



GREAT LAKES
**BUSINESS INTELLIGENCE
& BIG DATA SUMMIT**

Thursday, March 7, 2019 | Metro Detroit

Welcome!
The presentation will begin shortly



@GreatLakesBI
#GreatLakesBI19

Hosted by:



MACHINE LEARNING

HYPE OR HERO?

PRESENTATION

- Who are we?
- Who is Holland?
- Machine learning steps
- Machine learning on the Gartner hype curve
- Holland experience or story with Machine Learning
 - What business problems were we trying to solve
 - Azure machine learning for outbound
 - Azure custom vision for outbound
 - Customer churn using R
 - Deployment to SQL
 - What did we learn
 - What would I do different
 - What next
- Questions?

WHO ARE WE?

- Jeff Kwiatkowski & Joel Smith
- Director of Business Intelligence and Project Management teams for Holland
- Data Scientist
- We are a Microsoft shop
 - O365
 - Power BI
 - Azure
- Previous life -
 - Big data
 - IoT
 - Qlik, Tableau, Business Objects, Cognos

WHO IS HOLLAND

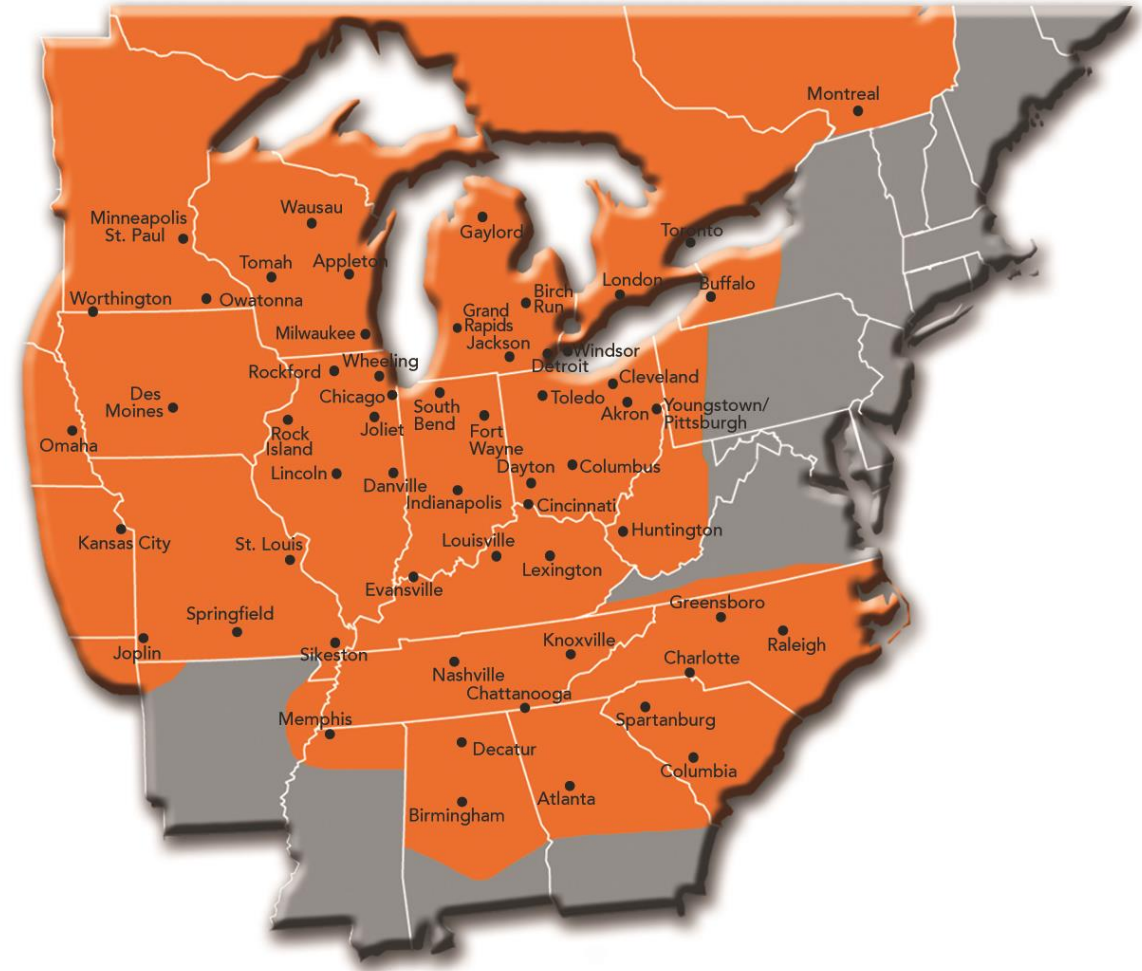
Less than load (**LTL**) **freight** is the transportation of products or goods that **does** not require a full truckload due to the smaller nature of the parcel.
... **LTL** carriers specialize in optimizing their loads; moving more goods for more shippers in an efficient manner.



WE MAKE NEXT-DAY HAPPEN.

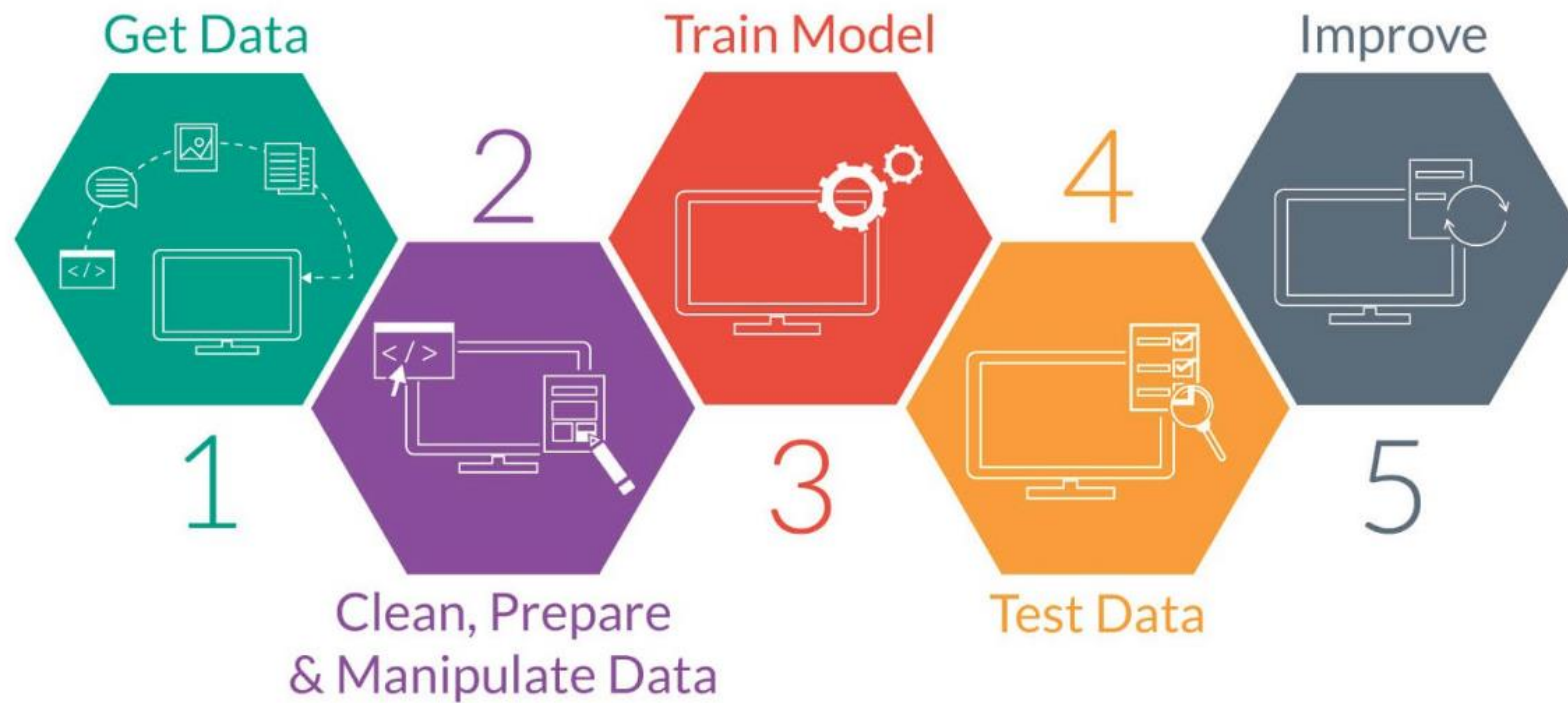
WHO IS HOLLAND

- Revenue (FY 2017) \$1.16 Billion
- Employees 8,000
- Terminals 53
- Tractors 4,200
- Trailers 7,000
- Forklifts 1,400
- Annual Miles 300 million
- Length of haul 428 miles
- Weight per shipment 1,635 lbs
- Avg. Transit Days 1.3 days

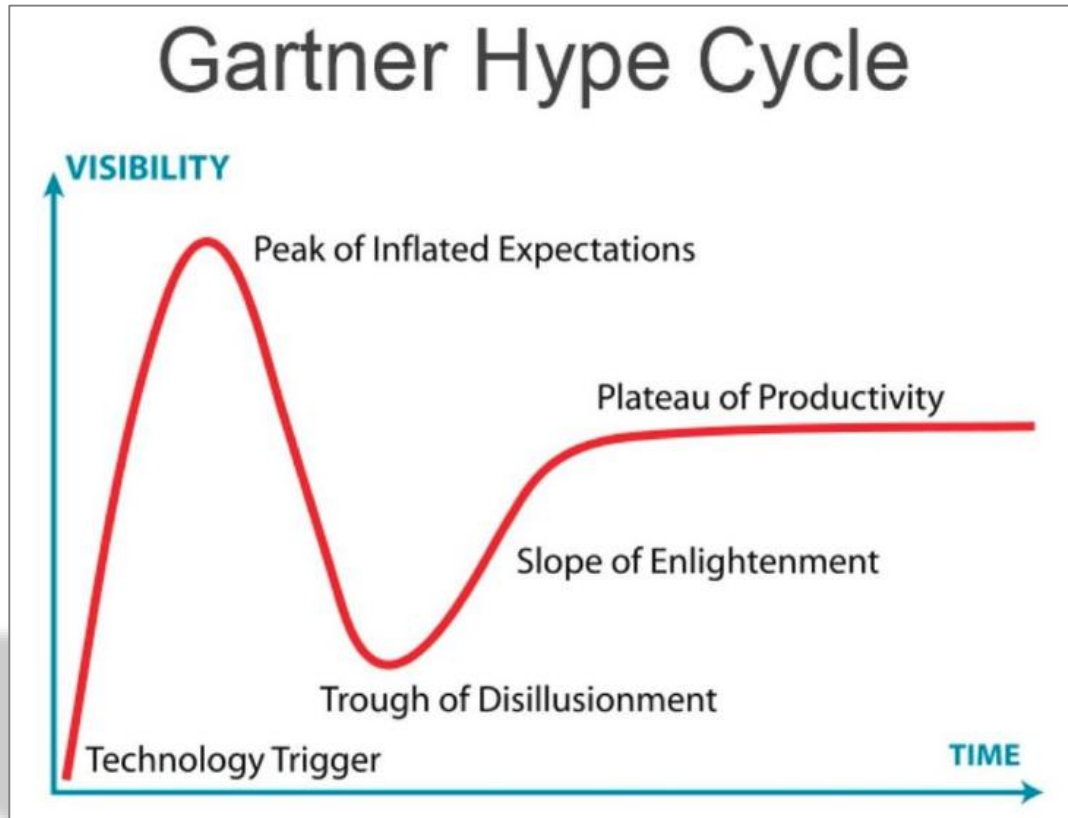


WE MAKE NEXT-DAY HAPPEN.

MACHINE LEARNING



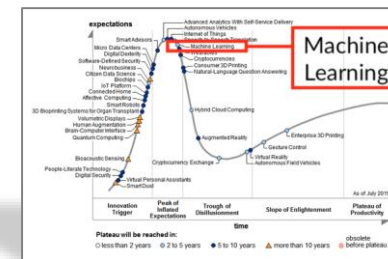
LOOKING BACK – HYPE CYCLES EMERGING TECHNOLOGIES



Machine Learning

2015 – Peak of Inflated Expectations
2 - 5 years to plateau

2017 – Peak of Inflated Expectations
2 - 5 years to plateau



Your Audience Tunes Out After 10 Minutes. Here's How To Keep Their Attention



Carmine Gallo Senior Contributor ⓘ

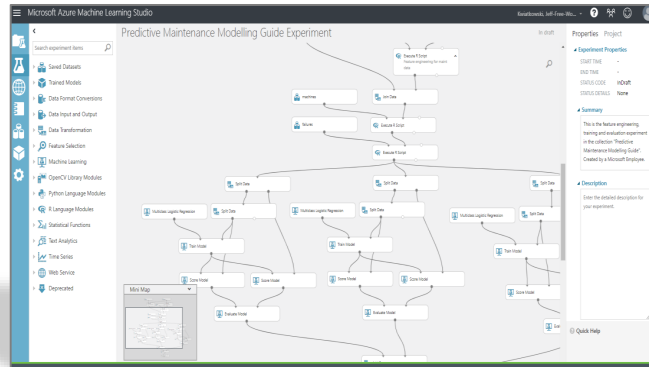
Leadership Strategy

I write about leadership communication to grow sales and build brands.


Cognitive scientists have a reasonably good idea of when audiences will stop listening to a presentation. It occurs at the 10-minute mark. According to molecular biologist, John Medina, people seem to get bored after approximately ten minutes—and it occurs in a class lecture or a business presentation.

OUR MACHINE LEARNING STORY AT HOLLAND

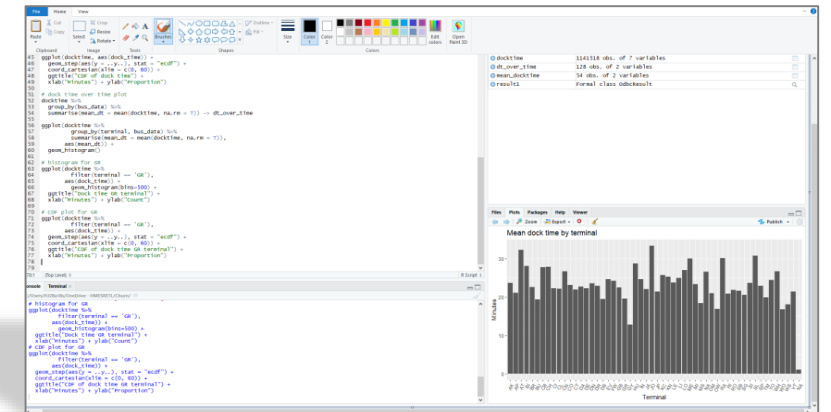
Outbound tonnage



Outbound trailer utilization



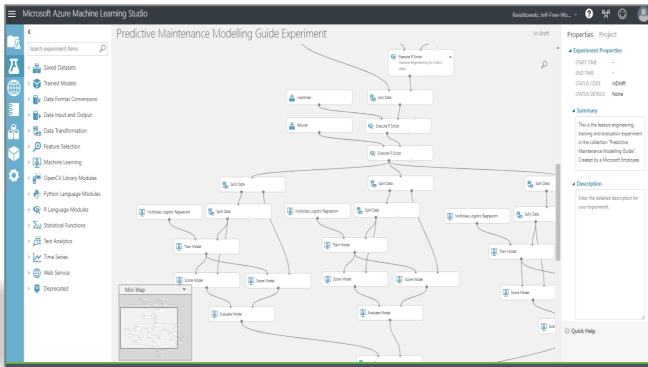
Customer churn



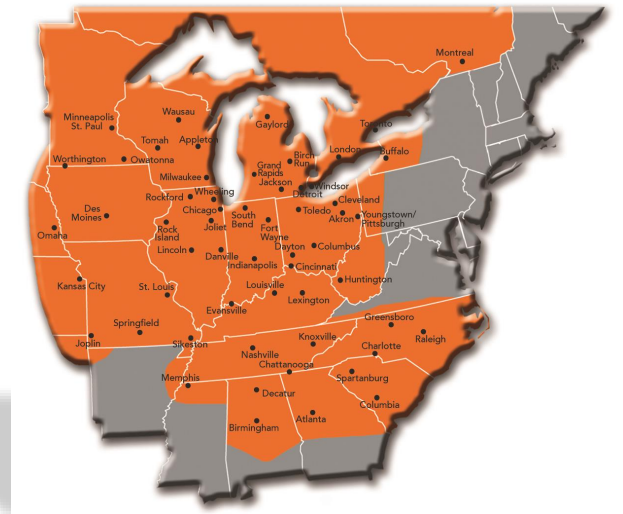
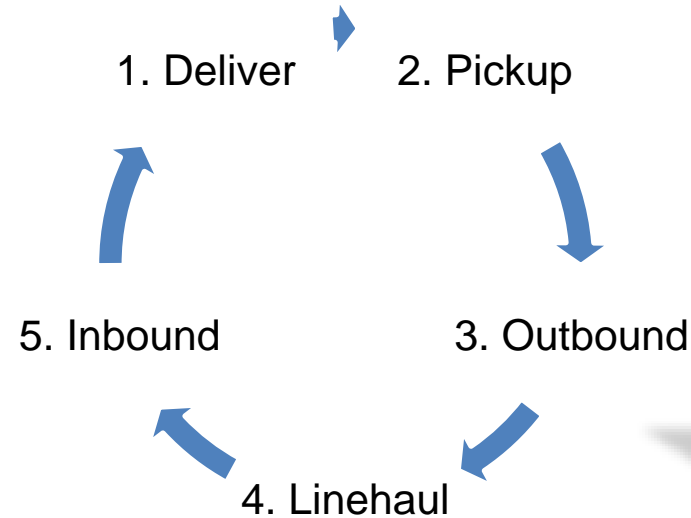
OUTBOUND TONNAGE

How much weight will ship each night from location to location?

Outbound tonnage

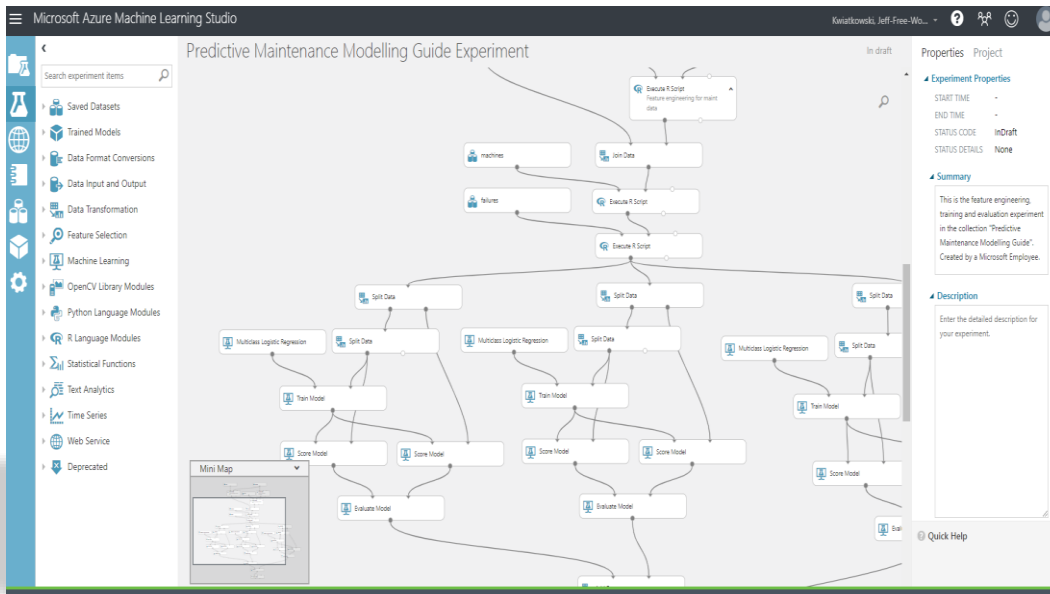


In 24 hours




SEEMS PRETTY SIMPLE

Azure machine learning studio




<https://studio.azureml.net/>



OUTBOUND TRAILER UTILIZATION


Are we filling our trailers?

Outbound trailer utilization



Upload Images

Upload your own labeled images, or use Custom Vision Service to quickly tag any unlabeled images.



Train

Use your labeled images to teach Custom Vision Service the concepts you want it to learn.



Evaluate

Use simple REST API calls to quickly tag images with your new custom computer vision model, or simply export the model to device to run real-time image understanding.



Active learning

Images evaluated through your custom vision model become part of a feedback loop you can use to keep improving your classifier.



<https://azure.microsoft.com/en-us/services/cognitive-services/custom-vision-service/>

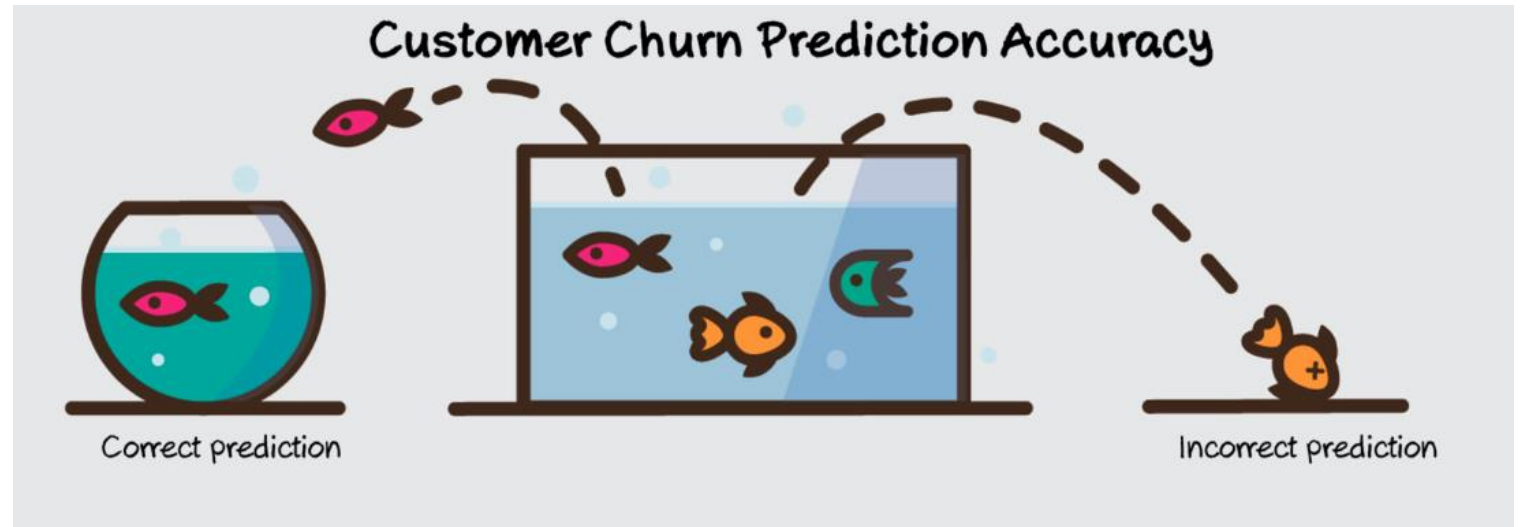
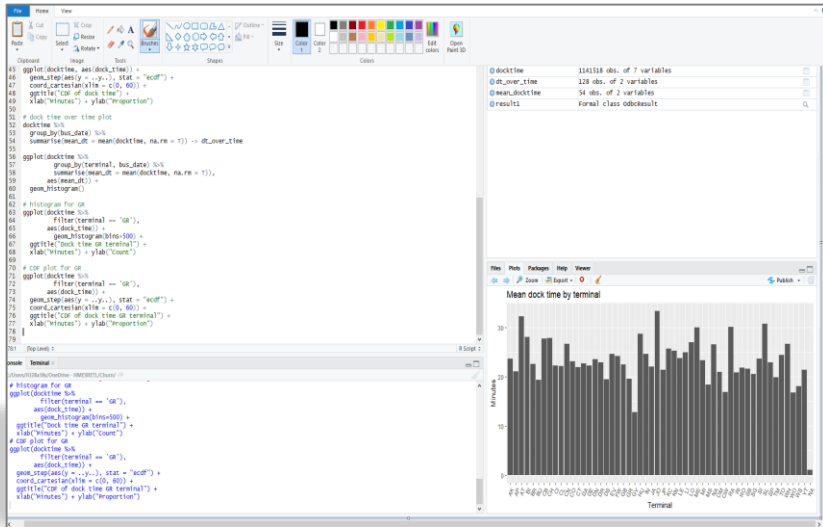


WE MAKE NEXT-DAY HAPPEN.

CUSTOMER CHURN


Are we going to lose a customer(s)

Customer Churn




TECHNOLOGY USED

- Boosted decision tree machine learning algorithm
 - Many small decision trees
 - Top-of-the-line algorithm – fast, accurate predictions
 - Reports most important variables
- Other options include:
 - Logistic regression
 - Support vector machine
 - Naïve Bayes
 - K-Nearest Neighbors




TIME TO VOTE

- One in production being used daily – knocked it out of the park
- One in semi-production – base hit, maybe a double.
- One not being used at all – Struck out

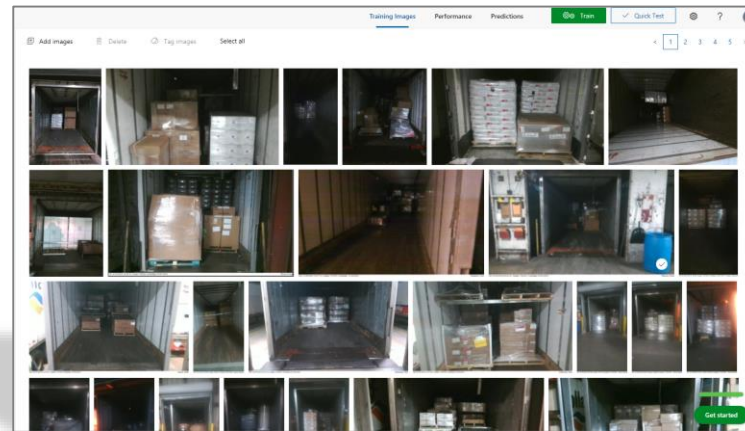


WHICH ONE WAS SUCCESSFUL?

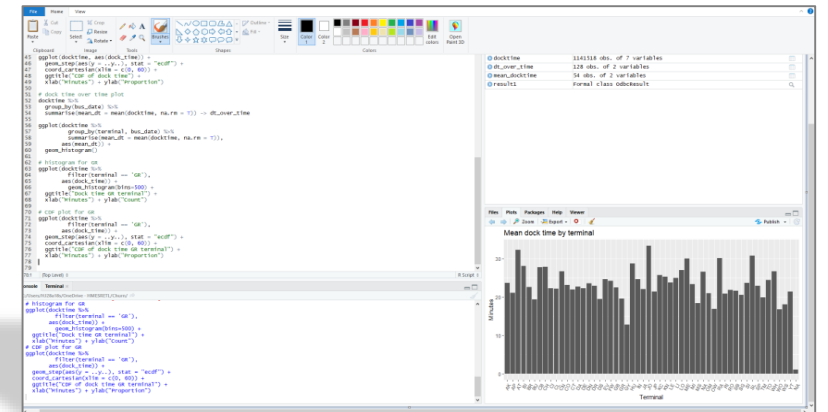
Azure machine learning studio



Azure custom vision API



R studio



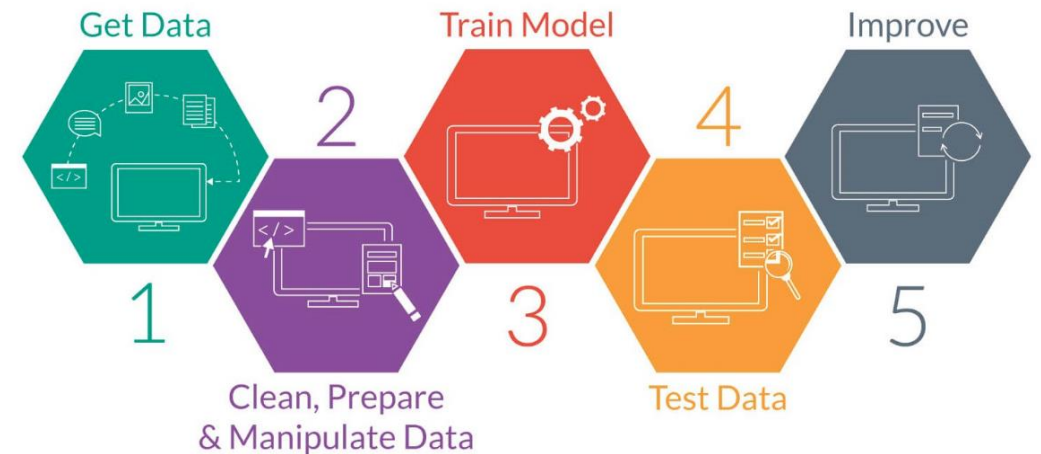
RESULTS

- One in production being used daily – Azure custom vision API (Trailer Capacity)
- One in semi-production – Customer churn – R studio (Customer Churn)
- One not being used at all – Azure ML studio (Outbound Tonnage)



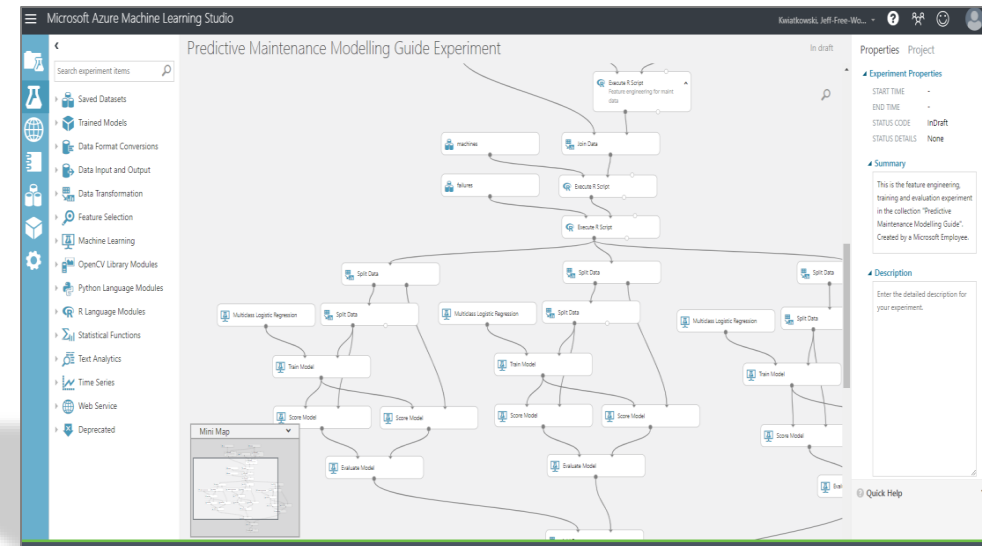
WHAT WE LEARNED

1. Find a good business partner?
2. What are you trying to predict?
3. What are you talking about?
4. Way to early, jump to early conclusions
5. Data, data, data



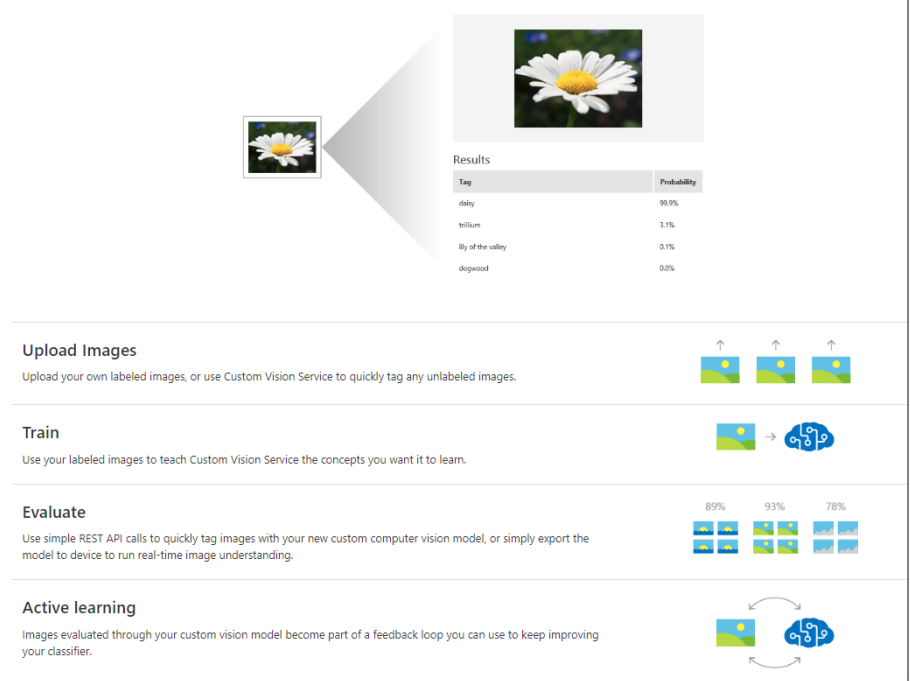
WHAT HAPPENED – OUTBOUND TONNAGE

1. Find a good business partner
 - Interest from business
 - **High level of integration**
 - **Feed data to model frequently**
2. What are you trying to predict
 - It was clear and defined
3. What are you talking about
 - What percent accurate did we need?
4. Way to early, jump to early
 - Took some feature engineering
5. Data
 - Difficult time getting the right data



WHAT HAPPENED – TRAILER CAPACITY

1. Find a good business partner
 - Engaged partner with business support
 - No change in process
2. What are you trying to predict
 - Well defined – how many feet
3. What are you talking about
 - Easy to define
4. Technology worked great!
5. Data was simple
 - Data was just a picture



The screenshot displays a custom vision model interface. At the top, a small image of a daisy is shown next to a larger, magnified version of the same image. Below the magnified image is a 'Results' table:

Tag	Probability
daisy	99.9%
trillium	3.1%
By of the valley	0.1%
dogwood	0.0%

Below the results table are four sections, each with a description and an icon:

- Upload Images:** Upload your own labeled images, or use Custom Vision Service to quickly tag any unlabeled images. (Icon: three upload arrows)
- Train:** Use your labeled images to teach Custom Vision Service the concepts you want it to learn. (Icon: image and brain)
- Evaluate:** Use simple REST API calls to quickly tag images with your new custom computer vision model, or simply export the model to device to run real-time image understanding. (Icon: three images with percentages: 89%, 93%, 78%)
- Active learning:** Images evaluated through your custom vision model become part of a feedback loop you can use to keep improving your classifier. (Icon: image and brain with a feedback loop arrow)




WHAT HAPPENED – CUSTOMER CHURN

1. Find a good business partner
2. What are you trying to predict
3. What are you talking about
4. Way to early, jump to early conclusions
5. I must have said “shape the data” a thousand times




WHAT ARE YOU PREDICTING

- Baseline – what’s a good customer
 - Predict account churn based on past history
 - Accounts that shipped at least 3 times a week for 12 consecutive weeks between the start of 2017 and halfway through 2018
 - Based on business input to focus on certain customers
- What is churn?
 - First attempt was using an excel sheet
 - Had to create “what is churn”




DATA

- Data shaping
 - More iterative than expected
- How to best translate output to end users




MOST IMPORTANT VARIABLES FROM THE MODEL

- Ratio of all shipments that are early
- Handling units
- Ratio of all shipments that are late
- Days between standard and actual service

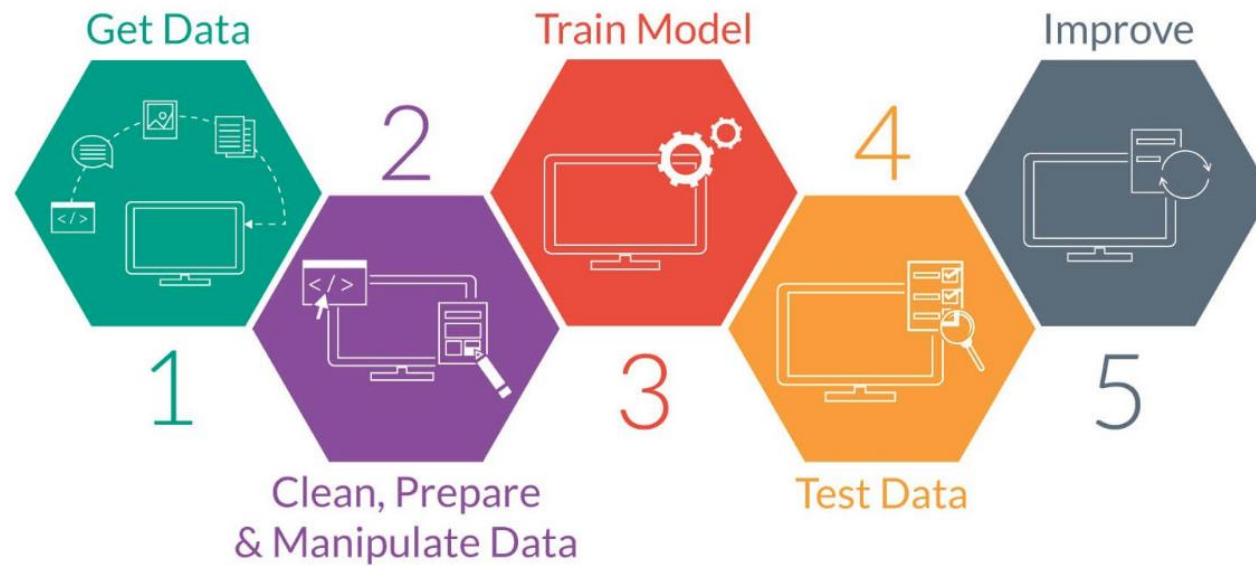


TECHNOLOGY USED

- Boosted decision tree machine learning algorithm
 - Many small decision trees
 - Top-of-the-line algorithm – fast, accurate predictions
 - Reports most important variables
- Other options include:
 - Logistic regression
 - Support vector machine
 - Naïve Bayes
 - K-Nearest Neighbors



WHAT WOULD I DO DIFFERENT?



QUESTIONS?





GREAT LAKES
**BUSINESS INTELLIGENCE
& BIG DATA SUMMIT**

Thursday, March 7, 2019 | Metro Detroit

Thank you for attending!



@GreatLakesBI
#GreatLakesBI19

Hosted by:

