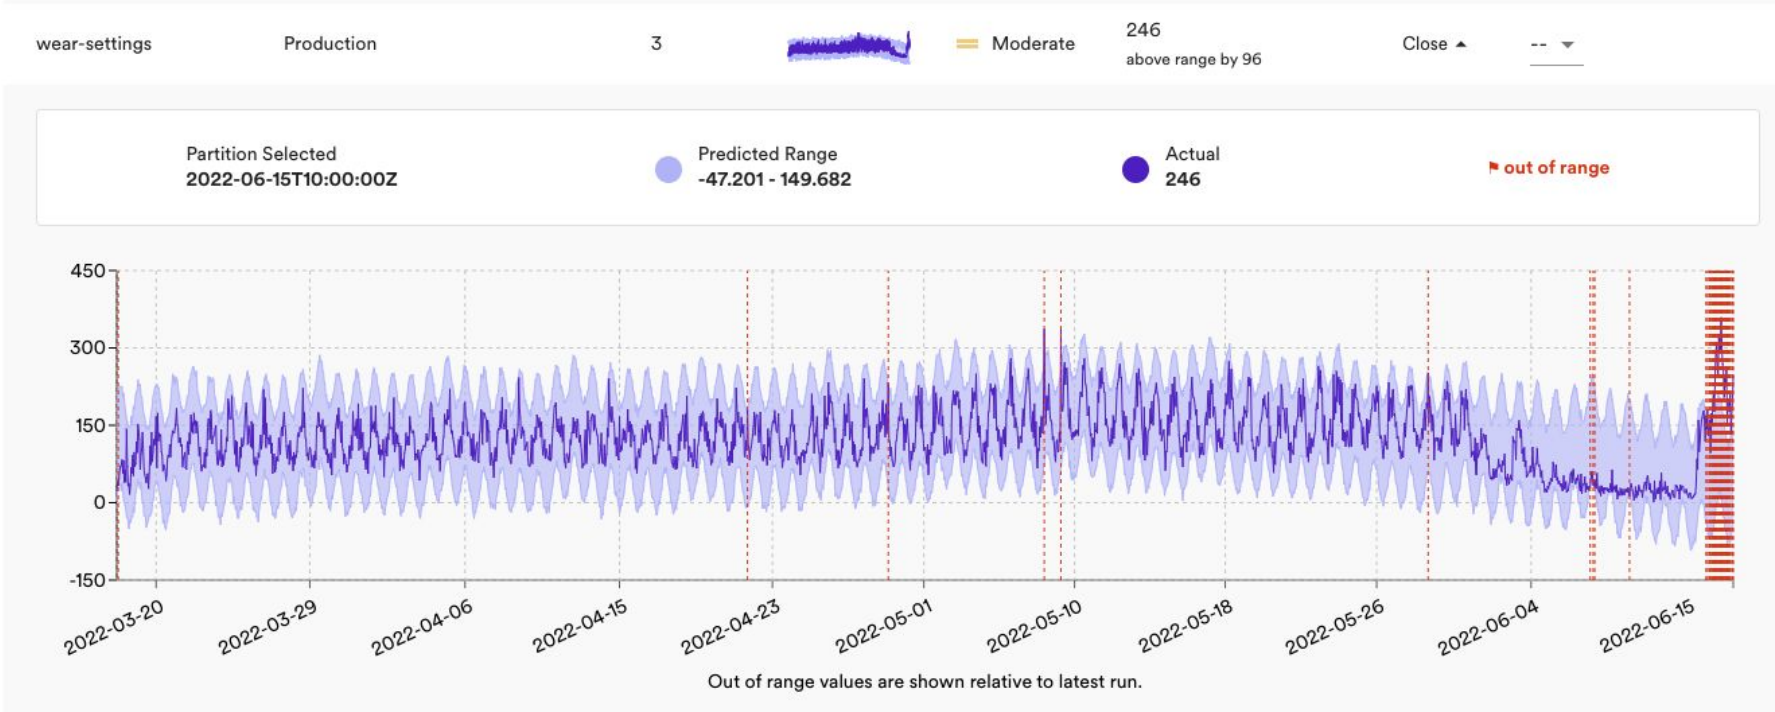

The difficult stuff

Scale, complexity, edge cases

Case by case

or

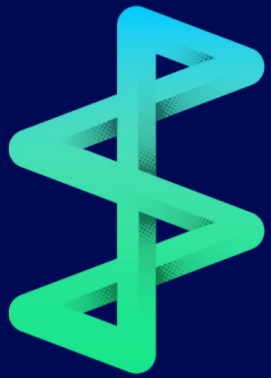

Scio

BEAM SUMMIT Austin, 2022

Spotify Wrapped

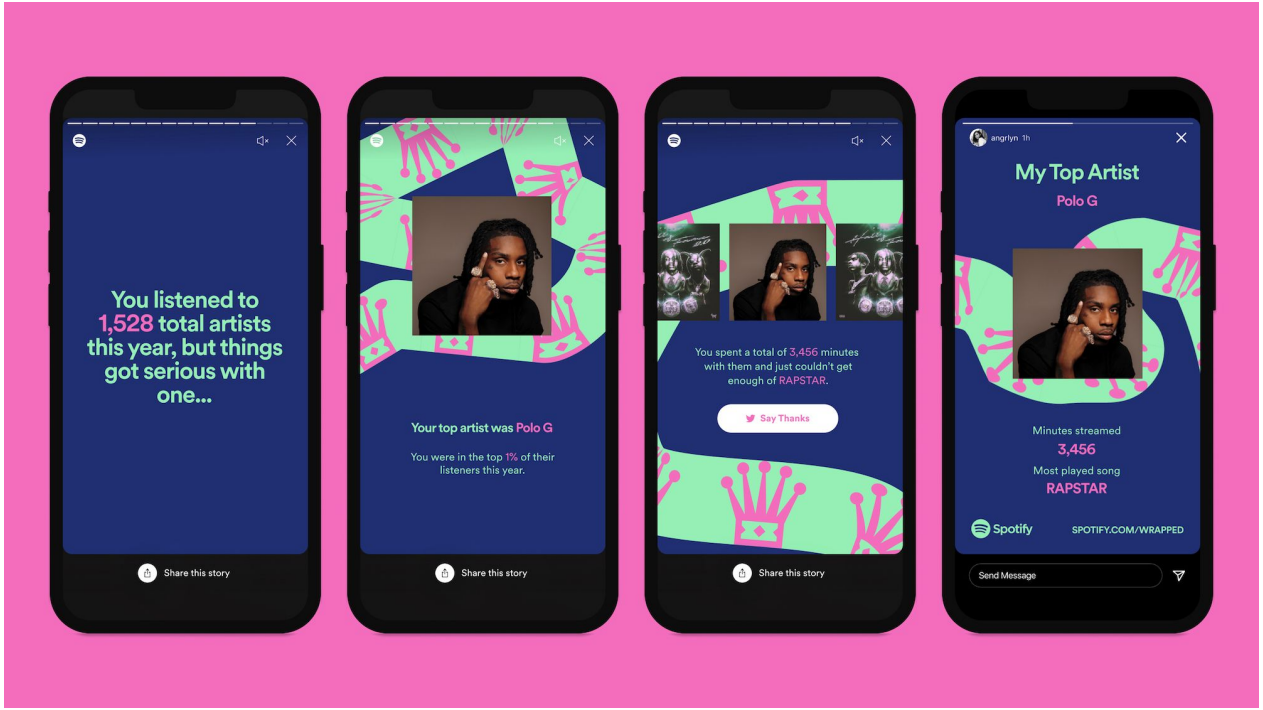

Scio + Sort Merge Bucket

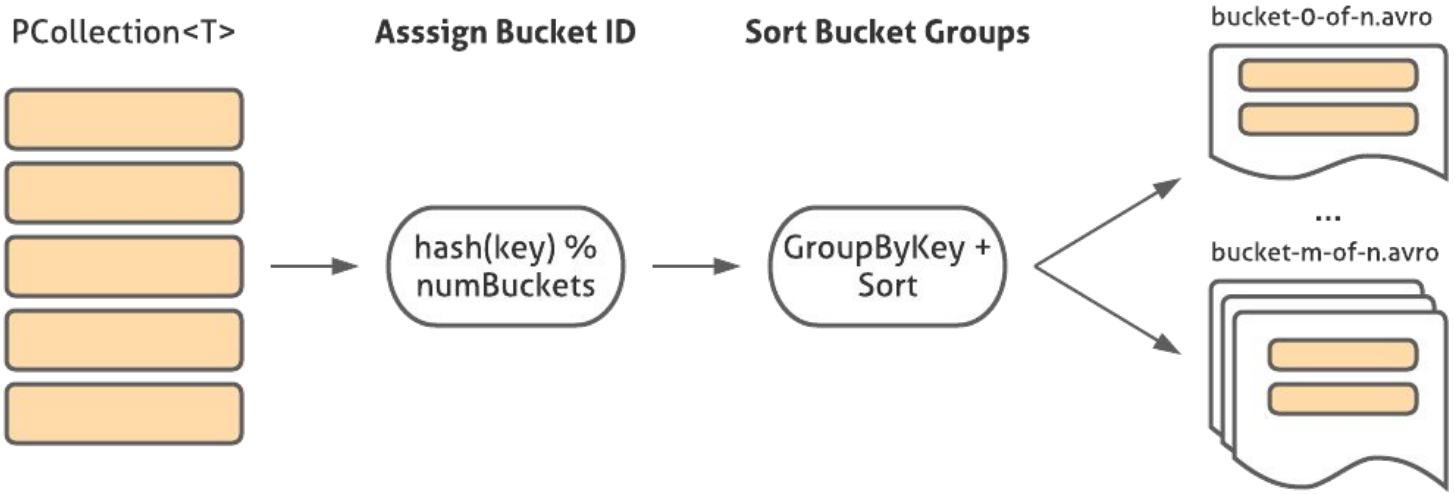

TL;DR

We use Beam at the highest level of abstraction that fits the use case

- Beam SDK
- Scio Scala API
- Plug-and-play images

Thanks!

Check out the Scio workshop on Wednesday

BEAM SUMMIT Austin, 2022

BEAM SUMMIT Austin, 2022

Beam @Twitter
Evaluation, Adoption, Migration and future.

BEAM SUMMIT Austin, 2022

Beam : Adoption, Current state and future @Twitter

Lohit VijayaRenu
@lohitvijayarenu

BEAM SUMMIT Austin, 2022

Data Processing @Twitter

Twitter Timelines
Recommendations
Analytics
Ad products
Trends, Search, Explore
Many more... or Everything

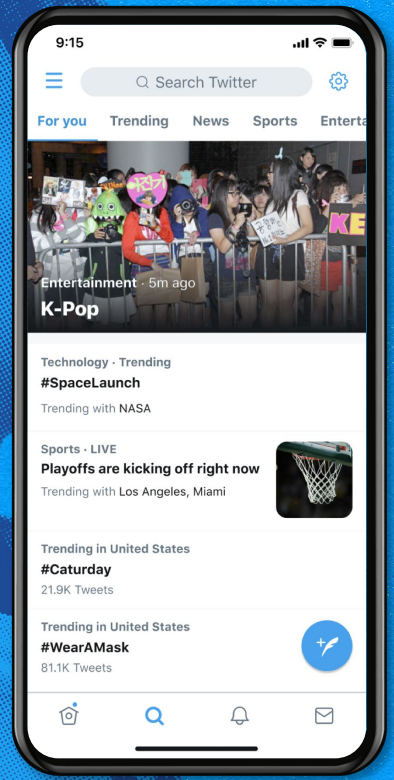

Technology

Streaming : Apache Storm, Apache Heron

Batch : Scalding, Spark, Apache Tez, Apache Hadoop

SQL : Presto, HIVE

Cloud : Google Cloud Platform (BQ, DF, GCS…)

Open Source

Twitter initiated projects : TwitterOSS

Contributions & Adoption : Apache Software, Linux Foundation, Python Foundation, Scala Center…

Every day challenges

Data pipelines	50k+
Data Volume processed	200+PB
Data across storage systems	1+EB
Events processed	7+Tri

Continuous improvement

Data Processing requirements

- Stream vs Batch (Unification)
- Modern execution framework
- Newer technologies (Spark, Tez, Flink, Beam)
- Newer API (Scio, Beam, Spark, SQL, Streaming SQL)
- Easier adoption (Metrics, configuration, debugging tools, deployment and support)

Data Processing Evaluation

Dimensions

API
Unified and modern API, API Support, Language Support, Conversion tools from existing to new API.

Platform Offering
Platform availability, support and stability. Evaluation of different runners.

Platform Integration
Integrate with other tools, SQL, tabular, Data formats, Industry adoption.

Twitter Integration
Security, Orchestration, Deployment, Workbook integration, Chargeback, Monitoring, Cost.

Model use cases
Production vs Ad hoc stream and batch processing, ML workloads, Analytics. Right tool for customer use case.

Why Beam is attractive

- **Unified API**, Modern Execution frameworks
- **Flexibility of different runners** and how it affects company strategy
- Attractive for **multi cloud support**
- Different **programming languages**
- Strong open source **community and support**

Streaming Adoption
Ad Engagement Analytic Platform

- Ad Engagement pipelines built on **lambda architecture**
- Stream processing **millions of events per second**
- Migrate **Apache Heron pipelines to Apache Beam**
- Utilize **same API** for both batch and streaming components
- Increased **developer velocity** and cleaner abstraction

Ad Engagement Analytics Platform

Batch Adoption
Experimentation Pipelines

- **Modernizing** Twitter Experimentation Pipelines
- Scalding based hard to **maintain, debug and scale**
- Easier **programming paradigm**
- Increase **developer productivity**
- Pipeline runtime from **days to hours**

Challenges

- **Language** :SCIO, Java, Python
- **Migration** : Variety and Scale
- **Custom libraries** : Use case specific logic
- **Long term support** : Compare against other APIs
- **Twitter Integration** : Metadata, deployment, monitoring…

Current Use cases

- Machine Learning & Feature Engineering pipelines
- Curated data and metrics calculation
- Data Replication and Ingestion framework
- Real Time Analytics and Monitoring
- Ad Analytics platform
- Twitter Health monitoring pipelines
- Product learning platform

Future for Beam @Twitter

- Migration of all pipelines to **Apache Beam**
- **Unifying streaming/batch** and increase streaming use cases
- Integration tooling for **data delivery, metadata and monitoring**
- **Self serve** deployment and management
- Excited about **community engagement and contributions**

More at Beam Summit

- Talk to us about **opportunities**
- Tuesday, 19 14:00
  Log ingestion and data replication at Twitter by Praveen Killamsetti & Zhenzhao Wang
- Tuesday, 19 17:15
  Apache Beam backend for open source Scalding by Navin Viswanath

Thank you!

Twitter Career
Twitter Engineer Blog
Twitter Open Source

BEAM SUMMIT Austin, 2022