Real World: Oscars

Predicting the Oscars

Poul Petersen
CIO, BigML

Where to Start?

"Let's predict customer churn!"

???

"Here are the customers we predict will leave our service"

Where to Start?

Step 1

"Let's detect fraud!

Step 2

???

Finish

"Here are the transactions we should stop immediately.

Getting Started Guide

State the problem as an ML Task

- Remember: ML finds **patterns** in **data** enabling **predictions** about future events
- This means you need **data**
    - What **data** depends on what you want to **predict**
    - And the data you have or can collect
- Data needs to have **patterns** related to what you want to **predict**
    - Not magic: still can't predict random events, lotteries, etc
- Your problem statement needs to be specific
    - Not "Let's predict churn"
    - But "Let's predict churn by looking at the profile data of all previous customers of our service who have/have not churned"
    - This can be tricky…

Where to Start?

"Let's predict the Oscars!"

???

"Here are the predicted winners"

Predicting the Oscars

BigML Scoresheet

2018

- 6 out of 6 right!
- 8 out of 8 actually, but probability of the predictions was “too low”
  - Adapted Screenplay
  - Original Screenplay

2019

- 4 our of 8 major awards correctly predicted
- Probabilities were lower this year
- This is still **significantly better** than guessing

Yay - Machine Learning!

Jennifer Riggins @jkriggins Follow

There's no doubt that as I skim the #Oscars results, @bigmlcom and @pchh nailed the prediction game! The first with #MachineLearning and the latter with great assumption making about the voters

Andy Thurai @AndyThurai Follow

What if there is a way to predict #Oscars before the awards? Wait! @bigmlcom just did that. With a 100% accuracy using #MachineLearning #AI, 3 days before. Impressive!!! #ML

KDnuggets @kdnuggets Follow

#Oscar prediction from @bigmlcom is 100% accurate - power of #MachineLearning confirmed

Mallik @S_J_Mallik Follow

Replying to @kdnuggets @efernandez @bigmlcom

If that is true why can't we predict lottery numbers? We have winning history data available online.

Contrapositive Time

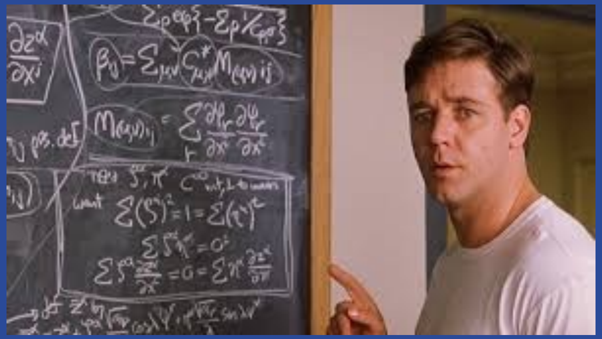

Assertion:

We can predict the Oscars ⇒ we can predict lottery numbers

Contrapositive:

We can’t predict lottery numbers ⇒ we can’t predict the Oscars

Contradiction!

Can’t Predict Lottery Numbers?

Nope Sorry! Two problems:

- The motion is chaotic (*that is extremely non-linear*)
- Even *small changes* in the initial conditions *greatly change* the outcome
- And you can’t measure the initial conditions with infinite precision (thanks Heisenberg)

Question:

Why can we predict the Oscars?

Wait, But I Read…

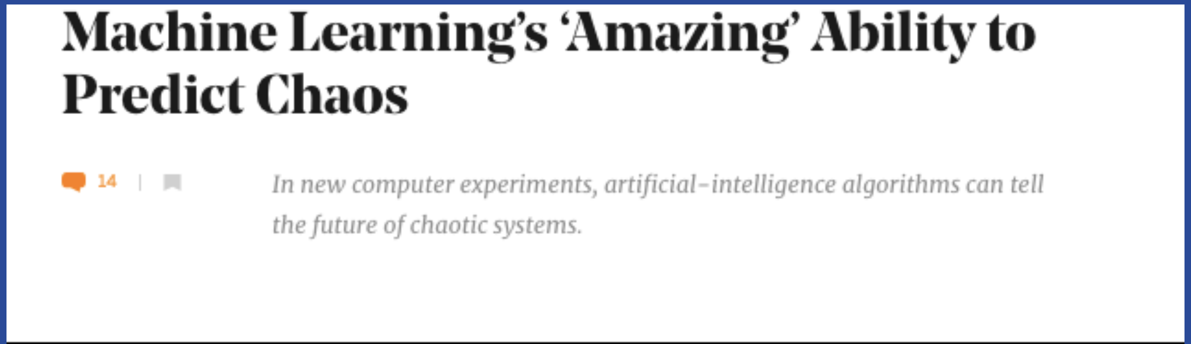
Machine Learning's 'Amazing' Ability to Predict Chaos

14

In new computer experiments, artificial-intelligence algorithms can tell the future of chaotic systems.

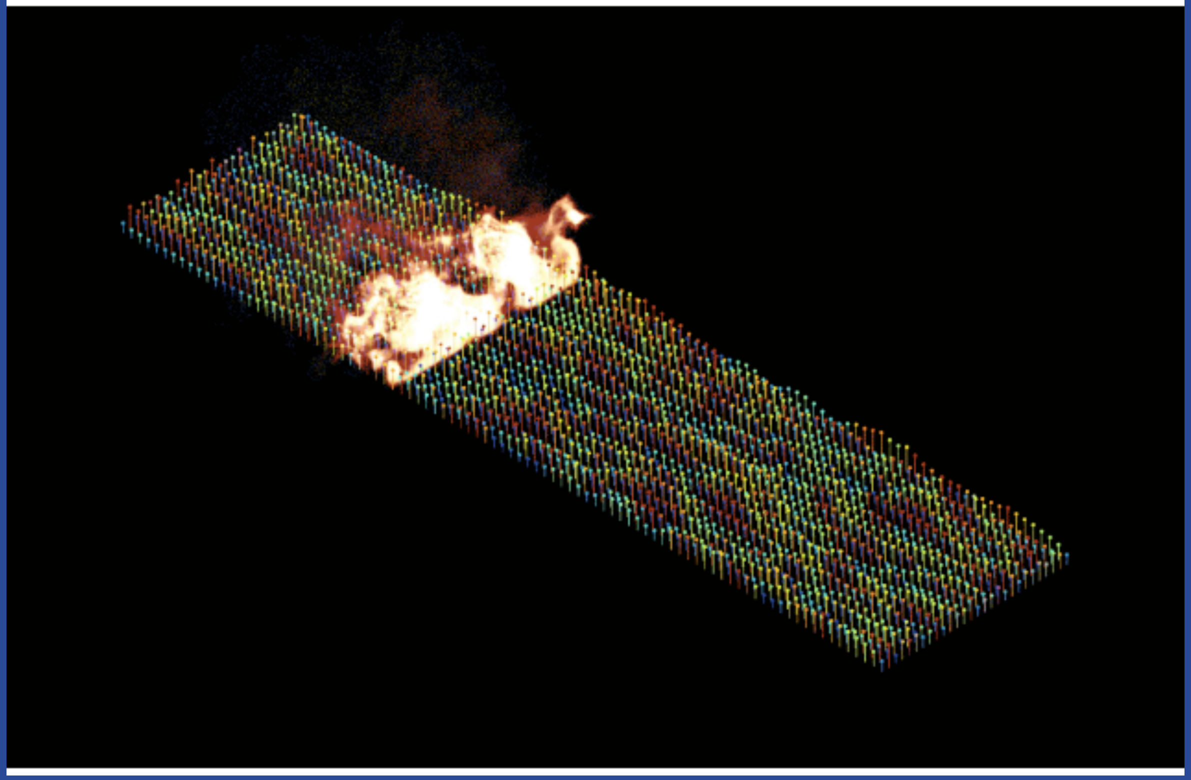

- ML to predict the propagation of a flame front
- This is also a chaotic system
- Succeeding in predicting out to 8 Lyapunov times
- Still a short amount of time
- A really short amount of time
- Lottery balls are allowed to "mix" for many, many Lyapunov times
- Does not contradict statement about predicting the lottery!

How an Oscar is Won

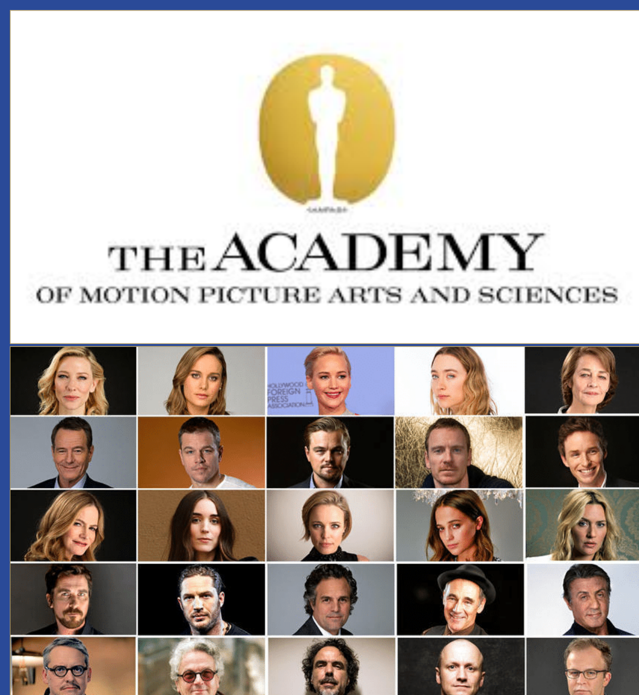

7,000+ members

voting intention?

Question:

Don’t we have the same problem as the lottery with predicting intention?

7,000 Chaotic Systems?

- Personal tastes
- Political considerations
- Values
- Cultural upbringing
- Critical education
- Pet peeves
- Corruption

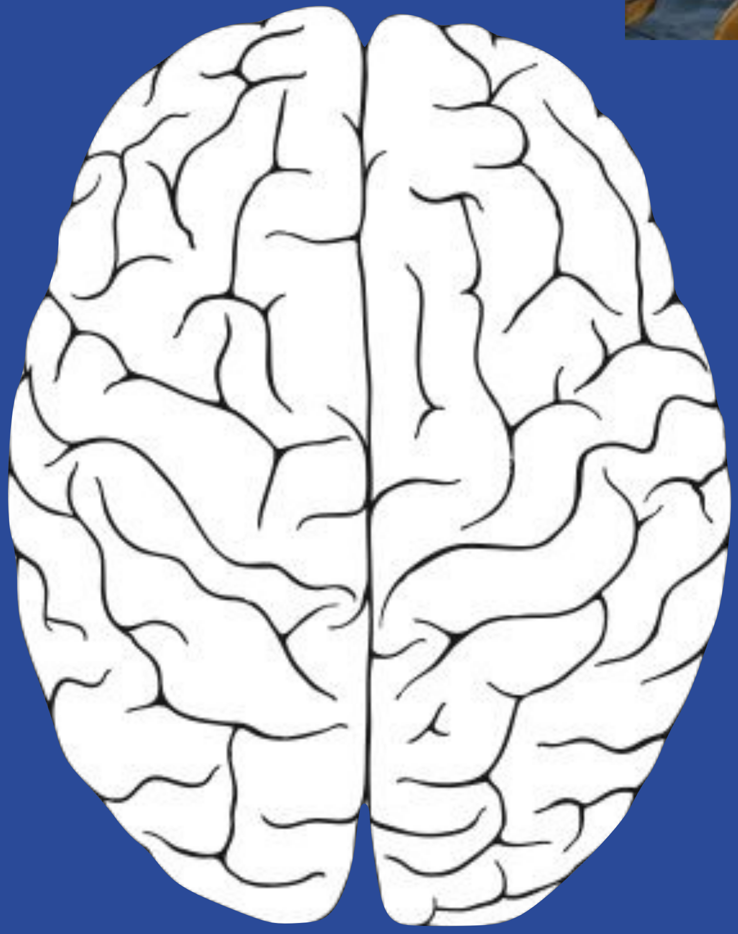

Movie Watching Robots!

- Program robots to record audio and video
- Train them to react like a human to the movie
- Run a simulation and collect votes from the robots!

Please don't try to solve everything with AI/ML

Ranking ML Applications

FEASIBILITY
(inc data availability / dec complexity)

POSTPONABLE

NO-BRAINERS
START HERE

NO-GO

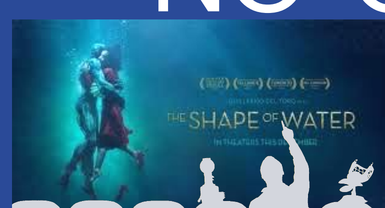

BRAINERS

ROI
(impact and cost)

What if we Just Guess?

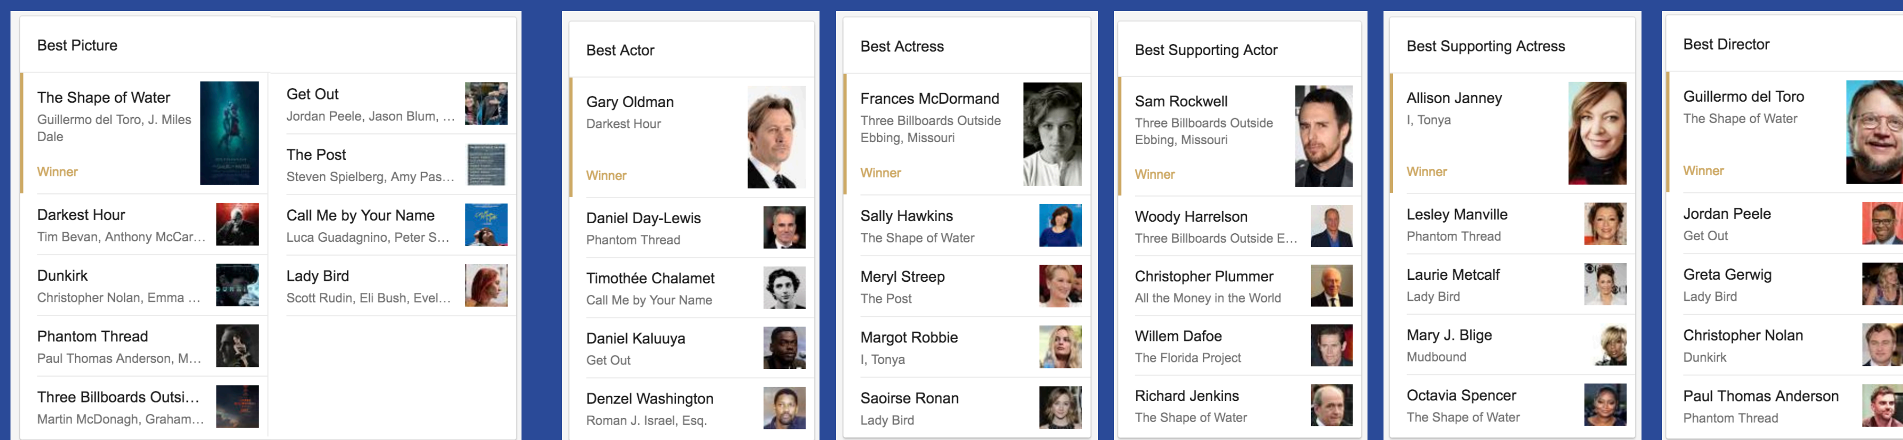

1 of 9 **1 of 5** **1 of 5** **1 of 5** **1 of 5** **1 of 5**

28,125
COMBINATIONS

0.00003556
PROBABILITY

Consider Predicting Coin Tosses

Flip 1	Flip 2	Flip 3	Flip 4	Flip 5	Flip 6
					**Prob 50%**
Tails	Heads	Tails	Tails	Tails	?Heads?

After observing the first five flips as above, what is the probability that flip 6 is Heads?

As stated, these events are independent, so the previous flips do not matter.

Consider Predicting Coin Tosses

Flip 1	Flip 2	Flip 3	Flip 4	Flip 5	Flip 6
					Prob <50%
Tails	Heads	Tails	Tails	Tails	?Heads?

What if this wasn't a fair coin?
Idea: movies are not "equally" likely to be Best Picture…

Let’s Predict Best Picture

Bafta

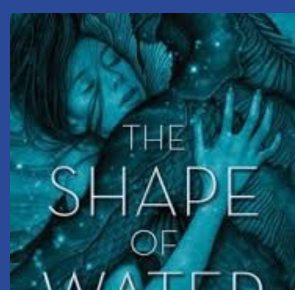

Win

Golden

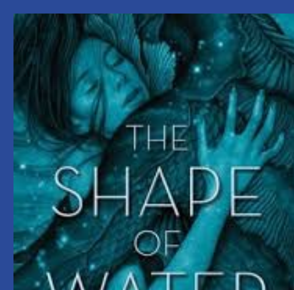

Win

Directors Guild

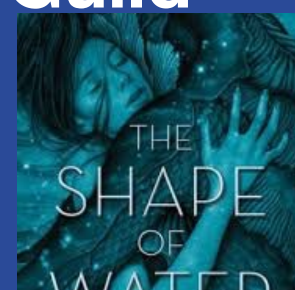

Win

Writers Guild

Lose

London Critics

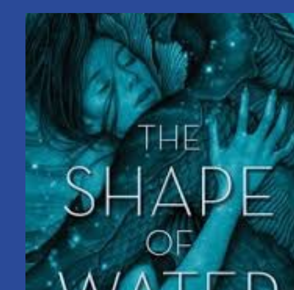

Win

Oscar

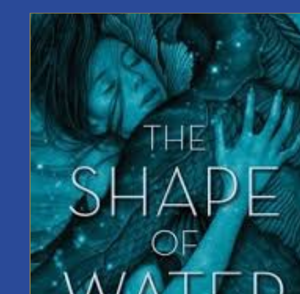

?Win?

- **These events are *not* independent**
- **Similar, but not identical, factors contribute to each win…**
- **We can expect a higher probability for Shape of Water to win**

The Features

Data pulled from IMDB…

MOVIES

- year
- movie
- movie_id
- certificate
- duration
- genre
- rate
- metascore
- synopsis
- votes
- gross
- release_date
- user_reviews
- critic_reviews
- popularity
- awards_wins
- awards_nominations
- release_date.year
- release_date.month
- release_date.day-of-month
- release_date.day-of-week

AWARDS

- Oscar_Best_Picture_nominated
- Oscar_Best_Director_nominated
- Oscar_Best_Actor_nominated
- Oscar_Best_Actress_nominated
- Oscar_Best_Supporting_Actor_nominated
- Oscar_Best_Supporting_Actress_nominated
- Oscar_Best_AdaScreen_nominated
- Oscar_Best_OriScreen_nominated
- Oscar_nominated
- Oscar_nominated_categories
- Golden_Globes_won
- Golden_Globes_won_categories
- Golden_Globes_nominated
- Golden_Globes_nominated_categories
- BAFTA_won
- BAFTA_won_categories
- BAFTA_nominated
- BAFTA_nominated_categories
- Screen_Actors_Guild_won
- Screen_Actors_Guild_won_categories
- Screen_Actors_Guild_nominated
- Screen_Actors_Guild_nominated_categories
- Critics_Choice_won
- Critics_Choice_won_categories
- Critics_Choice_nominated
- Critics_Choice_nominated_categories
- Directors_Guild_won
- Directors_Guild_won_categories
- Directors_Guild_nominated
- Directors_Guild_nominated_categories
- Producers_Guild_won
- Producers_Guild_won_categories
- Producers_Guild_nominated
- Producers_Guild_nominated_categories
- Art_Directors_Guild_won
- Art_Directors_Guild_won_categories
- Art_Directors_Guild_nominated
- Art_Directors_Guild_nominated_categories
- Writers_Guild_won
- Writers_Guild_won_categories
- Writers_Guild_nominated
- Writers_Guild_nominated_categories
- Costume_Designers_Guild_won
- Costume_Designers_Guild_won_categories
- Costume_Designers_Guild_nominated
- Costume_Designers_Guild_nominated_categories
- Online_Film_Television_Association_won
- Online_Film_Television_Association_won_categories
- Online_Film_Television_Association_nominated
- Online_Film_Television_Association_nominated_categories
- Online_Film_Critics_Society_won
- Online_Film_Critics_Society_won_categories
- Online_Film_Critics_Society_nominated
- Online_Film_Critics_Society_nominated_categories
- People_Choice_won
- People_Choice_won_categories
- People_Choice_nominated
- People_Choice_nominated_categories
- London_Critics_Circle_Film_won
- London_Critics_Circle_Film_won_categories
- London_Critics_Circle_Film_nominated
- London_Critics_Circle_Film_nominated_categories
- American_Cinema_Editors_won
- American_Cinema_Editors_won_categories
- American_Cinema_Editors_nominated
- American_Cinema_Editors_nominated_categories
- Hollywood_Film_won
- Hollywood_Film_won_categories
- Hollywood_Film_nominated
- Hollywood_Film_nominated_categories
- Austin_Film_Critics_Association_won
- Austin_Film_Critics_Association_won_categories
- Austin_Film_Critics_Association_nominated
- Austin_Film_Critics_Association_nominated_categories
- Denver_Film_Critics_Society_won
- Denver_Film_Critics_Society_won_categories
- Denver_Film_Critics_Society_nominated
- Denver_Film_Critics_Society_nominated_categories
- Boston_Society_of_Film_Critics_won
- Boston_Society_of_Film_Critics_won_categories
- Boston_Society_of_Film_Critics_nominated
- Boston_Society_of_Film_Critics_nominated_categories
- New_York_Film_Critics_Circle_won

OBJECTIVE

- Oscar_Best_Picture_won
- Oscar_Best_Director_won
- Oscar_Best_Actor_won
- Oscar_Best_Actress_won
- Oscar_Best_Supporting_Actor_won
- Oscar_Best_Supporting_Actress_won

Engineered Features:

Award items field
Nomination Counts
Awards Counts

Oscars Dataset

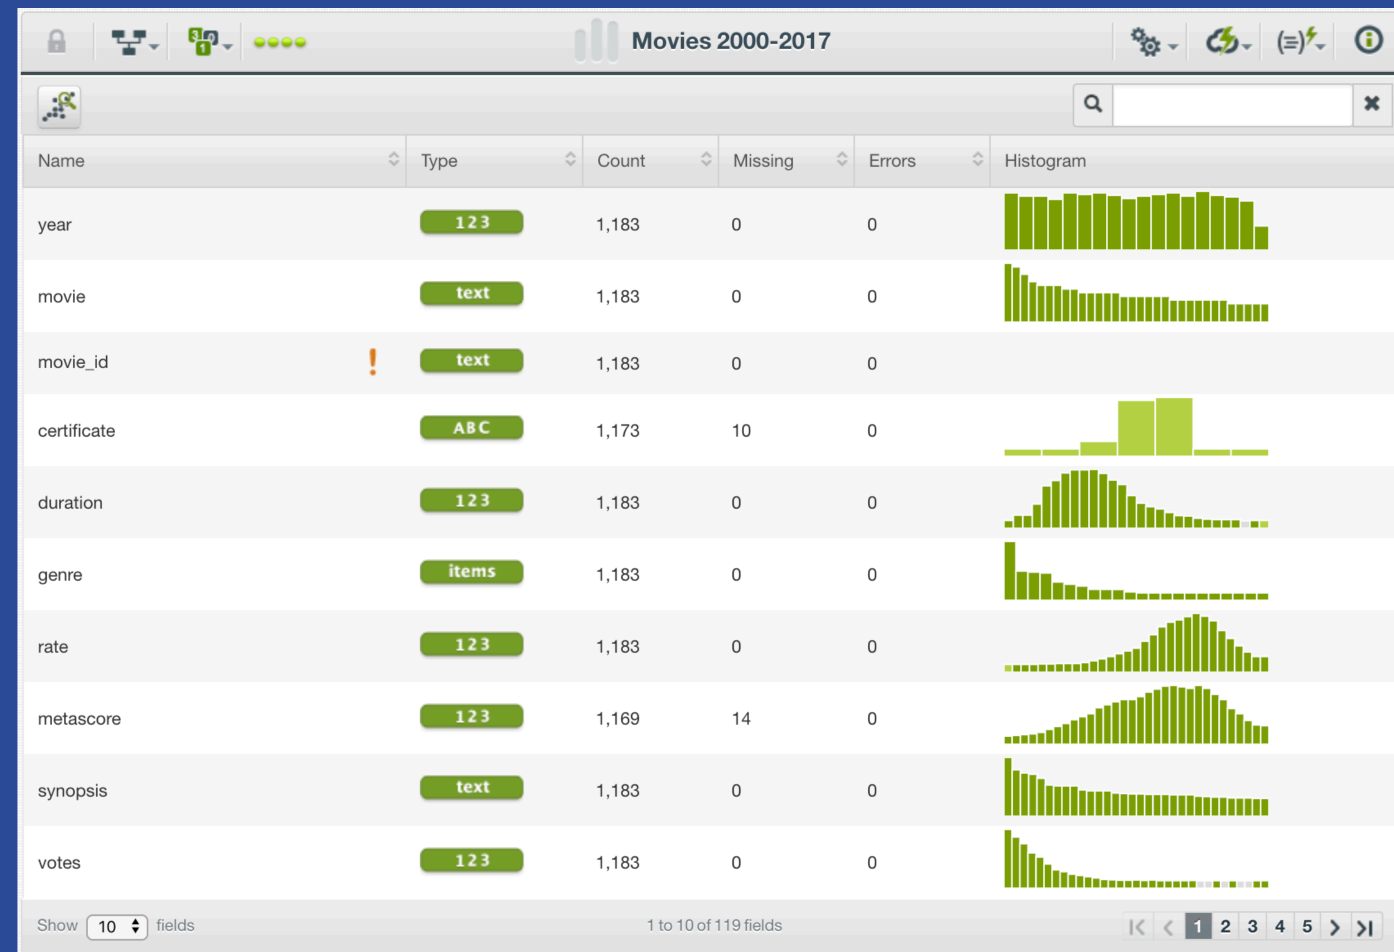

DATASET is publicly available:

https://bigml.com/user/academy_awards/gallery/dataset/5a94302592fb565ed400103b

Oscars Example

Tidbits and Lessons Learned….

- When specifying the problem, be as specific as possible
  - **Not:** "Let's predict the Oscars"
  - **Instead:** "Let's Predict the Oscars by correlating a series of award wins with the final Oscar win."
- The statement of the problem will guide the data required
- Be aware of the cost of collecting the data versus the ROI:

Effort of a ML Application

Task	Effort
State the problem as an ML task	**~10% effort**
Data wrangling	
Data transformations	**~80% effort**
Feature engineering	
Modeling and Evaluations	
Predictions	**~5% effort**
Measure Results	**~5% effort**

This is an area where **bigml**® is currently innovating

This is only such low effort **because** of platforms like **bigml**®

Oscars Example

Tidbits and Lessons Learned….

- When specifying the problem, be as specific as possible
  - **Not:** “Let’s predict the Oscars”
  - **Instead:** “Let’s Predict the Oscars by correlating a series of award wins with the final Oscar win.”
- The statement of the problem will guide the data required
- Be aware of the cost of collecting the data versus the ROI:
  - IMDB data is readily available
  - Start small and go straight to the desired result
- We’re done right?
  - Nope. You can’t escape **Feature Engineering**
  - Items: BAFTA_won_categories = list of nominations
  - Aggregations: Nomination and Award counts
- You can’t escape **Feature Selection**
  - Full user reviews costly to collect and not useful

Wait: ***How were you confident in the predictions?***

Evaluating the Models

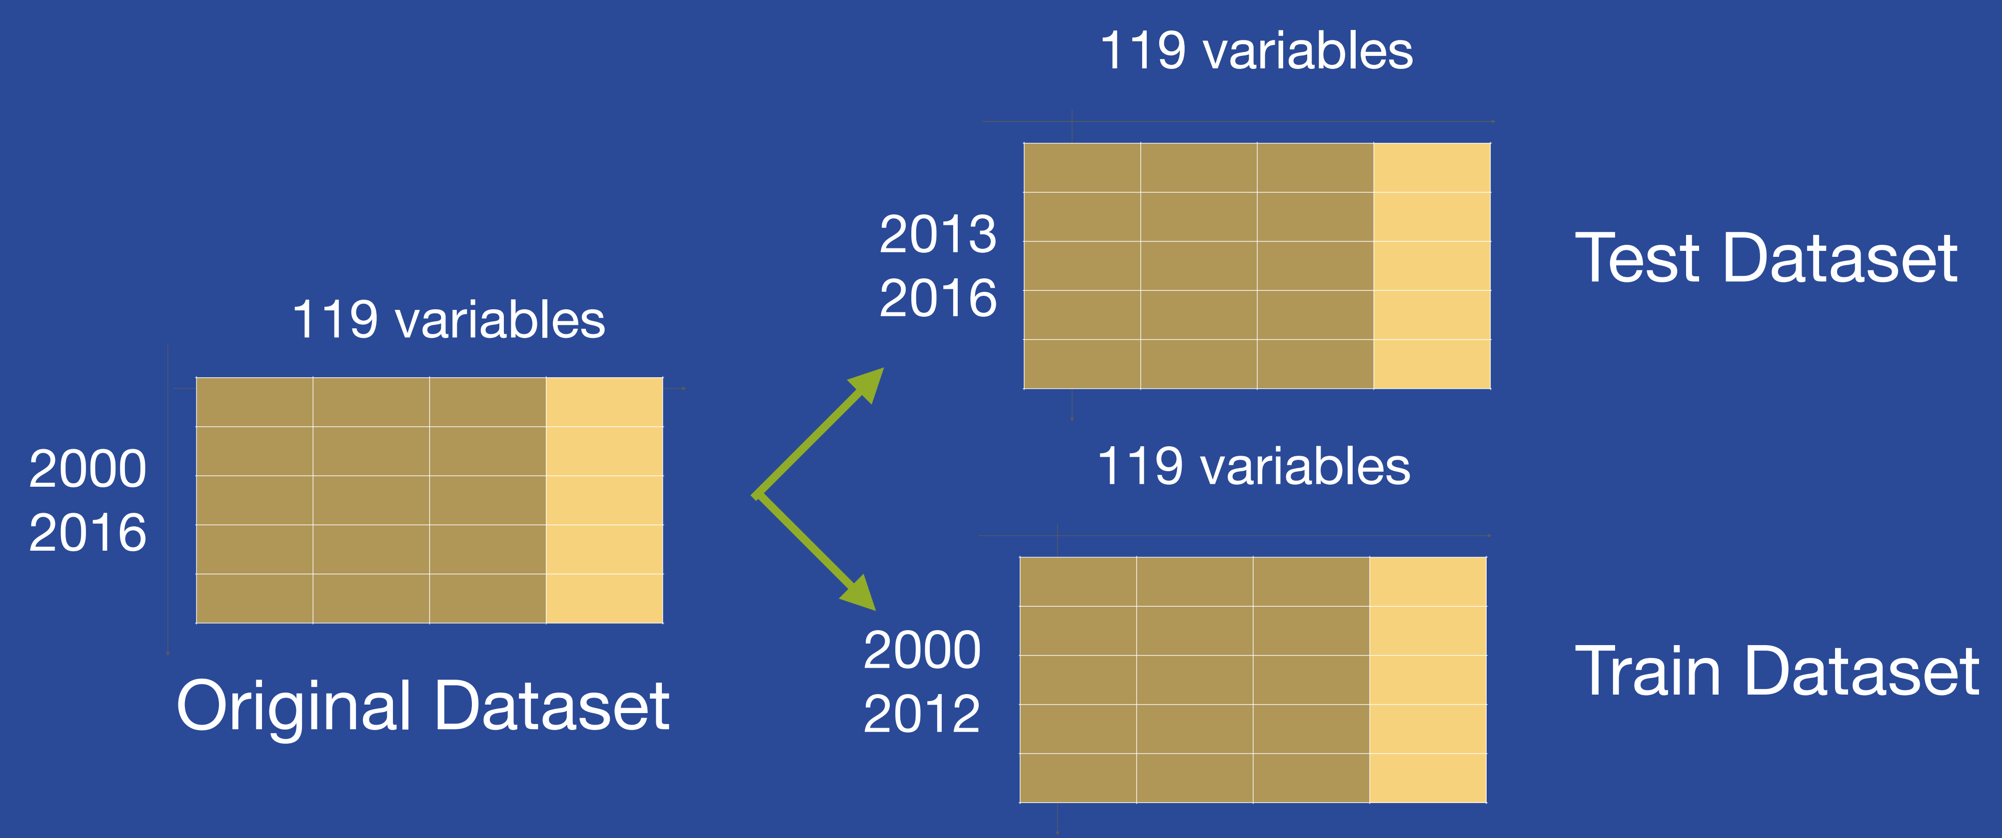

- Ultimately, we want to use all the history to predict the winner for the current year
- In order to evaluate success, we use a model built from 2000-2012 data to predict the winners for 2013-2016
- Built a separate Deepnet for each award category
- Evaluation obtained a ROC AUC over 0.98 across all award categories

Reality Check

Three Important Concepts in Applying ML…

- All Machine Learned models are **wrong**
- Real-world Machine Learning is **iterative**
- End-to-end Machine Learning is **compositional**

Tenets of Machine Learning

- All Machine Learned models are **wrong**, but some are useful
  - Better features always beat better algorithms
  - Good algorithms already exist and are good enough
  - Tools like OptiML exist which can help optimize performance
  - The data is **never** good enough
- Real-world Machine Learning is **iterative**
  - Automation is better than hand tuning - you need an **API**!
  - When data changes quickly, training speed is more important than accuracy
  - **Repeatability** is superior to a single strong result
- End-to-end Machine Learning is **compositional**
  - Problems are solved with workflows of algorithms
  - A ML solution is not real until it is in production
  - ML is here: Now we need 100,000x people applying ML

Your Turn!

- What are some problems you can solve with ML?
- Do you have the data
  - Where is it? Can you get it?
  - Does it need cleaning (hint: yes)
  - What ML tasks will be involved?
- Remember: go straight to the result
  - Prove it before you build it
  - Use Models and Logistic Regressions to start
  - Spend time on features and introspection

bigml®