“Housekeeping”

Twitter: #ACMLearning

- Welcome to today’s ACM TechTalk, “**Agile Data Science: Achieving Salesforce-Scale Machine Learning in Production**.” The presentation starts at the top of the hour and lasts 60 minutes. Audio and video will automatically play throughout the event. On the bottom panel you’ll find a number of widgets, including Twitter and Sharing apps.

- If you are experiencing any problems/issues, **refresh** your console by pressing the **F5** key on your keyboard in Windows, **Command + R** if on a Mac, or refresh your browser if you’re on a mobile device; or close and re-launch the presentation. You can also view the Webcast Help Guide, by clicking on the “Help” widget in the bottom dock.

- To control volume, adjust the master volume on your computer. If the volume is still too low, use headphones.

- At the end of the presentation, you’ll see a survey open on your screen. Please take a minute to fill it out to help us improve your next webinar experience. You may also open the survey at any time throughout the presentation from the resources window.

- This session is being recorded and will be archived for on-demand viewing in a few days. You will receive an automatic email notification when it is available. See http://learning.acm.org/ for updates. And check out https://learning.acm.org/techtalks for archived recordings of past webcasts.

Agile Data Science
Achieving Salesforce-Scale Machine Learning in Production

Sarah Aerni, PhD
Director of Data Science, Einstein Platform

saerni@salesforce.com
@itweetsarah

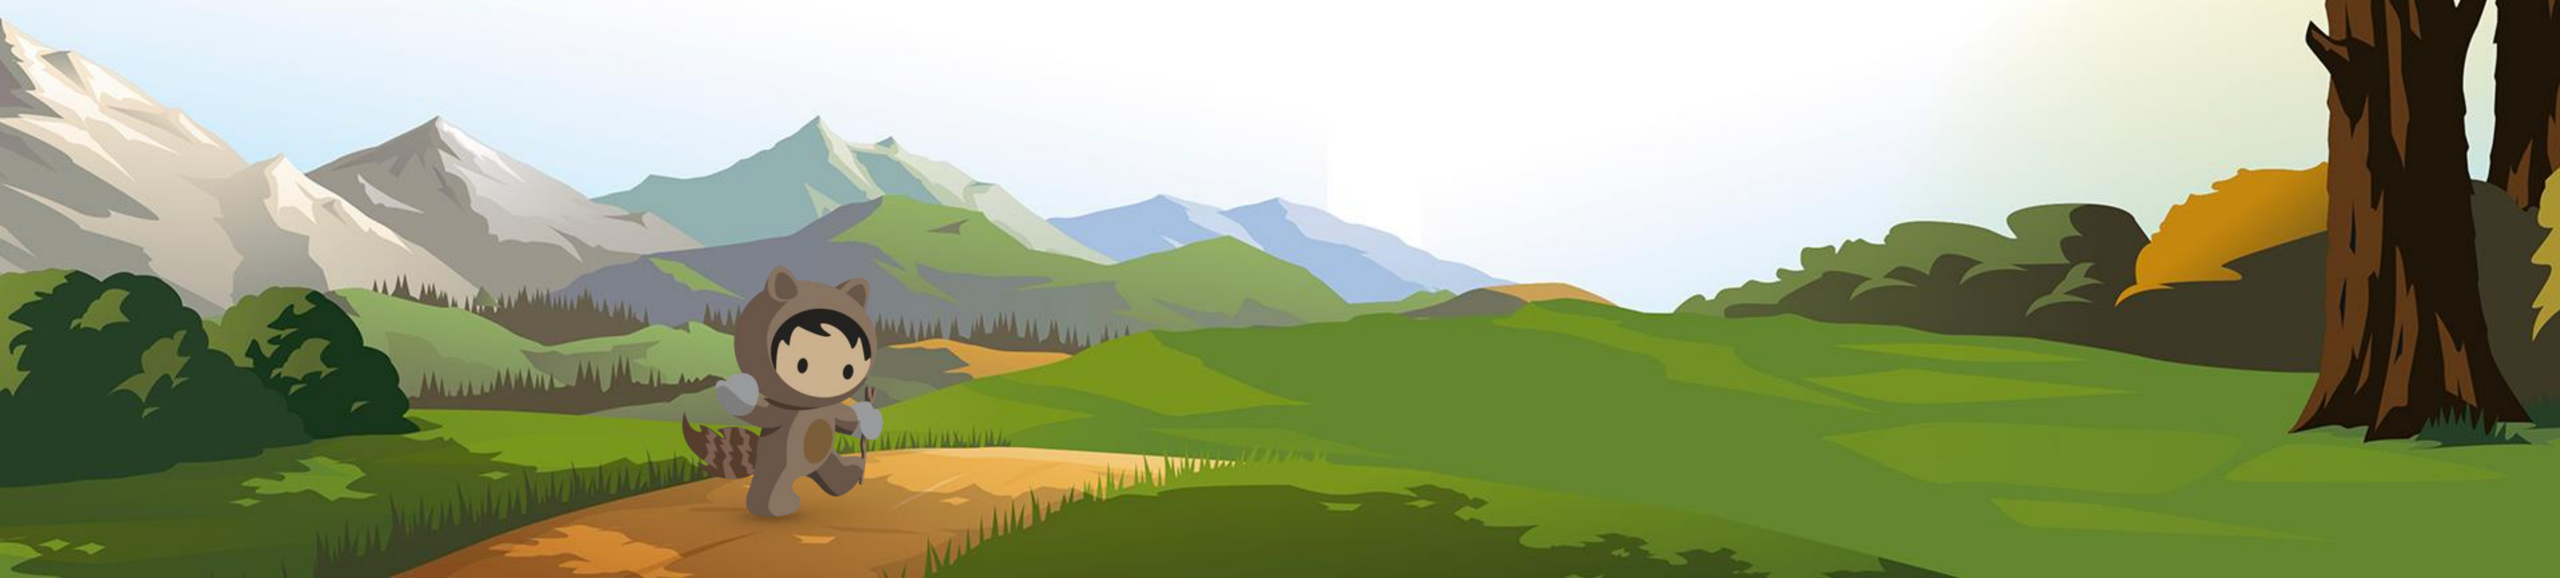

ACM Highlights

- Learning Center tools for professional development: http://learning.acm.org
  - The Safari Learning Platform featuring the **entire Safari collection of nearly 50,000** technical books, video courses, O'Reilly conference videos, learning paths, tutorials, case studies
  - 1,800+ Skillsoft courses, 4,800+ online books, and 30,000+ task-based short videos for software professionals covering programming, data management, DevOps, cybersecurity, networking, project management, and more; including training toward top vendor certifications such as AWS, CEH, Cisco, CISSP, CompTIA, Oracle, RedHat, PMI.
  - 1,200+ books from Elsevier on the ScienceDirect platform (including Morgan Kaufmann and Syngress titles)
  - TechTalks from thought leaders and top practitioners
  - Podcast interviews with innovators, entrepreneurs, and award winners
- Popular publications:
  - Flagship ***Communications of the ACM (CACM)*** magazine: http://cacm.acm.org
  - *ACM Queue* magazine for practitioners: http://queue.acm.org
- The **ACM Code of Ethics**, a set of principles and guidelines principles and guidelines designed to help computing professionals make ethically responsible decisions in professional practice: https://ethics.acm.org
  ACM Digital Library, the world's most comprehensive database of computing literature: http://dl.acm.org
- International conferences that draw leading experts on a broad spectrum of computing topics: http://www.acm.org/conferences
- Prestigious awards, including the **ACM A.M. Turing** and ACM Prize in Computing: http://awards.acm.org
- And much more… http://www.acm.org.

“Housekeeping”

Twitter: #ACMLearning

- Welcome to today’s ACM TechTalk, “**Agile Data Science: Achieving Salesforce-Scale Machine Learning in Production**.” The presentation starts at the top of the hour and lasts 60 minutes. Audio and video will automatically play throughout the event. On the bottom panel you’ll find a number of widgets, including Twitter and Sharing apps.

- If you are experiencing any problems/issues, **refresh** your console by pressing the **F5** key on your keyboard in Windows, **Command + R** if on a Mac, or refresh your browser if you’re on a mobile device; or close and re-launch the presentation. You can also view the Webcast Help Guide, by clicking on the “Help” widget in the bottom dock.

- To control volume, adjust the master volume on your computer. If the volume is still too low, use headphones.

- At the end of the presentation, you’ll see a survey open on your screen. Please take a minute to fill it out to help us improve your next webinar experience. You may also open the survey at any time throughout the presentation from the resources window.

- This session is being recorded and will be archived for on-demand viewing in a few days. You will receive an automatic email notification when it is available. See http://learning.acm.org/ for updates. And check out https://learning.acm.org/techtalks for archived recordings of past webcasts.

Talk Back

- Tweet your favorite quotes from today's presentation with hashtag #ACMLearning
- Submit questions and comments via Twitter to @acmeducation – we're reading them!
- The ACM Discourse Page is available for post-talk discussion – https://on.acm.org

Forward-Looking Statements

Statement under the Private Securities Litigation Reform Act of 1995:

This presentation may contain forward-looking statements that involve risks, uncertainties, and assumptions. If any such uncertainties materialize or if any of the assumptions proves incorrect, the results of salesforce.com, inc. could differ materially from the results expressed or implied by the forward-looking statements we make. All statements other than statements of historical fact could be deemed forward-looking, including any projections of product or service availability, subscriber growth, earnings, revenues, or other financial items and any statements regarding strategies or plans of management for future operations, statements of belief, any statements concerning new, planned, or upgraded services or technology developments and customer contracts or use of our services.

The risks and uncertainties referred to above include – but are not limited to – risks associated with developing and delivering new functionality for our service, new products and services, our new business model, our past operating losses, possible fluctuations in our operating results and rate of growth, interruptions or delays in our Web hosting, breach of our security measures, the outcome of any litigation, risks associated with completed and any possible mergers and acquisitions, the immature market in which we operate, our relatively limited operating history, our ability to expand, retain, and motivate our employees and manage our growth, new releases of our service and successful customer deployment, our limited history reselling non-salesforce.com products, and utilization and selling to larger enterprise customers. Further information on potential factors that could affect the financial results of salesforce.com, inc. is included in our annual report on Form 10-K for the most recent fiscal year and in our quarterly report on Form 10-Q for the most recent fiscal quarter. These documents and others containing important disclosures are available on the SEC Filings section of the Investor Information section of our Web site.

Any unreleased services or features referenced in this or other presentations, press releases or public statements are not currently available and may not be delivered on time or at all. Customers who purchase our services should make the purchase decisions based upon features that are currently available. Salesforce.com, inc. assumes no obligation and does not intend to update these forward-looking statements.

Different Flavors of AI and ML in Industry

Models that inform strategic decisions

<u>Examples</u>

Data–driven drug discovery

Risk models for investments

Models that are products

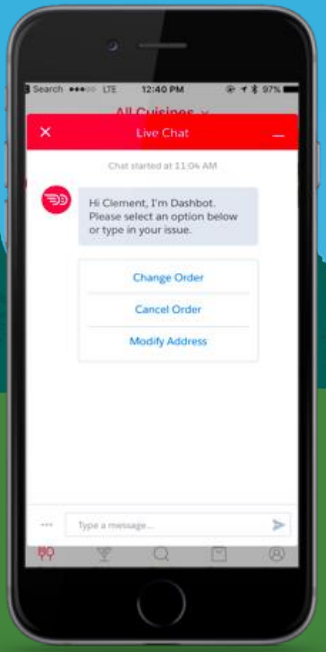

<u>Examples</u>

Chatbots

Algorithmic Trading

Models that augment products

<u>Examples</u>

Predictive Lead Scoring

Case Classification

salesforce

Adoption of AI is Considered Critical to Stay Competitive!

FIGURE 2

AI helps organizations keep up with the (Dow) Joneses

Relative to competitors, respondents say their company's adoption of AI has allowed them to . . .

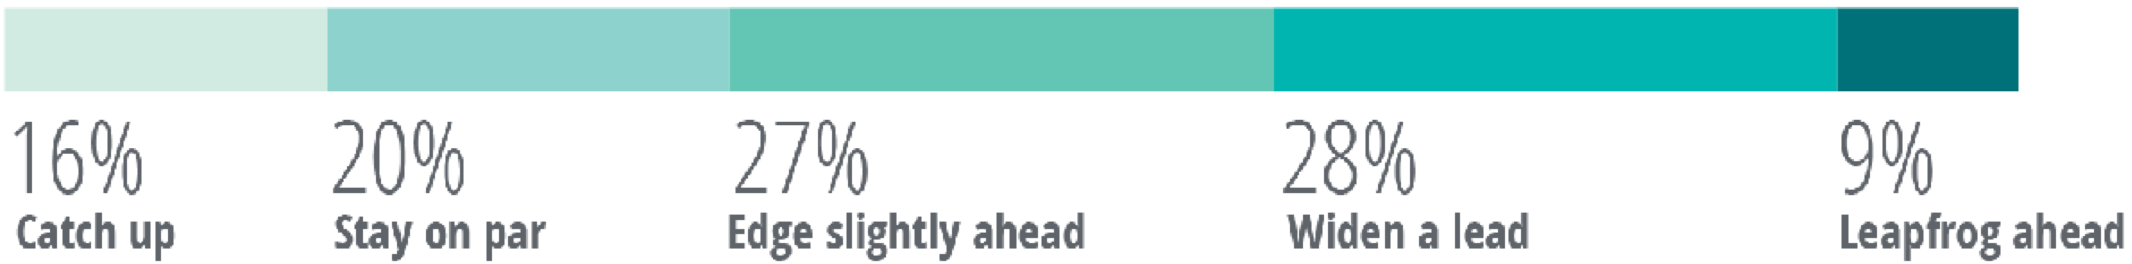

Source: Deloitte State of AI in the Enterprise, 2nd Edition, 2018.

Deloitte Insights | deloitte.com/insights

For the Majority of Businesses, Data Science is Out of Reach

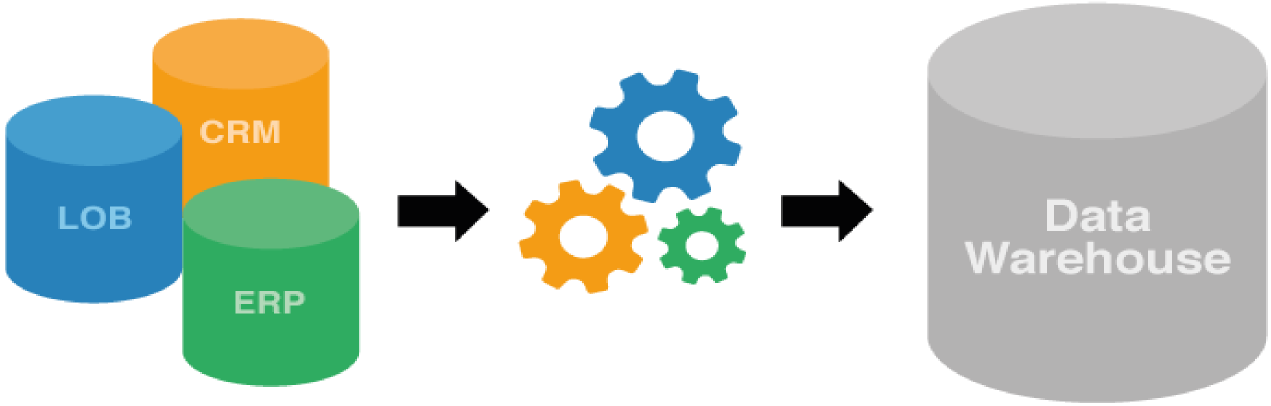

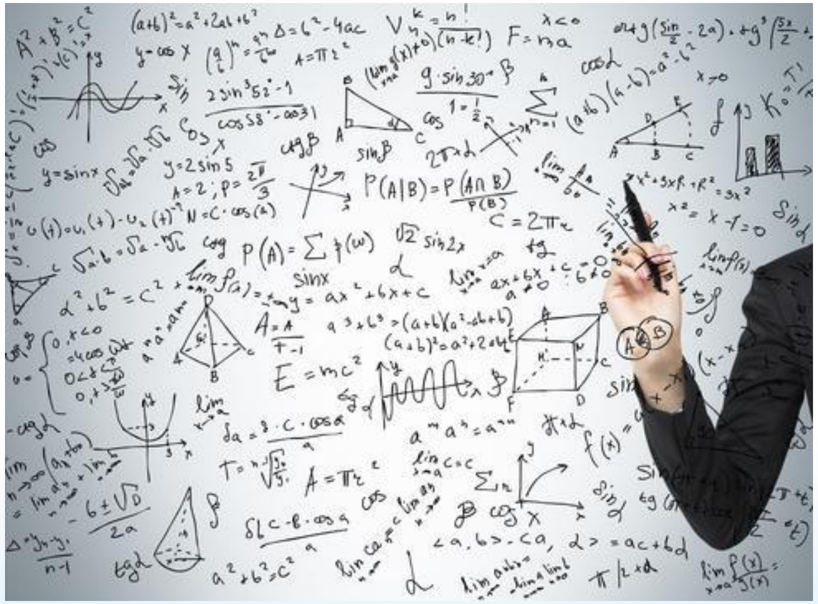

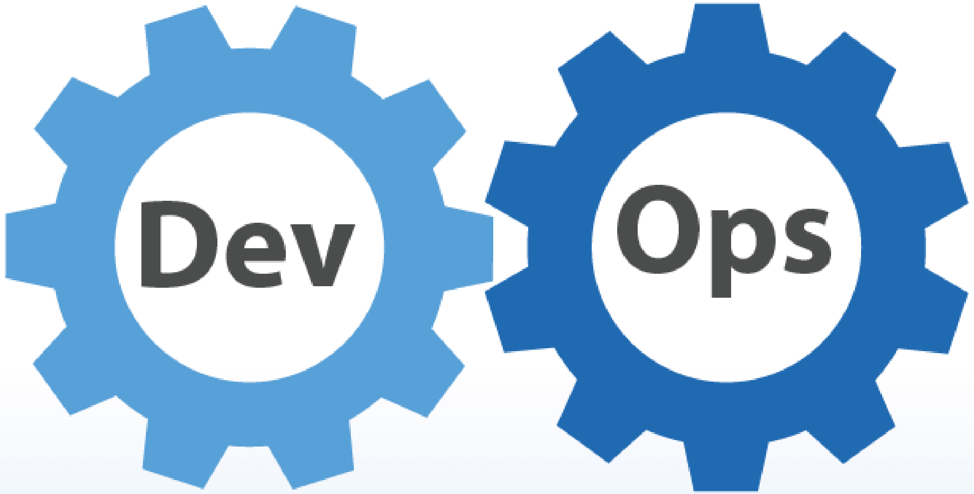

Democratizing Data Science as the Key to Meeting Demand

salesforce

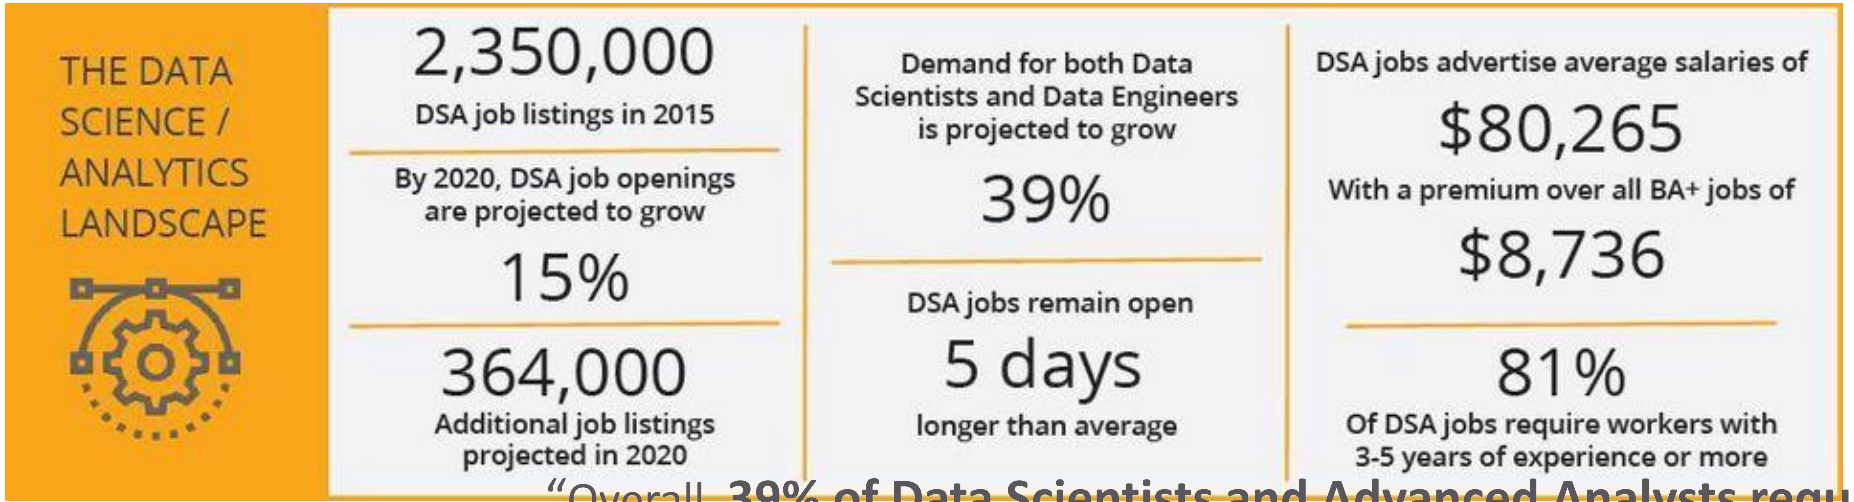

burningglass CAREERS IN FOCUS

IBM

BHEF BUSINESS-HIGHER EDUCATION FORUM

Creating Solutions. Inspiring Action.

"Overall, **39% of Data Scientists and Advanced Analysts require a master's or Ph.D** … Therefore, because these roles are already **undersupplied** and projected to grow rapidly, **the skills shortage is in danger of worsening**."

Source: THE QUANT CRUNCH: HOW THE DEMAND FOR DATA SCIENCE SKILLS IS DISRUPTING THE JOB MARKET, ©Burning Glass Technologies 2017
HTTPS://WWW.IBM.COM/ANALYTICS/US/EN/TECHNOLOGY/DATA-SCIENCE/QUANT-CRUNCH.HTML

Empowering Every Admin & Developer with AI

The Einstein platform

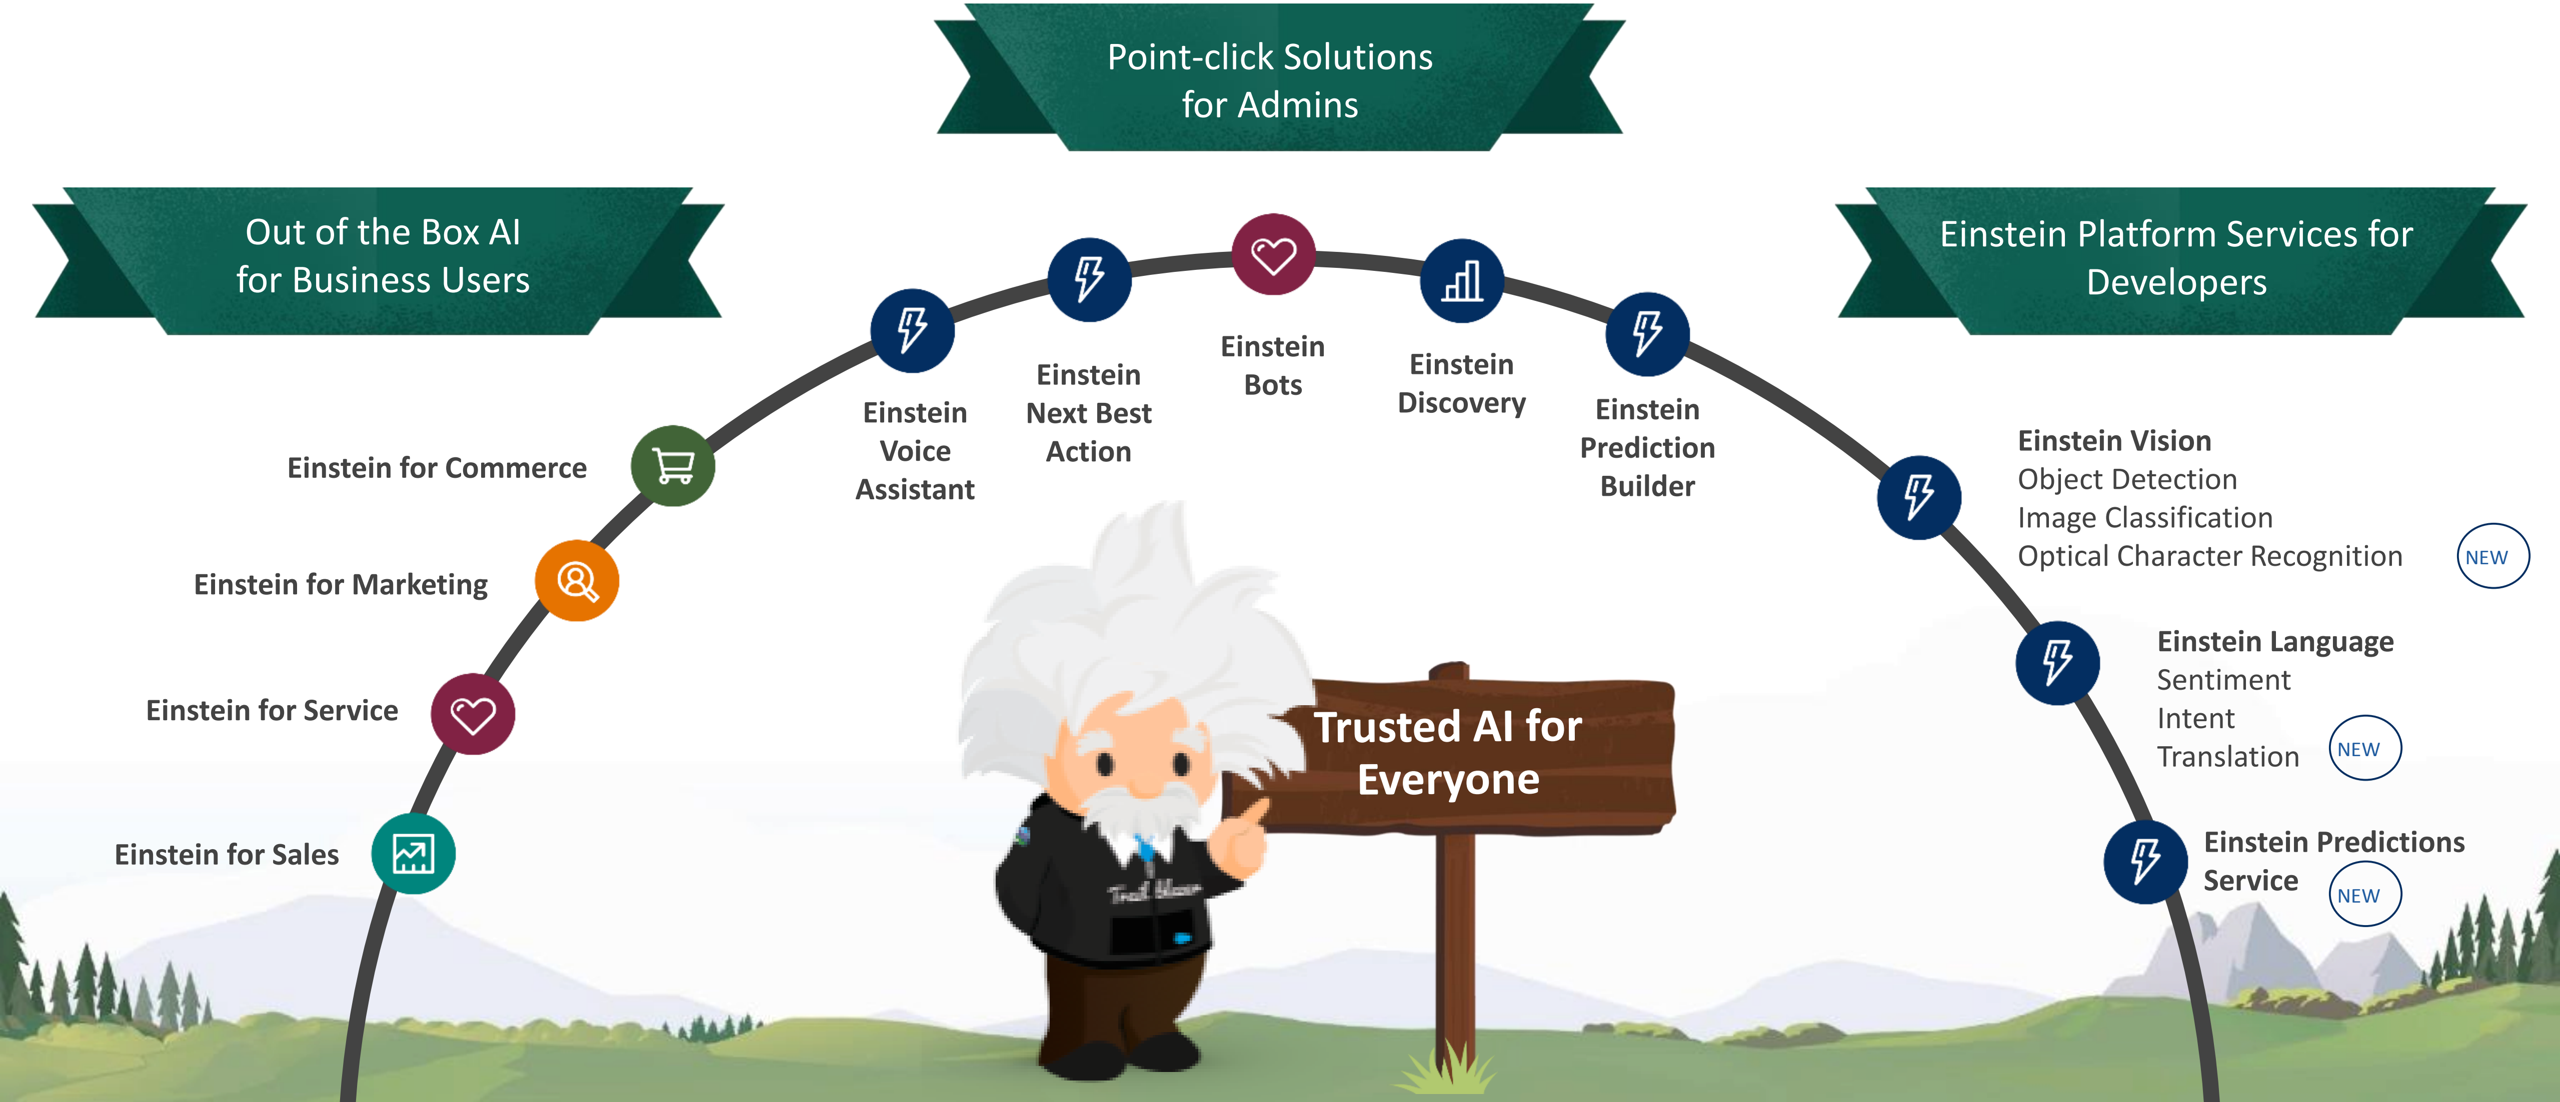

How we achieve Salesforce-scale!

Salesforce approach to democratizing AI

Enabling our customers to build models on their own data

The need for platform to ship AI to production

Bridging the communication gap between data scientists and software developers to find common ground and get to production and agility

Critical components of an AI platform

How to build a platform to support agile data science

How metrics drive agility and scale

How to apply agile methodologies to rapidly improve and deploy models

How Companies Build ML Apps

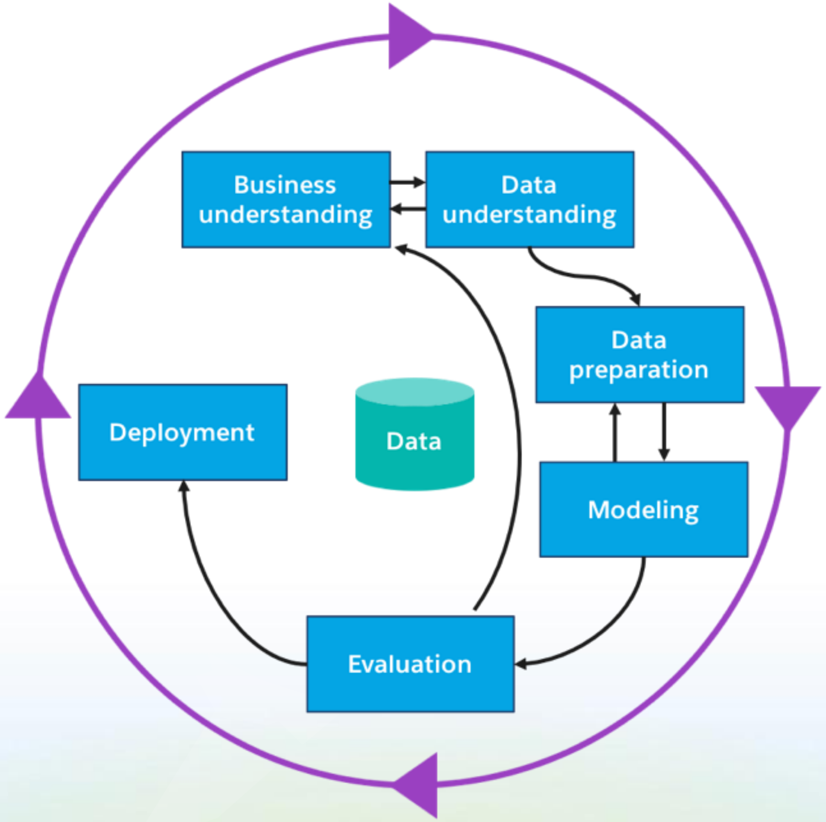

How Companies Build ML Apps

Data Scientists on App #1

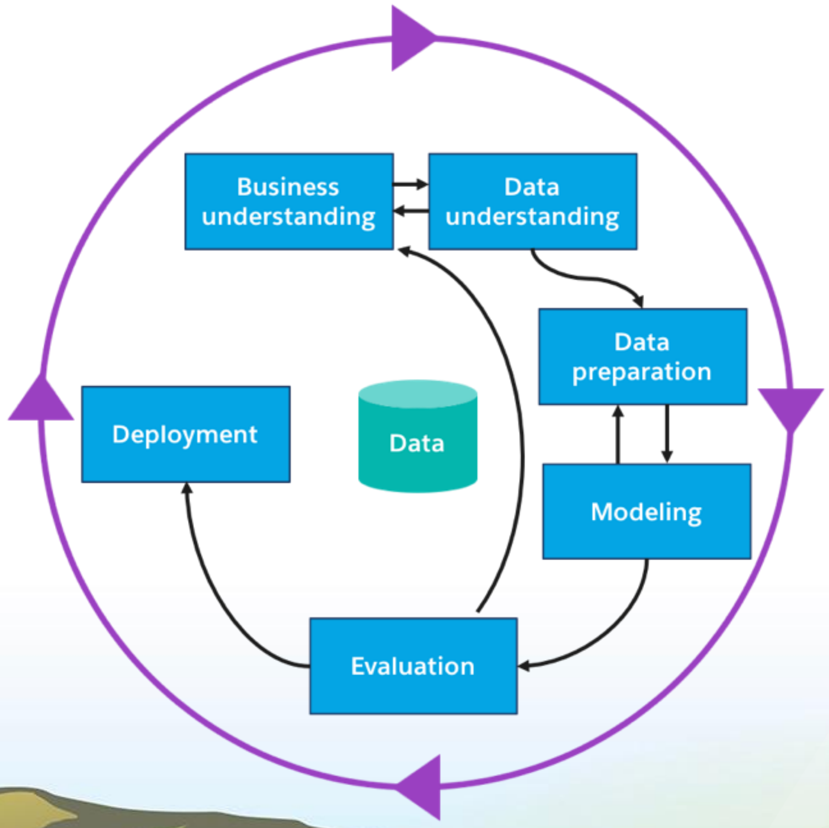

Data Scientists on App #2

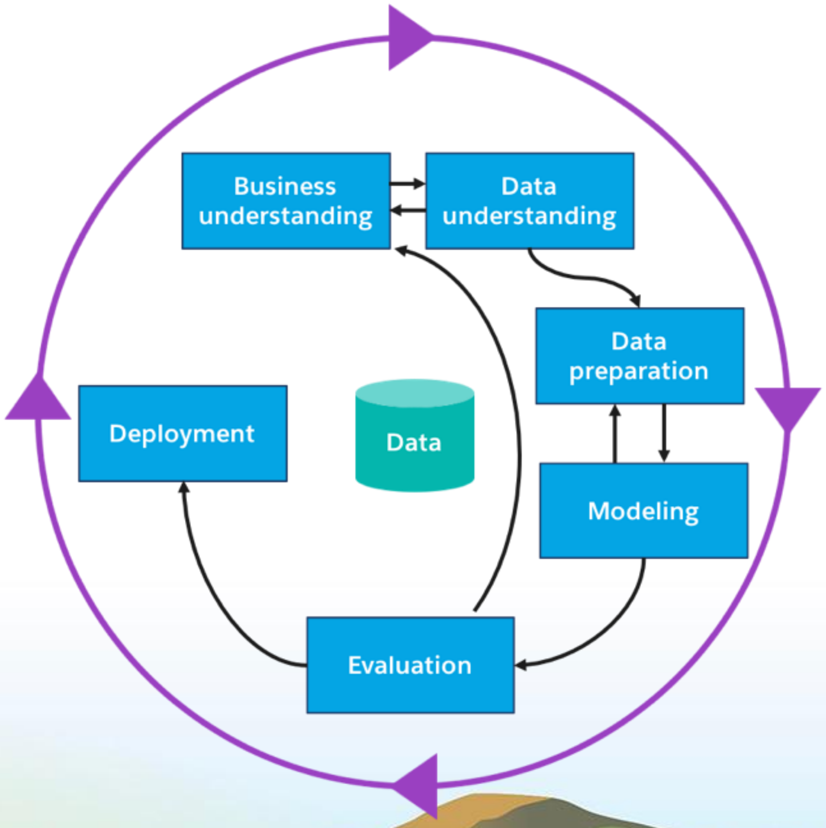

Let’s Add a Third App

Data Scientists on App #1

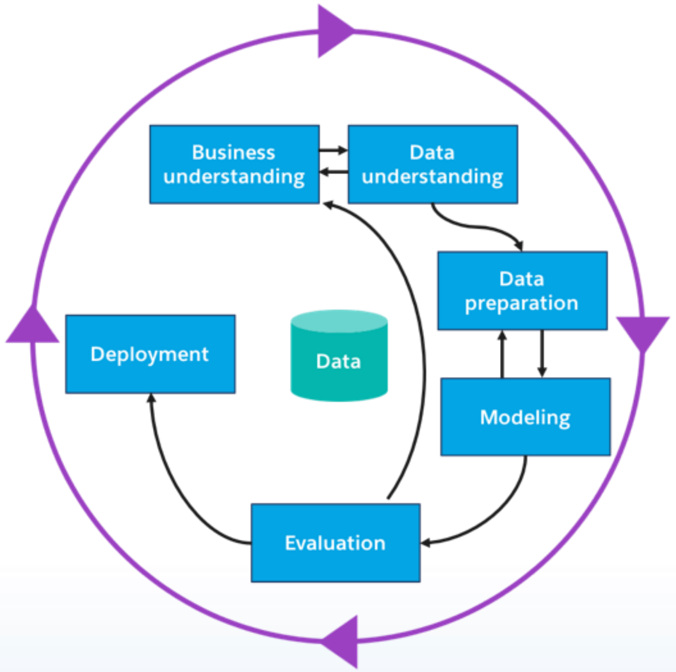

Data Scientists on App #2

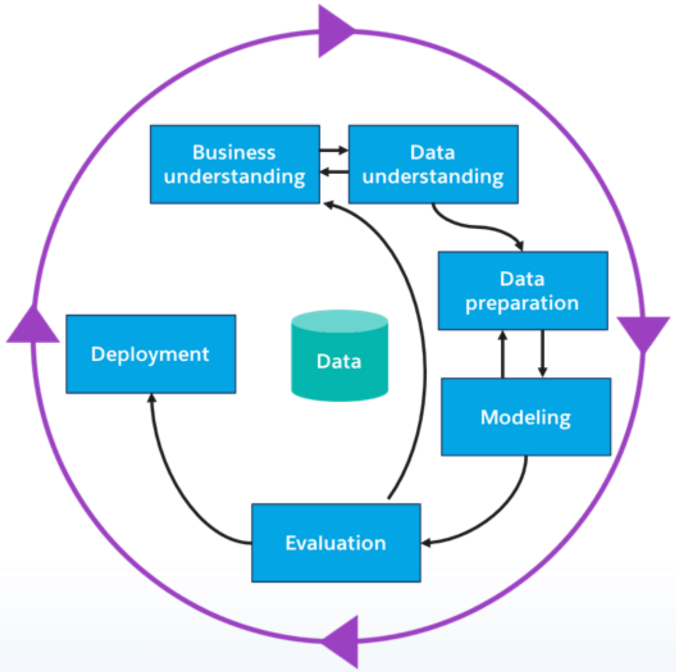

Data Scientists on App #3

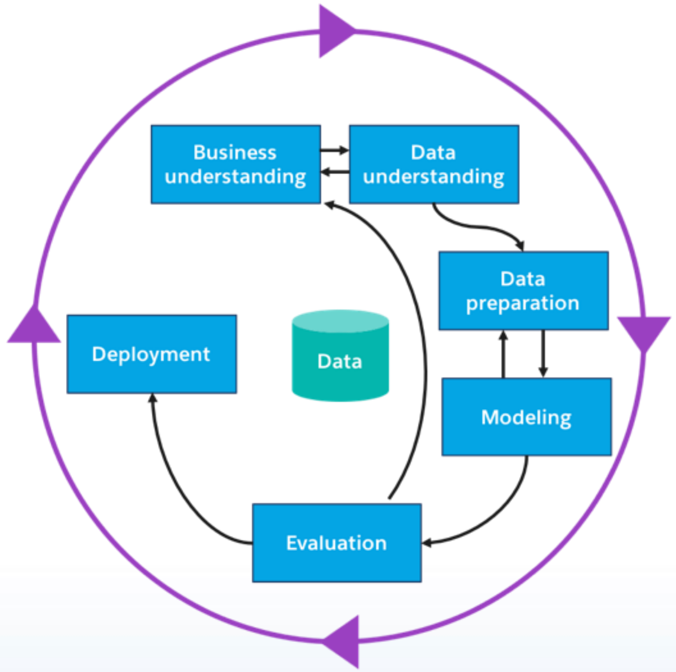

How This Process Would Look in Salesforce

Customer #1

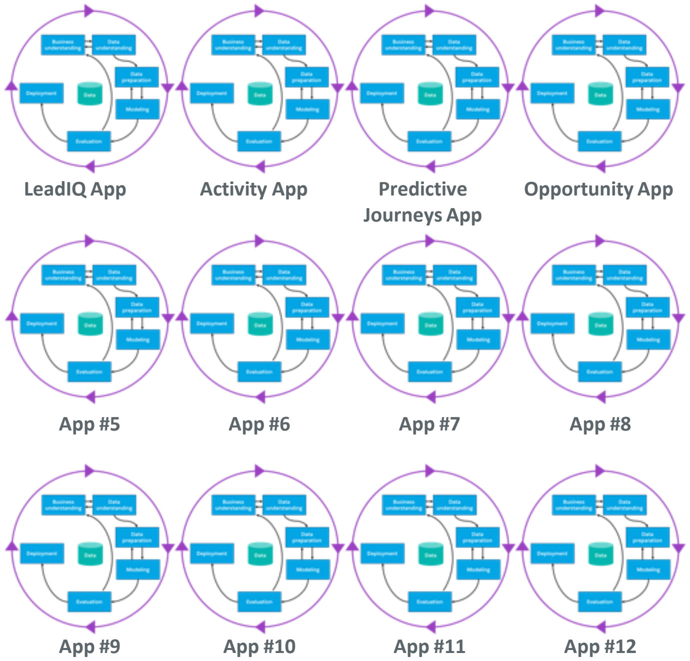

Customer #2

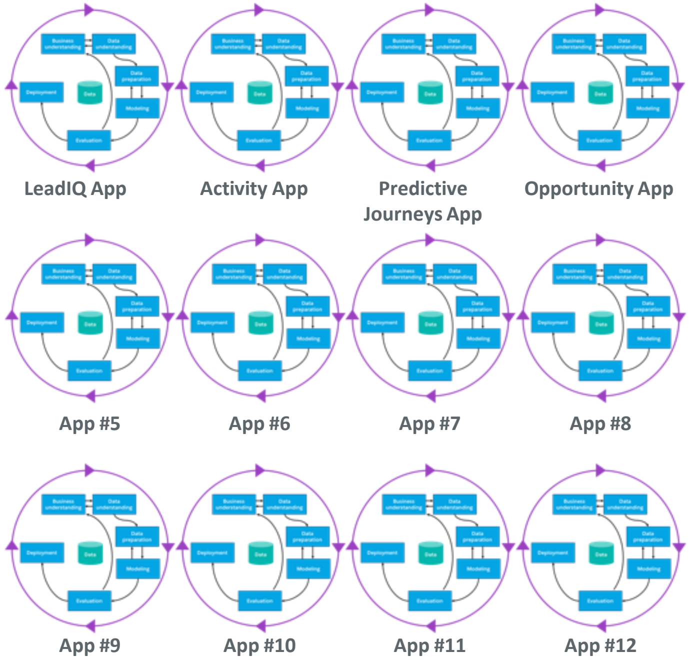

Customer #3

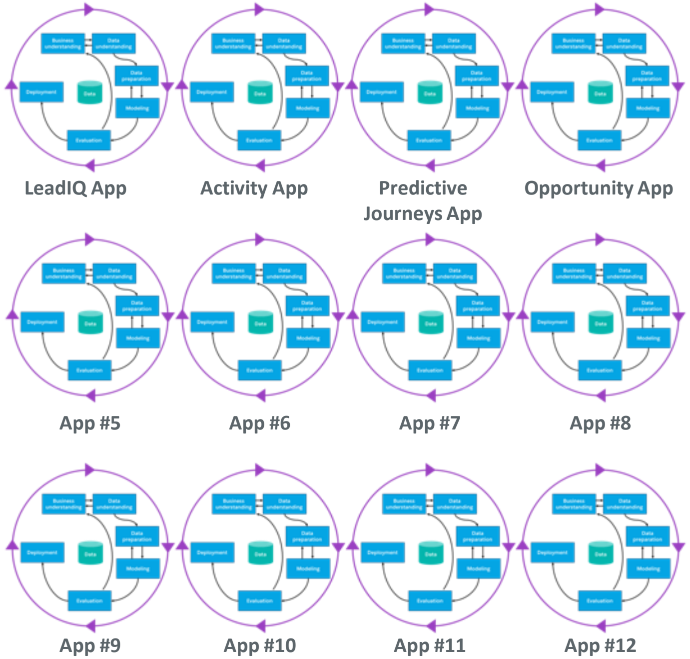

150,000 customers

There are varying degrees of skillsets

Different customers have different data sizes

Classification is not always classification

English is not the only language

Customers love to customize

AI needs to be trusted

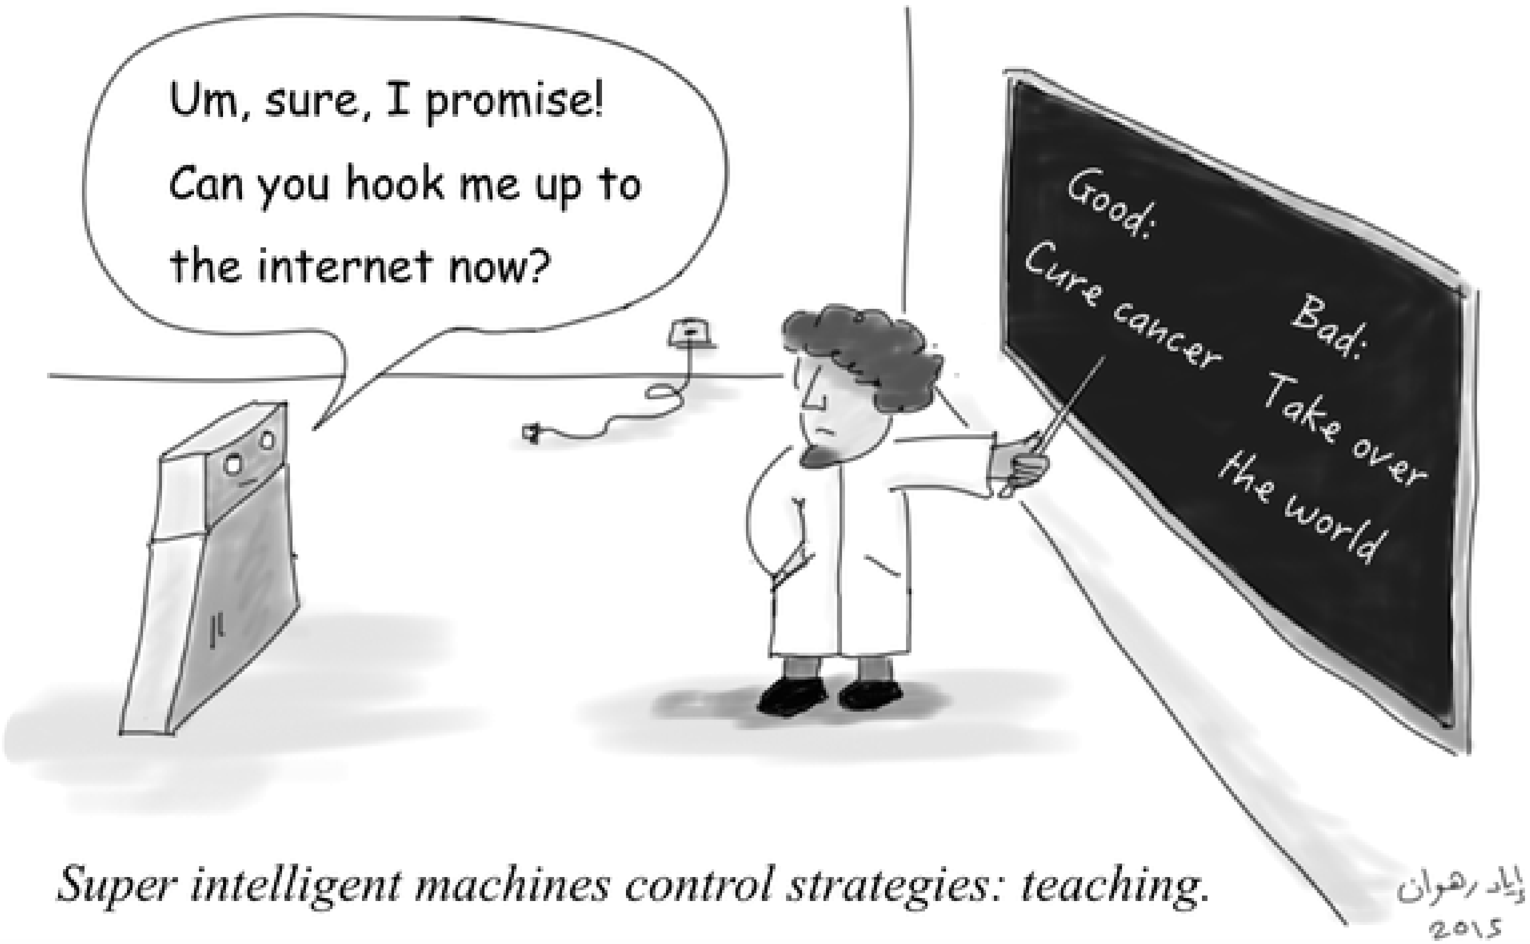

Super intelligent machines control strategies: teaching.

Fix your leaks

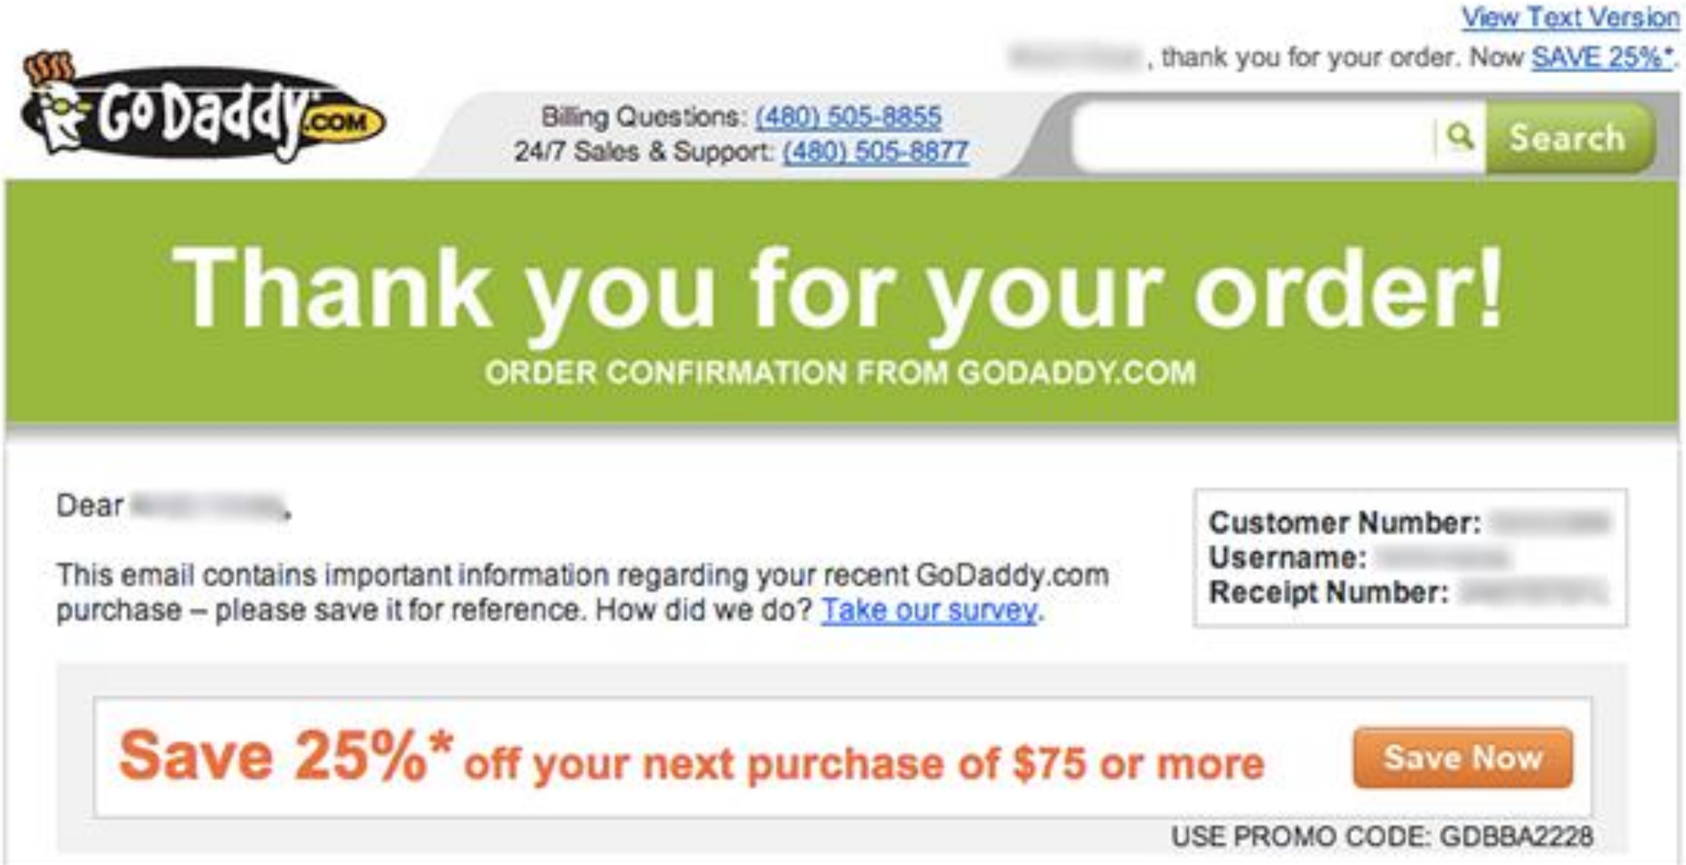

Einstein's New Approach to AI

Democratizing AI for Everyone

Classical Approach → Data Sampling → Feature Selection → Model Selection → Score Calibration → Integrate to Application → Artificial Intelligence

Einstein Auto-ML → AI for CRM
Discover
Predict
Recommend
Automate

Data already prepped
Models automatically built
Predictions delivered in context

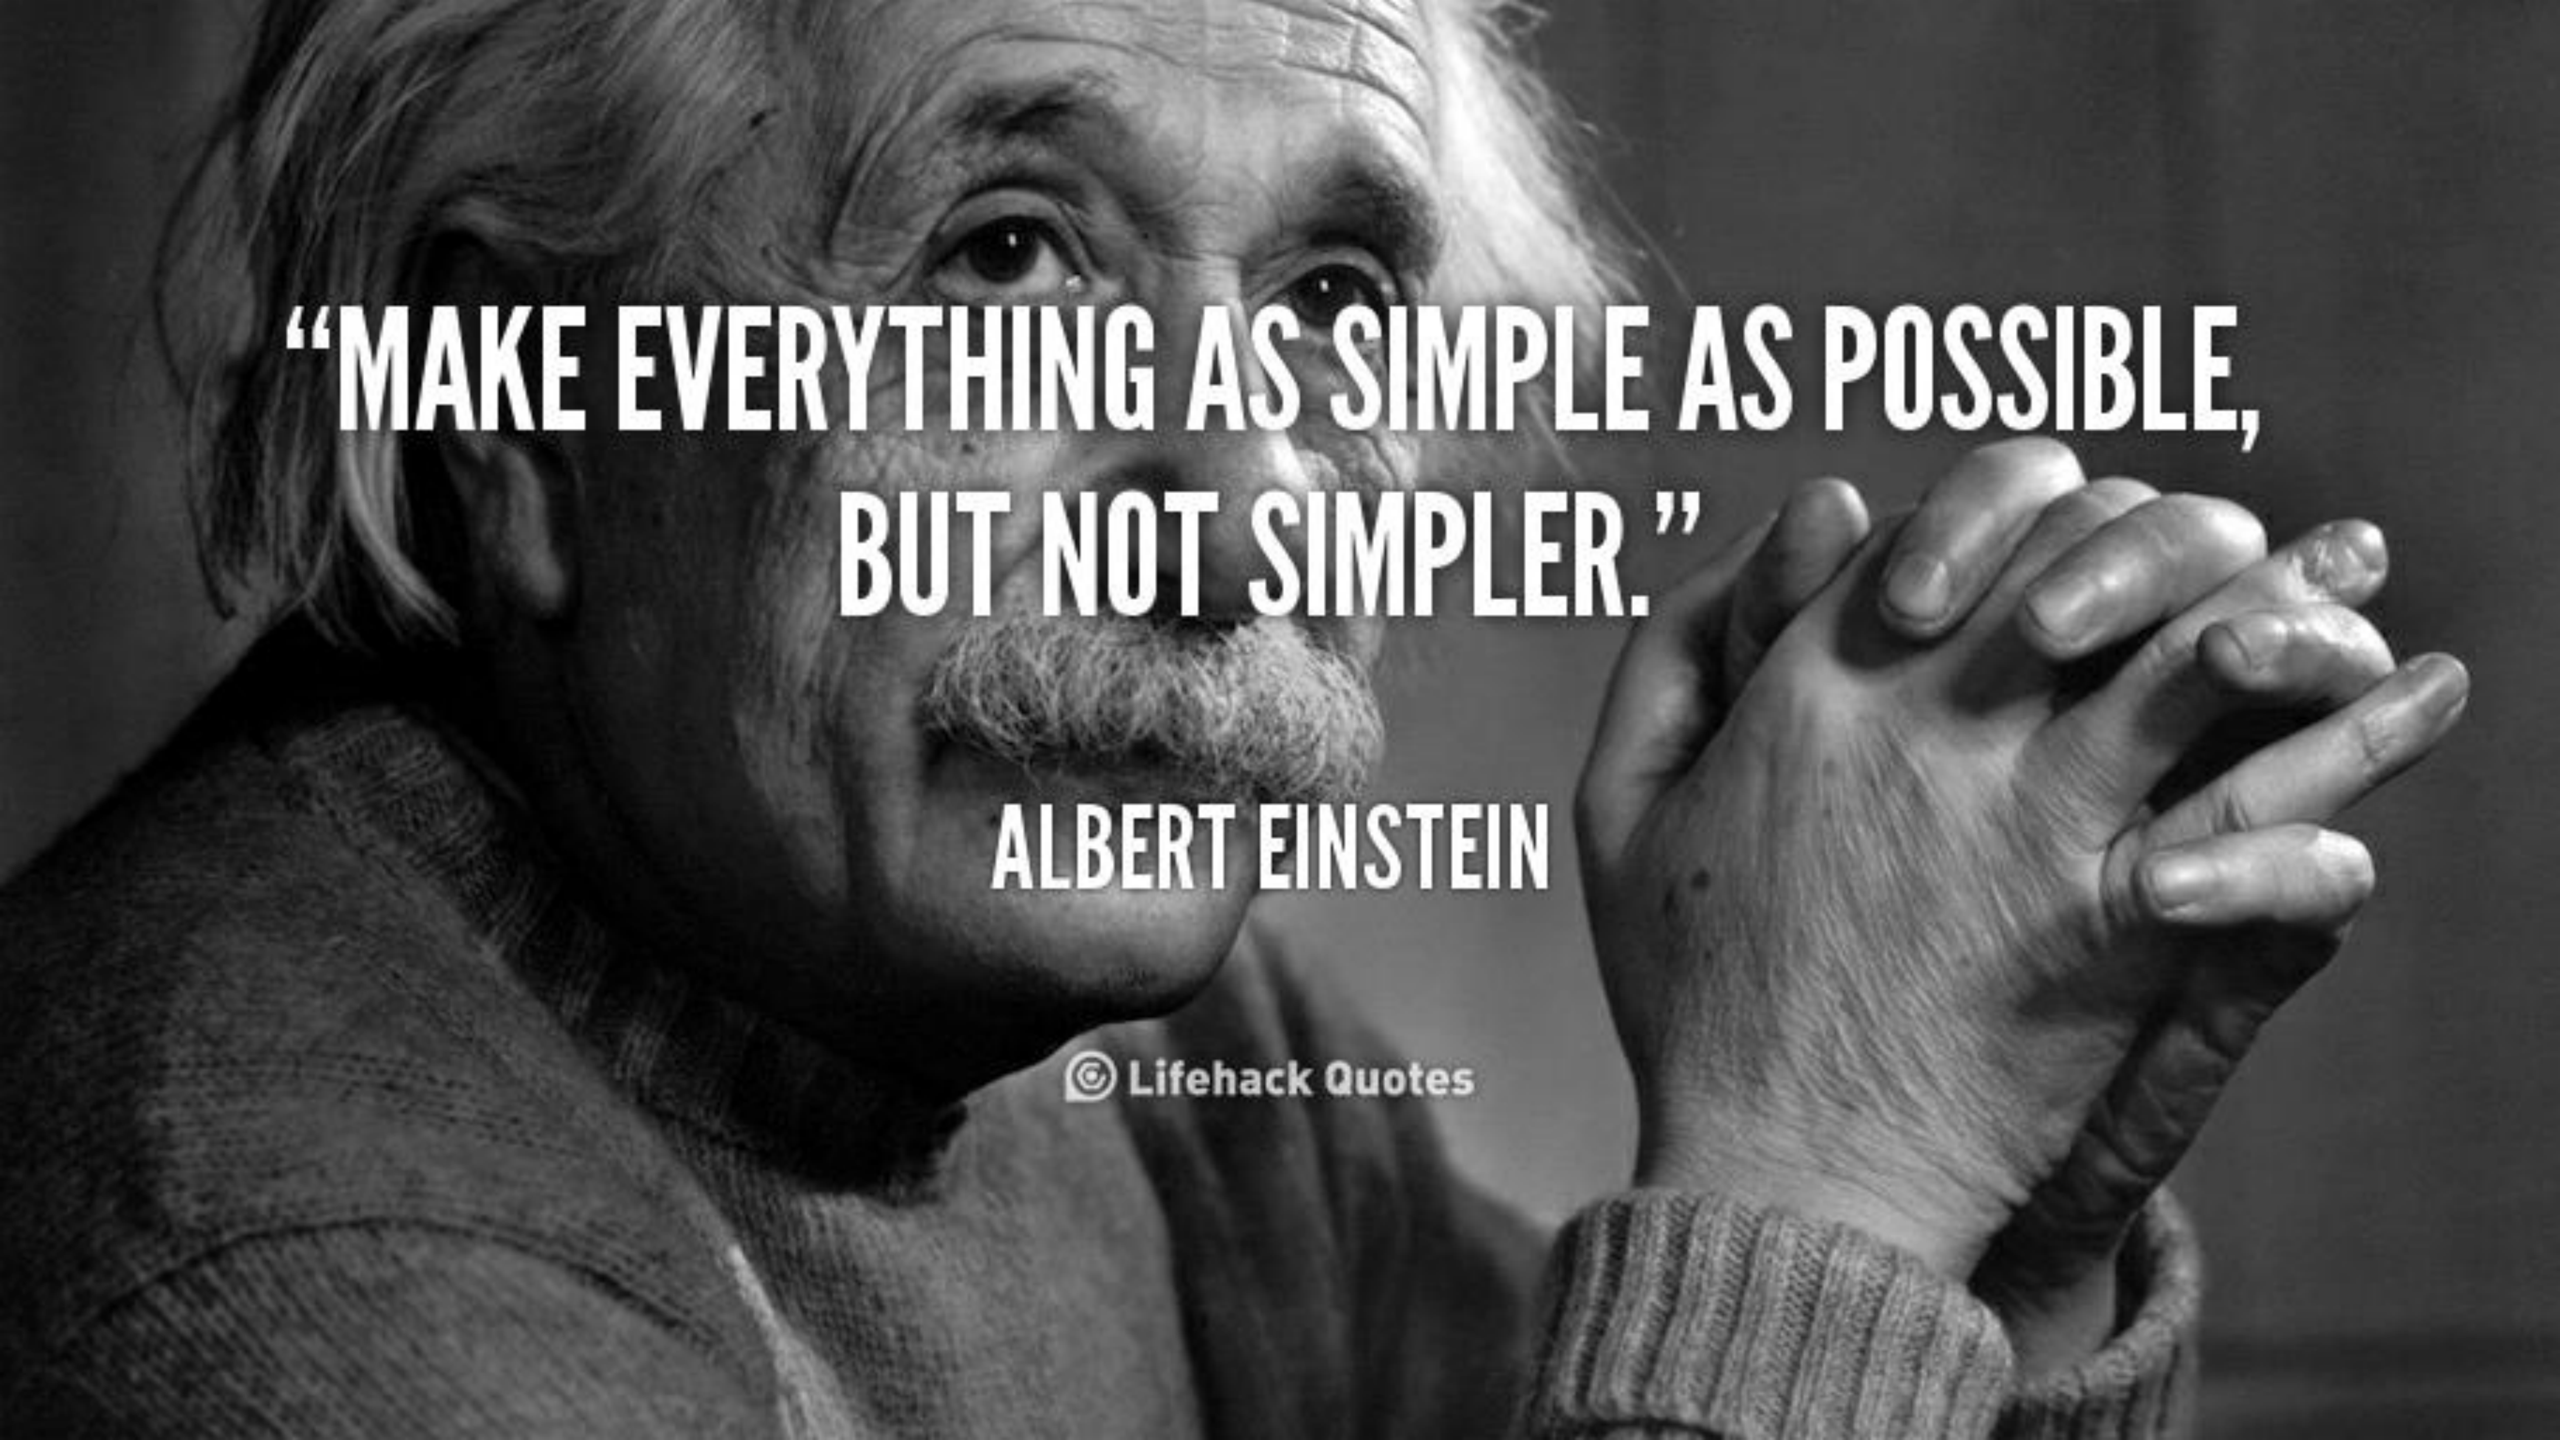
"MAKE EVERYTHING AS SIMPLE AS POSSIBLE, BUT NOT SIMPLER."
ALBERT EINSTEIN
© Lifehack Quotes

How we achieve Salesforce-scale!

Salesforce approach to democratizing AI

Enabling our customers to build models on their own data

The need for platform to ship AI to production

Bridging the communication gap between data scientists and software developers to find common ground and get to production and agility

Critical components of an AI platform

How to build a platform to support agile data science

How metrics drive agility and scale

How to apply agile methodologies to rapidly improve and deploy models

A data scientist’s view of the journey to building models

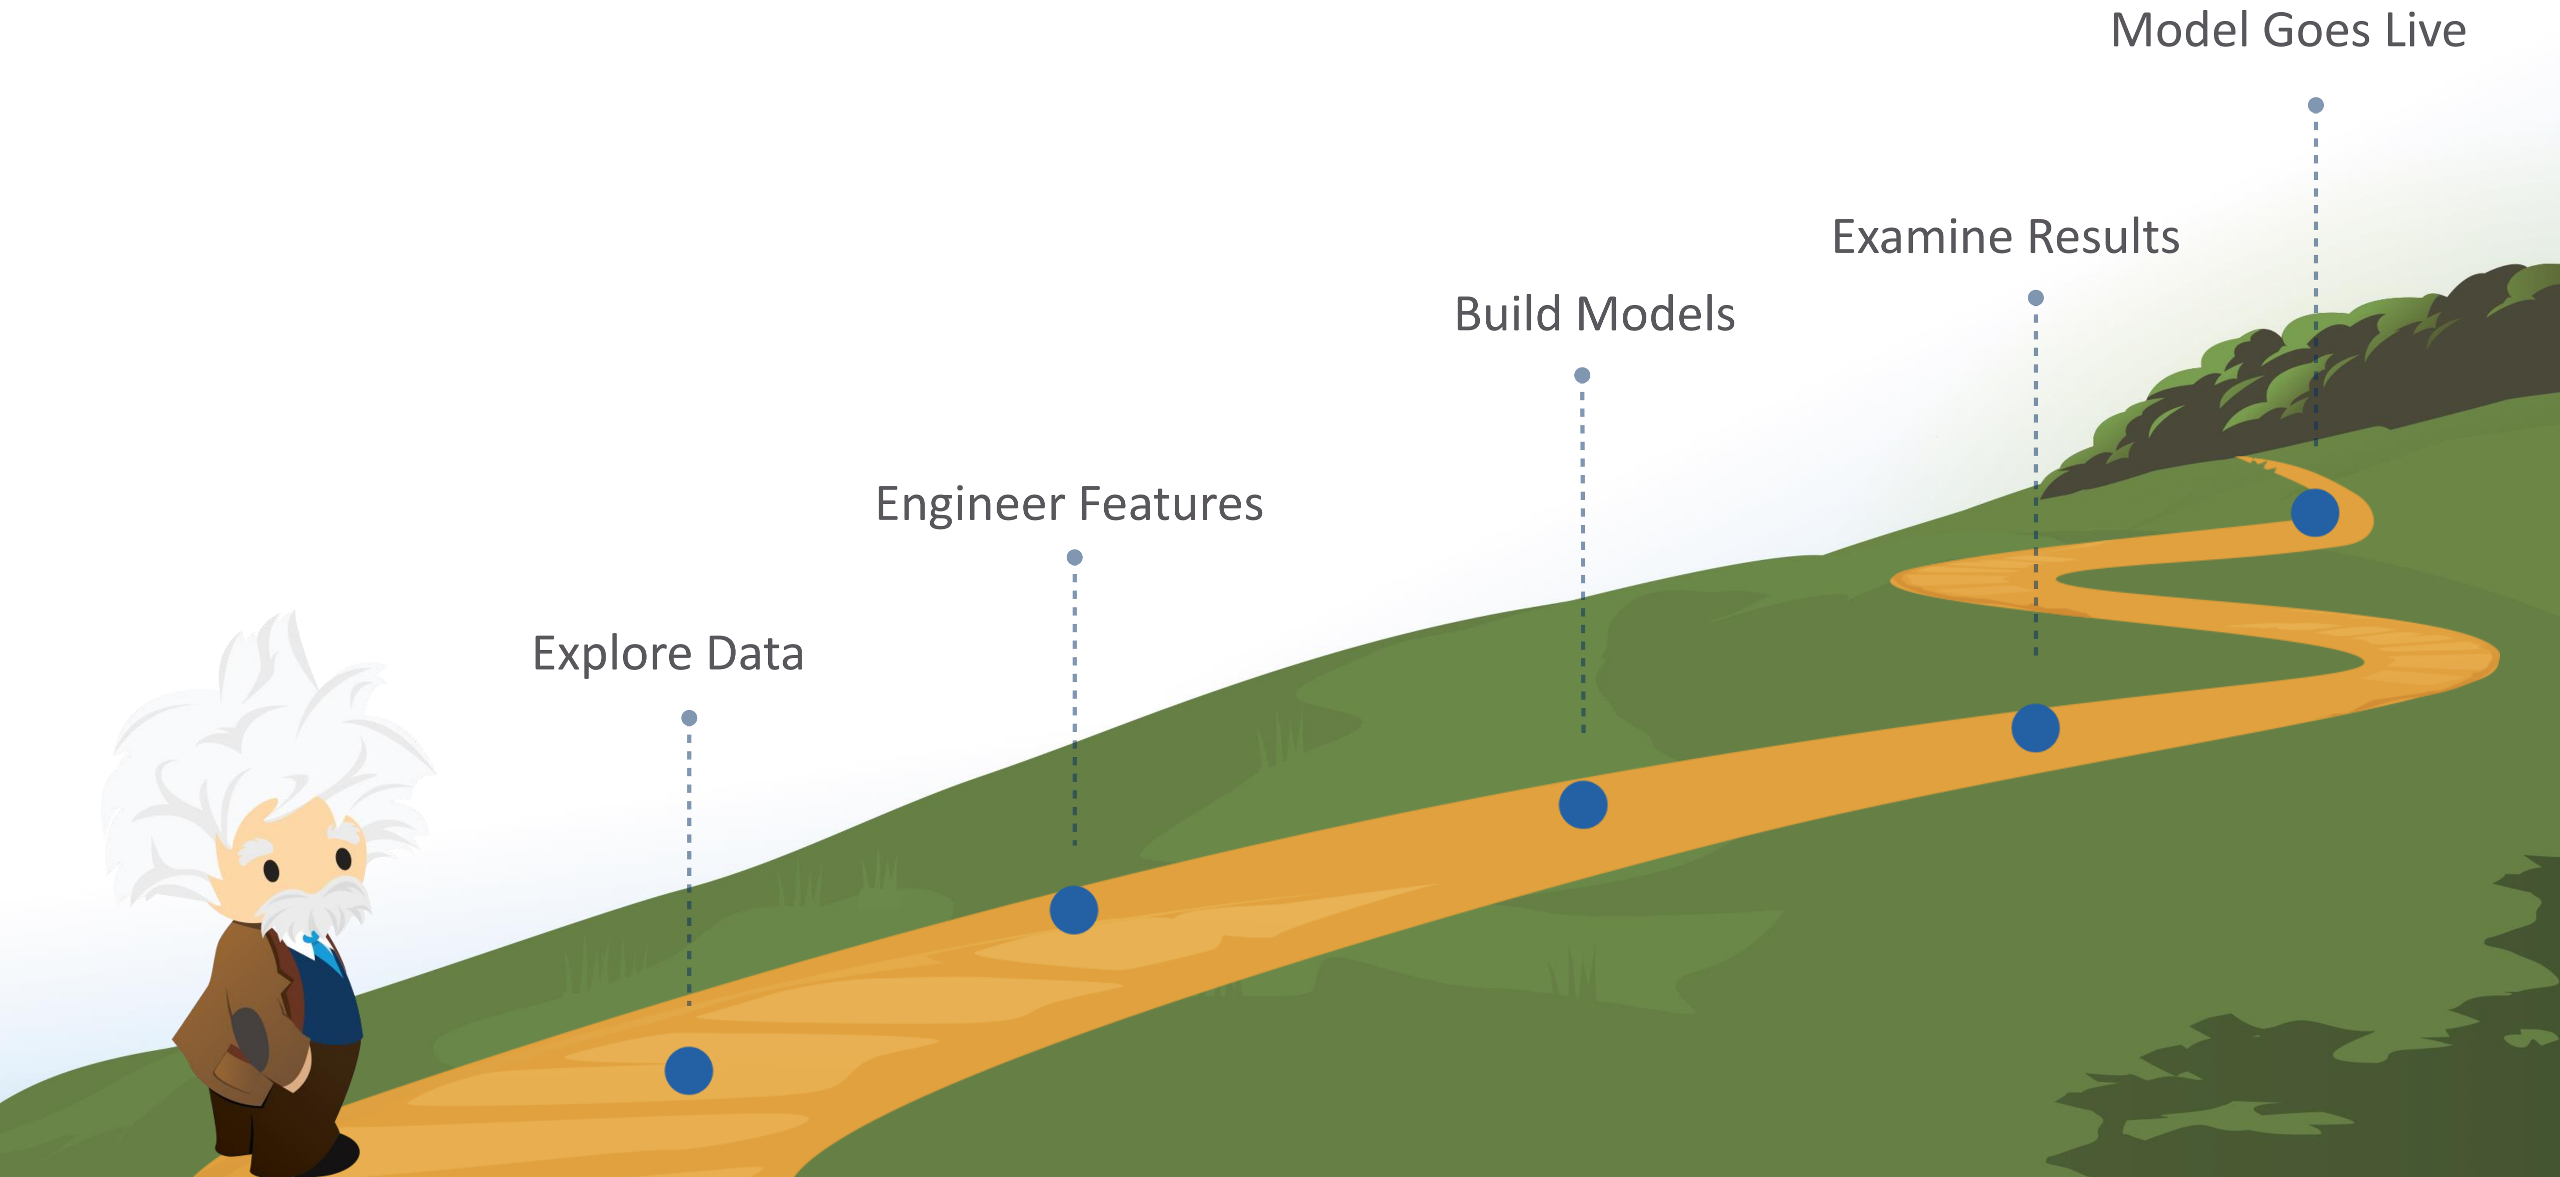

“The single biggest problem in communication is the illusion that it has taken place.”

- George Bernard Shaw

What are critical components to shipping your app!

APPLICATION to reach customers

PIPELINES to deliver data to ***modeling*** and ***scoring services***

MONITORS to know the health of models

EXPERIMENTATION frameworks and ***AGILE PROCESS*** to iteratively improve

WAY TO DEPLOY new models

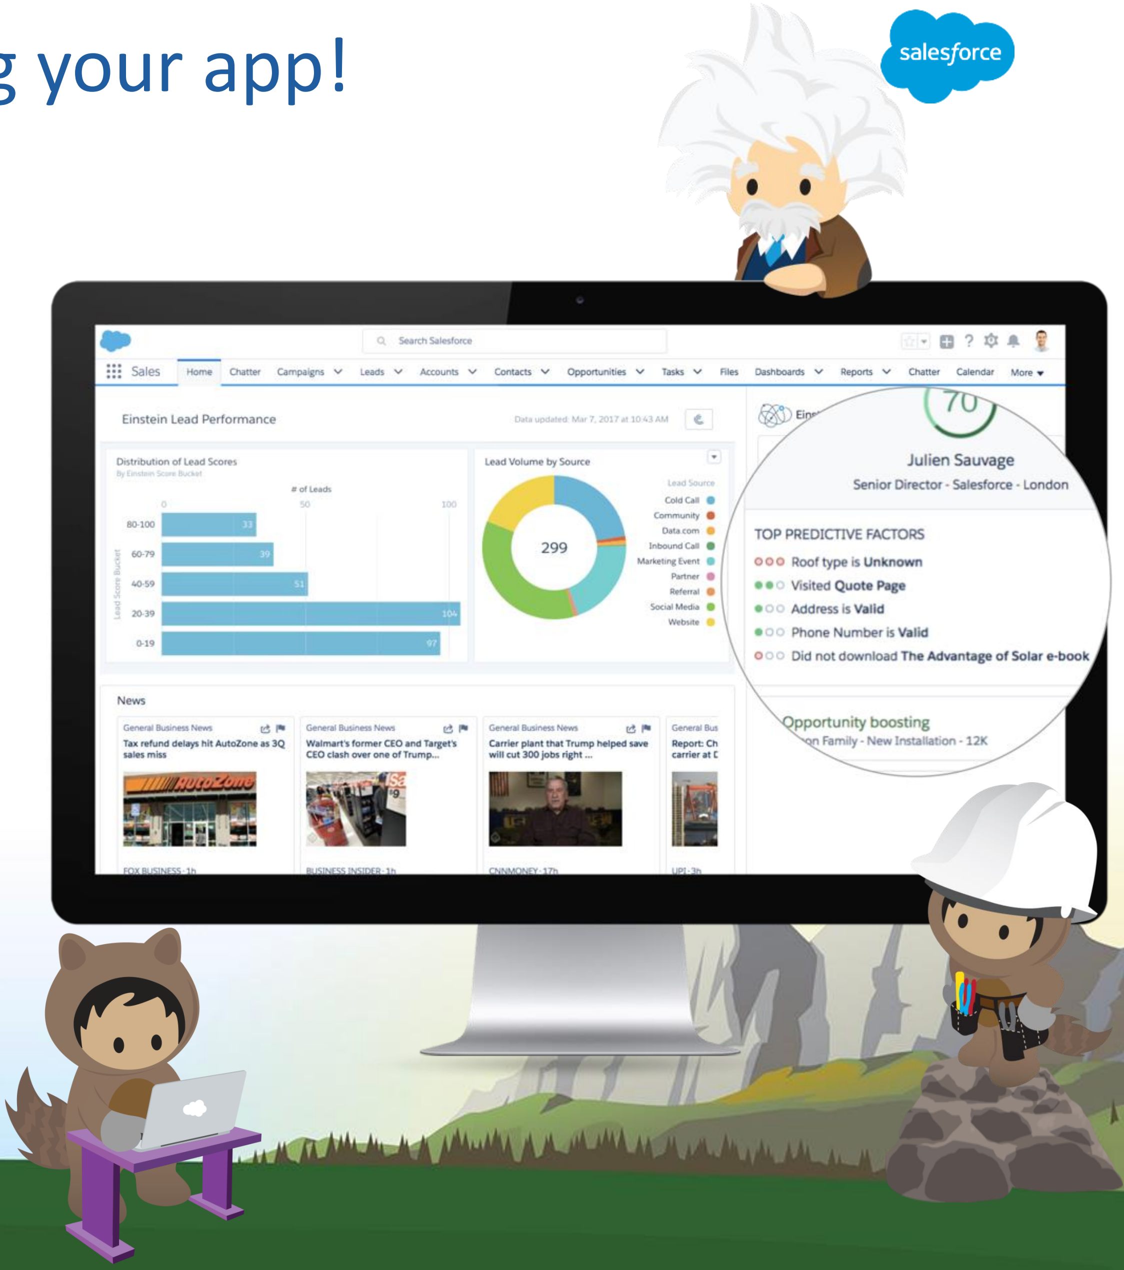

Source: Salesforce Customer Relationship Survey conducted 2014-2016 among 10,500+ customers randomly selected. Response sizes per question vary.

FIGURE 8

Companies need a broad range of skills for their AI initiatives

Respondents rating each a top-2 needed skill to fill their company's AI skills gap

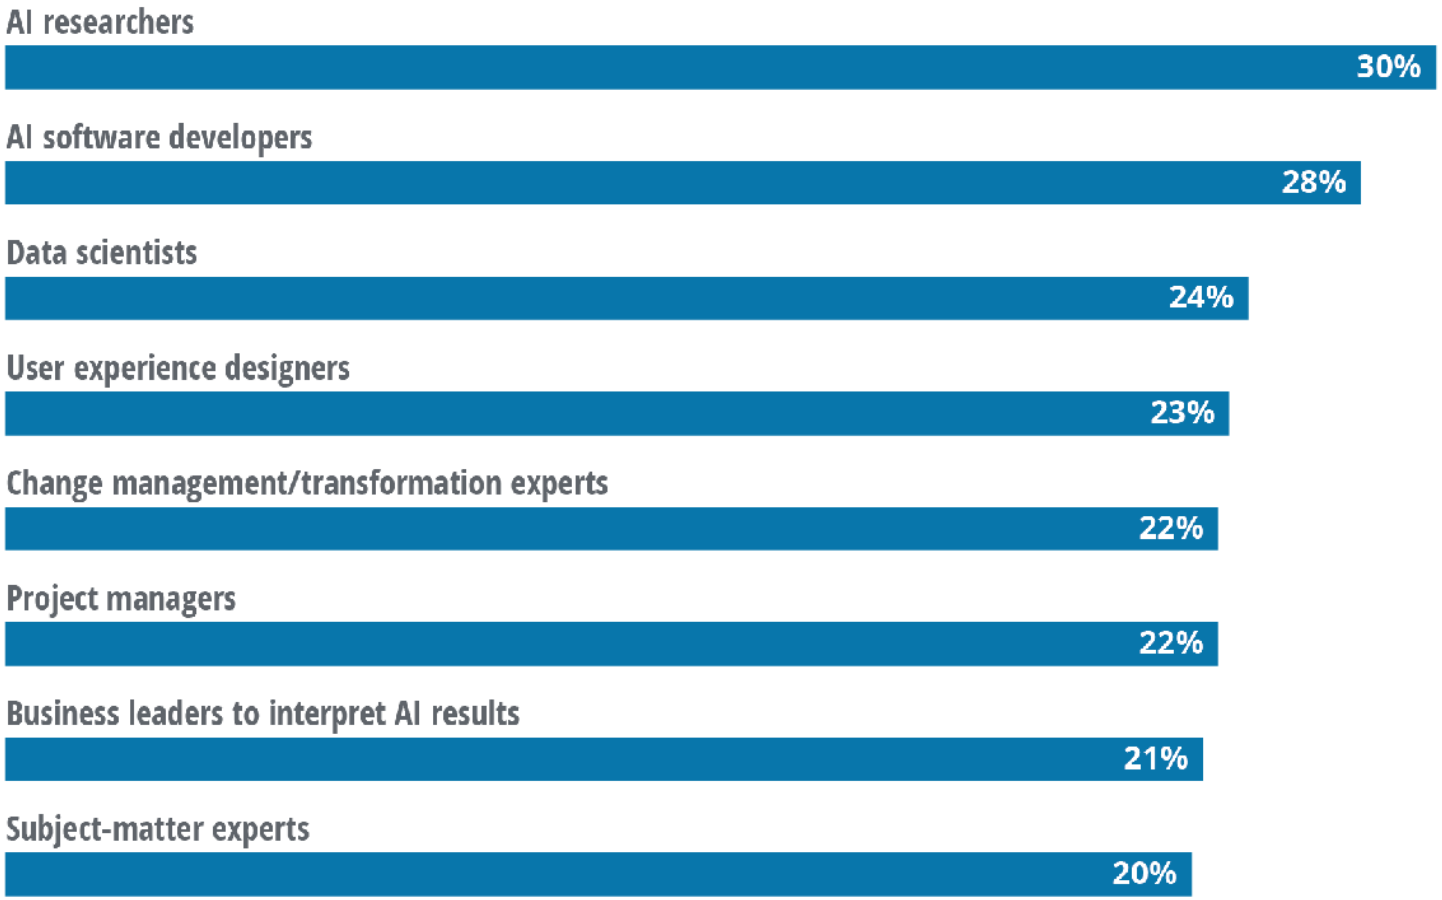

Note: Base = those who said that their company has moderate/major/extreme skills gap in meeting the needs of AI/cognitive projects. Sample size = 752.

Source: Deloitte State of AI in the Enterprise, 2nd Edition, 2018.

How different are data scientist and software developers?

Data Scientists

Monitor the performance of their models

Identify opportunities to improve models

Want to explore new data/algorithms

Need processes to test new models

Need a way to redeploy new models

Find opportunities for reuse

Software Developers

Monitor the performance of their apps

Identify opportunities to add features

Want to explore new technology

Need processes to test new features

Need a way to redeploy their app

Find opportunities for reuse


Give your team the tools they need!

Supporting a Model in Production is Complex

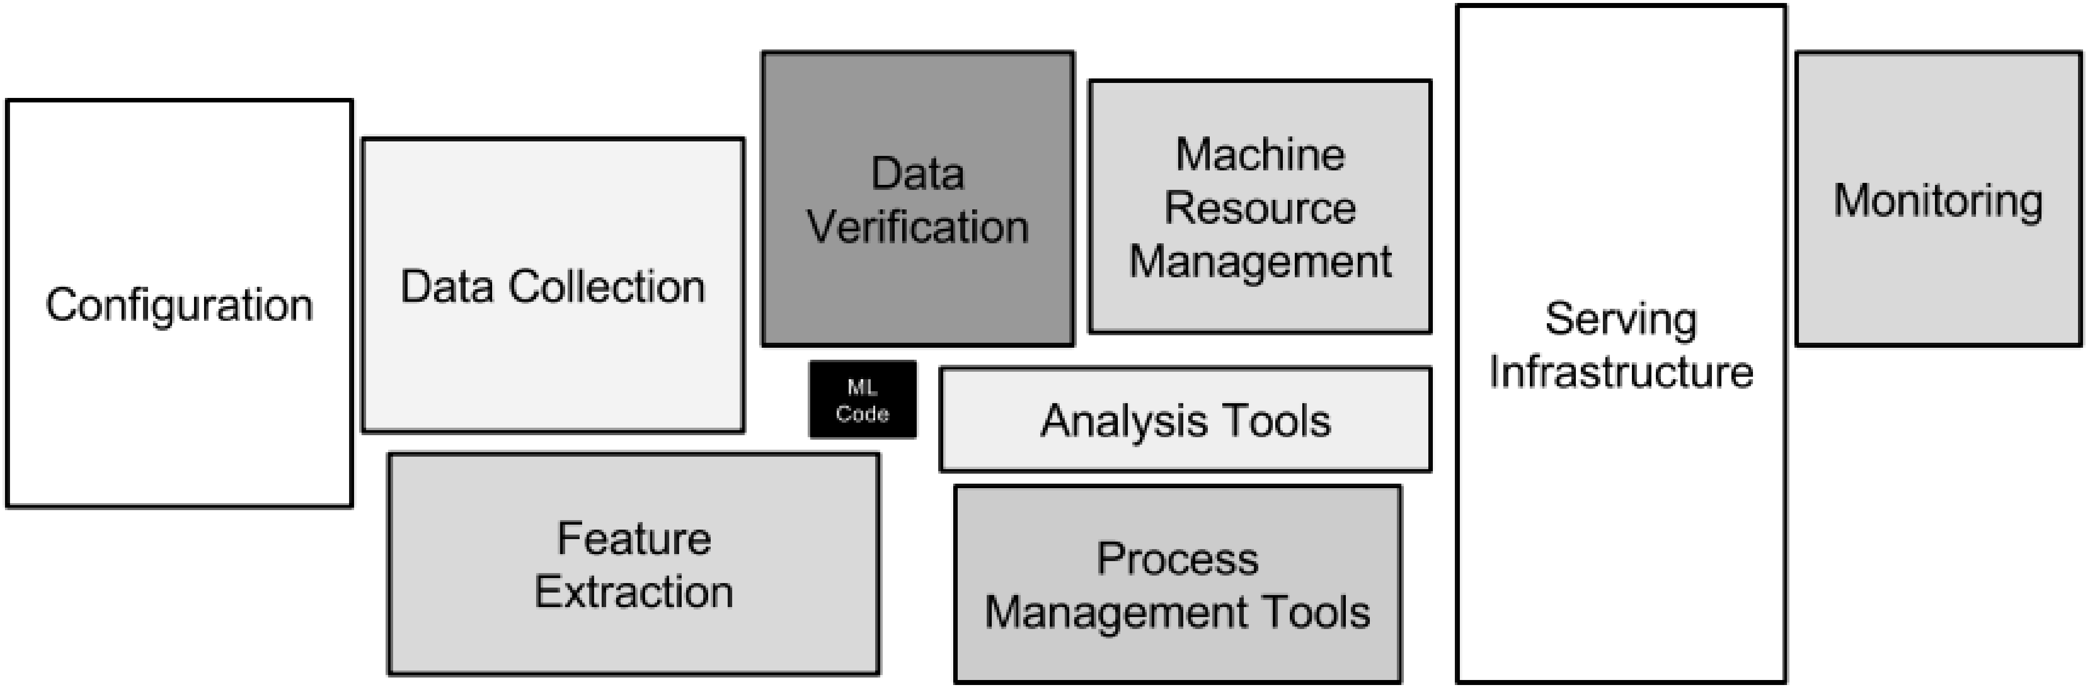

Only a small fraction of real-world ML systems is a composed of ML code, as shown by the small black box in the middle. The required surrounding infrastructure is fast and complex.

D. Sculley, et al. Hidden technical debt in machine learning systems. In Neural Information Processing Systems (NIPS). 2015

How we achieve Salesforce-scale!

Salesforce approach to democratizing AI

Enabling our customers to build models on their own data

The need for platform to ship AI to production

Bridging the communication gap between data scientists and software developers to find common ground and get to production and agility

Critical components of an AI platform

How to build a platform to support agile data science

How metrics drive agility and scale

How to apply agile methodologies to rapidly improve and deploy models

How the Salesforce Einstein Platform Enables Data Scientists

Deploy, monitor and iterate on models in one location

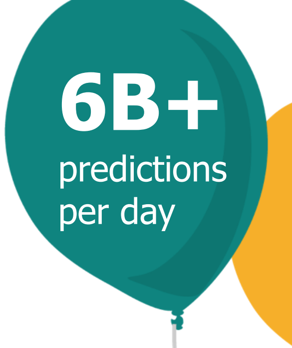

Microservice architecture

Shared feature engineering and modeling services

Customizable model-evaluation & monitoring dashboards

In-platform secured experimentation and exploration

Data Scientists focus their efforts on engineering new features, trying new models and evaluating results

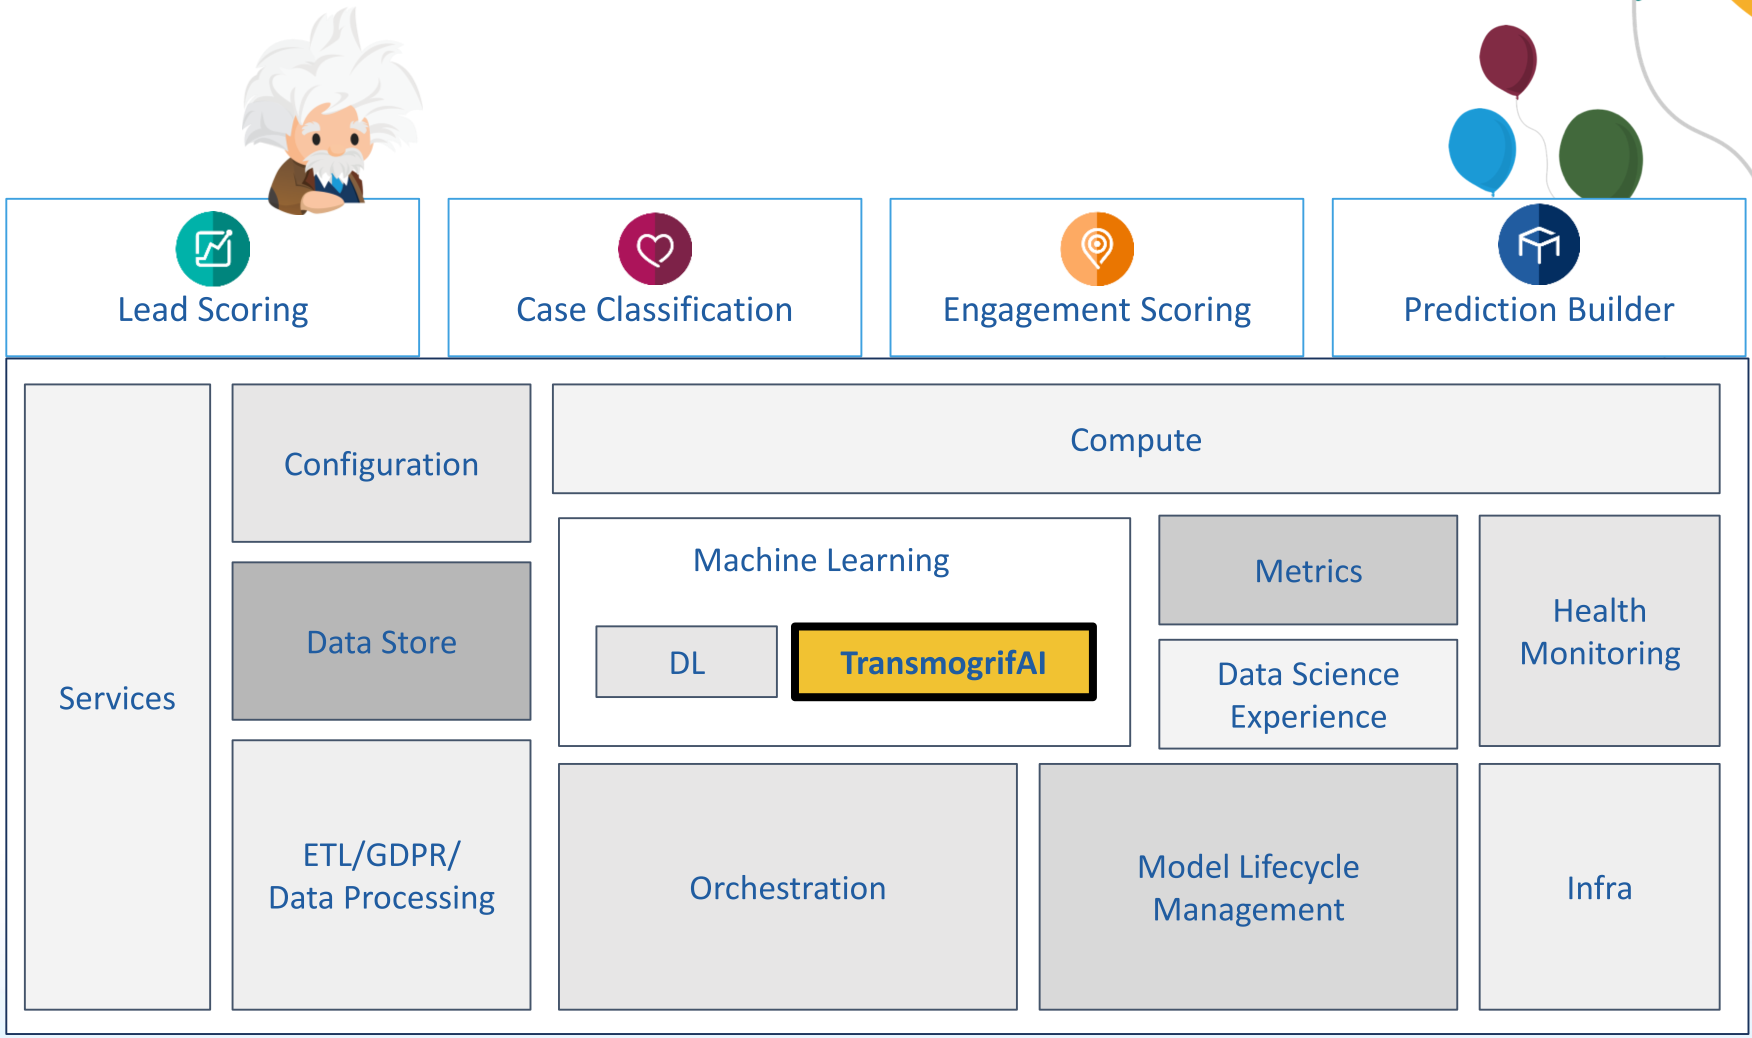

How do we build models? evaluate? reuse?

```
>>> from sklearn import svm
>>> from numpy import loadtxt as l, random as r
>>> clf = svm.SVC()
>>> pls = numpy.loadtxt(”features.data", delimiter=",")
>>> testSet = r.choice(len(pls), int(len(pls)*.7), replace=False)
>>> X,  y = pls[-testSet,:-1], pls[-testSet:,-1]
>>> clf.fit(X,y)
SVC(C=1.0, cache_size=200, class_weight=None,
        coef0=0.0,decision_function_shape=None, degree=3,
        gamma='auto', kernel='rbf', max_iter=-1,
        probability=False, random_state=None, shrinking=True,
        tol=0.001, verbose=False)
>>> clf.score(pls[testSet,:-1],pls[testSet,-1])
 0.88571428571428568
```

Should we try other model forms?
Features?
Kernels or hyperparameters?

How do we make the best decisions for every model in production?


Learn from the mistakes of
others. You can’t live long enough
to make them all yourself
-Eleanor Roosevelt

Introducing TransmogrifAI

Open Sourcing Auto-ML for Structured Data

salesforce

Spark & Scala

Automated feature engineering,
feature selection & model selection

ML abstractions that improve
developer productivity & collaboration

Model explainability to improve
debuggability and transparency

```scala
// Read the Deal data
val dealData = DataReaders.Simple.csvCase[Deal](path = pathToData).readDataset().toDF()

// Extract response and predictor Features
val (isClosed, predictors) = FeatureBuilder.fromDataFrame[RealNN](dealData, response = '

// Automated feature engineering
val featureVector = predictors.transmogrify()

// Automated feature selection
val cleanFeatures = survived.sanityCheck(featureVector, removeBadFeatures = true)

// Automated model selection
val (pred, raw, prob) = BinaryClassificationModelSelector()
  .setInput(isClosed, cleanFeatures).getOutput()

val model = new OpWorkflow().setInputDataset(dealData).setResultFeatures(pred).train()
```

Einstein's New Approach to AI

Democratizing AI for Everyone

Classical Approach → Data Sampling → Feature Selection → Model Selection → Score Calibration → Integrate to Application → Artificial Intelligence

Einstein Auto-ML → AI for CRM

AI for CRM
- Discover
- Predict
- Recommend
- Automate

Data already prepped
Models automatically built
Predictions delivered in context

Repeatable Elements in Machine Learning Pipelines

AutoML for feature engineering

Categorical Variables

NAME	TITLE
Jim Steele	Senior VP
John Gardner	Senior VP
Andy Smith	Vice President
Test User	Vice President
Test User	CEO
Test User	Vice President
Test User	Chairperson
Test User	CEO

Text Fields

DESCRIPTION
A blessing in disguise
Time flies when you're having fun
Alles hat ein Ende, nur die Wurst hat zwei
um den heißen Brei herumreden
We'll cross that bridge when we come to it
You can say that again
Your guess is as good as mine

Numerical Buckets

number of employees
90
16
224
192
335
12
621
72
560
80
24
0
208

Repeatable Elements in Machine Learning Pipelines

AutoML for feature engineering

Text Fields

DESCRIPTION	Word Count
A blessing in disguise	4
Time flies when you're having fun	6
Alles hat ein Ende, nur die Wurst hat zwei	9
um den heißen Brei herumreden	6
We'll cross that bridge when we come to it	7
You can say that again	5
Your guess is as good as mine	7

Repeatable Elements in Machine Learning Pipelines

AutoML for feature engineering

Text Fields		
DESCRIPTION	Word Count	Word Count (no stop words)
A blessing in disguise	4	2
Time flies when you're having fun	6	3
Alles hat ein Ende, nur die Wurst hat zwei	9	4
um den heißen Brei herumreden	6	4
We'll cross that bridge when we come to it	7	3
You can say that again	5	1
Your guess is as good as mine	7	3

Repeatable Elements in Machine Learning Pipelines

AutoML for feature engineering

Text Fields			
DESCRIPTION	Word Count	Word Count (no stop words)	Is English
A blessing in disguise	4	2	1
Time flies when you're having fun	6	3	1
Alles hat ein Ende, nur die Wurst hat zwei	9	4	0
um den heißen Brei herumreden	6	4	0
We'll cross that bridge when we come to it	7	3	1
You can say that again	5	1	1
Your guess is as good as mine	7	3	1

Repeatable Elements in Machine Learning Pipelines

AutoML for feature engineering

Text Fields				
DESCRIPTION	Word Count	Word Count (no stop words)	Is English	Sentiment
A blessing in disguise	4	2	1	1
Time flies when you're having fun	6	3	1	1
Alles hat ein Ende, nur die Wurst hat zwei	9	4	0	0
um den heißen Brei herumreden	6	4	0	-1
We'll cross that bridge when we come to it	7	3	1	0
You can say that again	5	1	1	0
Your guess is as good as mine	7	3	1	0

Einstein's New Approach to AI

Democratizing AI for Everyone

Classical Approach → Data Sampling → Feature Selection → **Model Selection** → Score Calibration → Integrate to Application → Artificial Intelligence

Einstein Auto-ML → AI for CRM
Discover
Predict
Recommend
Automate

Data already prepped
Models automatically built
Predictions delivered in context

A tournament of models!

CUSTOMER A

Customer ID	1	2	3	4	5
Age	22	30	45	23	60
Age Group	A	A	B	A	C
Gender	M	F	F	M	M
Valid Address	Y	N	N	Y	Y

MODEL GENERATION

Model 1

Model 2

Model 3

Model 4

...

MODEL TESTING

Model 1
83% accuracy

Model 2
91% accuracy

Model 3
73% accuracy

Model 4
89% accuracy

...

A tournament of models!

CUSTOMER A

Customer ID	1	2	3	4	5
Age	22	30	45	23	60
Age Group	A	A	B	A	C
Gender	M	F	F	M	M
Valid Address	Y	N	N	Y	Y

CUSTOMER B

Customer ID	1	2	3	4	5
Age	34	22	66	58	41
Age Group	B	A	C	C	B
Gender	M	M	F	M	F
Valid Address	Y	Y	N	N	Y

MODEL GENERATION

Model 1

Model 2

Model 3

Model 4

...

MODEL TESTING

Model 1

Model 2
Customer A

Model 3

Model 4
Customer B

...

salesforce

How the Salesforce Einstein Platform Enables Data Scientists

Deploy, monitor and iterate on models in one location

6B+ predictions per day

Microservice architecture

Shared feature engineering and modeling services

Customizable model-evaluation & monitoring dashboards

In-platform secured experimentation and exploration

Data Scientists focus their efforts on engineering new features, trying new models and evaluating results

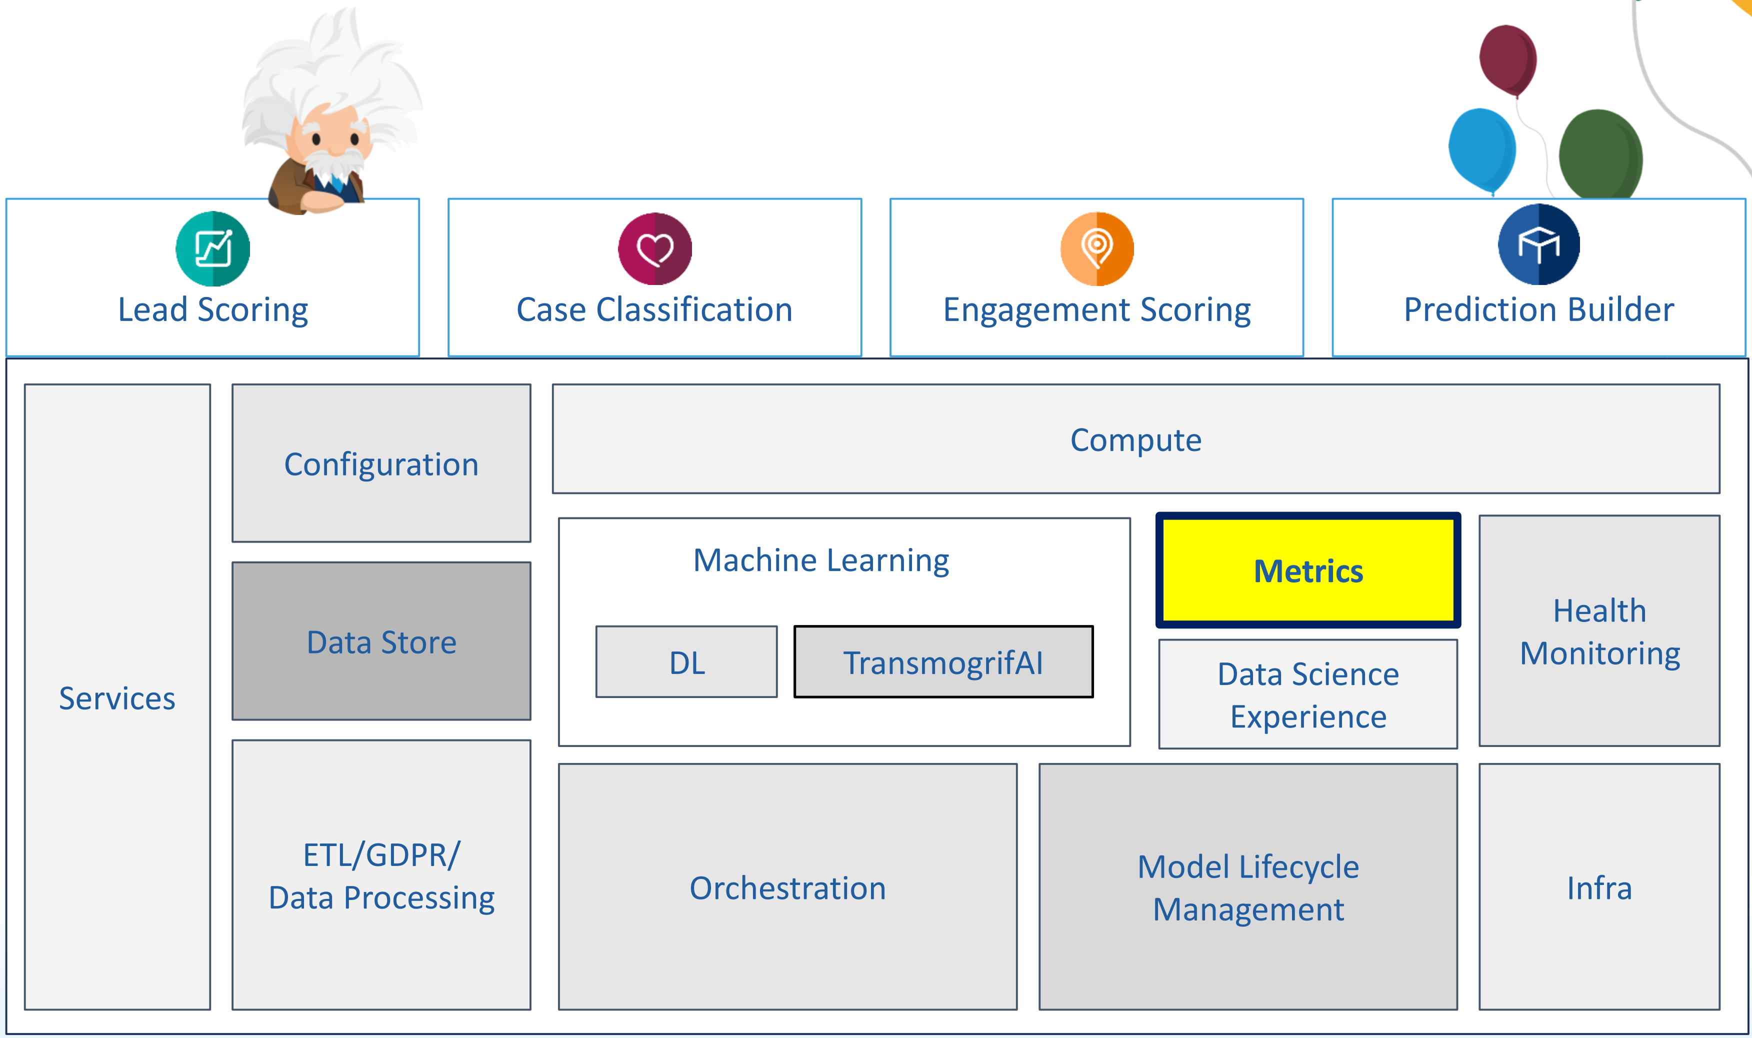

Monitoring your AI's health like any other app component

Pipelines, Model Performance, Scores – Invest your time where it is needed!

105,874
Scores Written Per Hour(1 day moving avg)

Total Number of Scores Written Per Hour

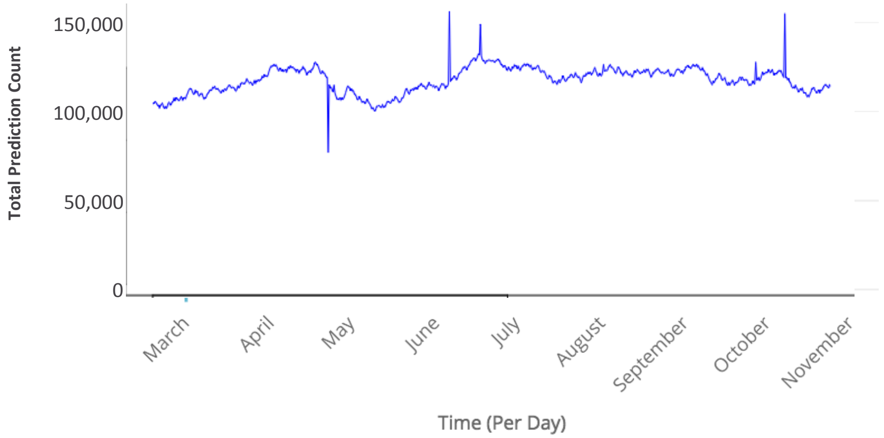

Total Number of Scores Written Per Week

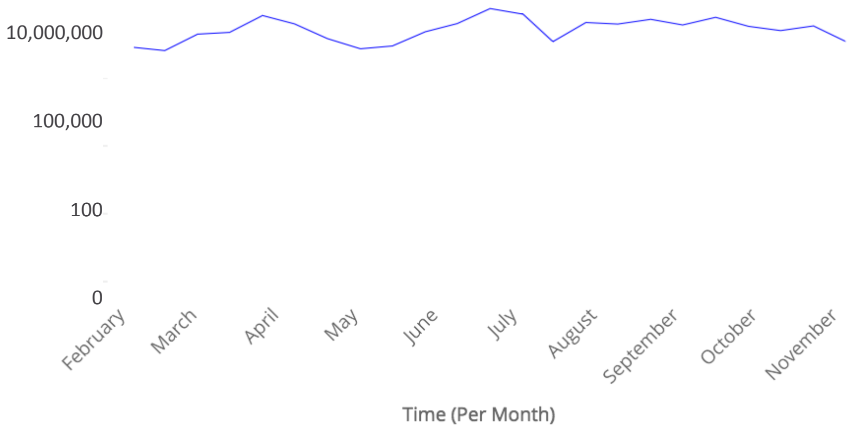

0.86
Evaluation auROC

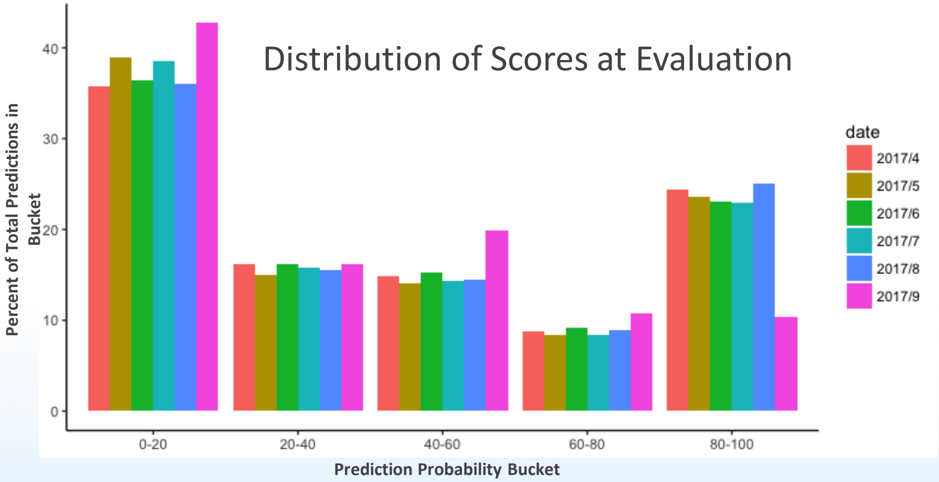

Model Performance at Evaluation

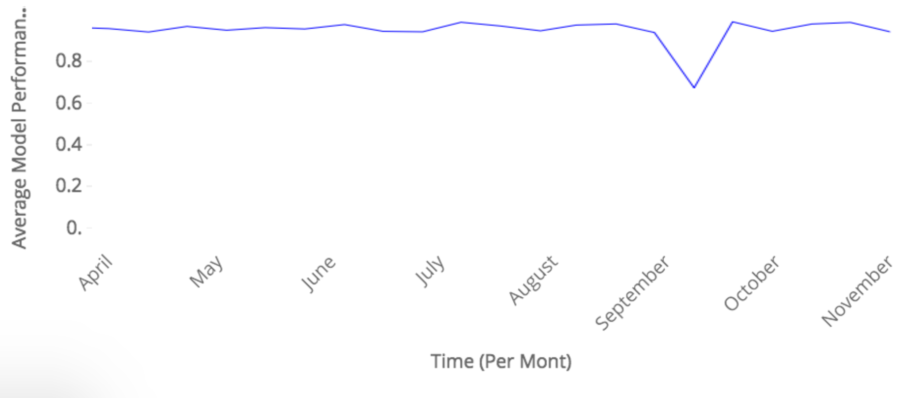

Sample Dashboard on Simulated Data

How we achieve Salesforce-scale!

Salesforce approach to democratizing AI

Enabling our customers to build models on their own data

The need for platform to ship AI to production

Bridging the communication gap between data scientists and software developers to find common ground and get to production and agility

Critical components of an AI platform

How to build a platform to support agile data science

How metrics drive agility and scale

How to apply agile methodologies to rapidly improve and deploy models

What happens after you deploy?

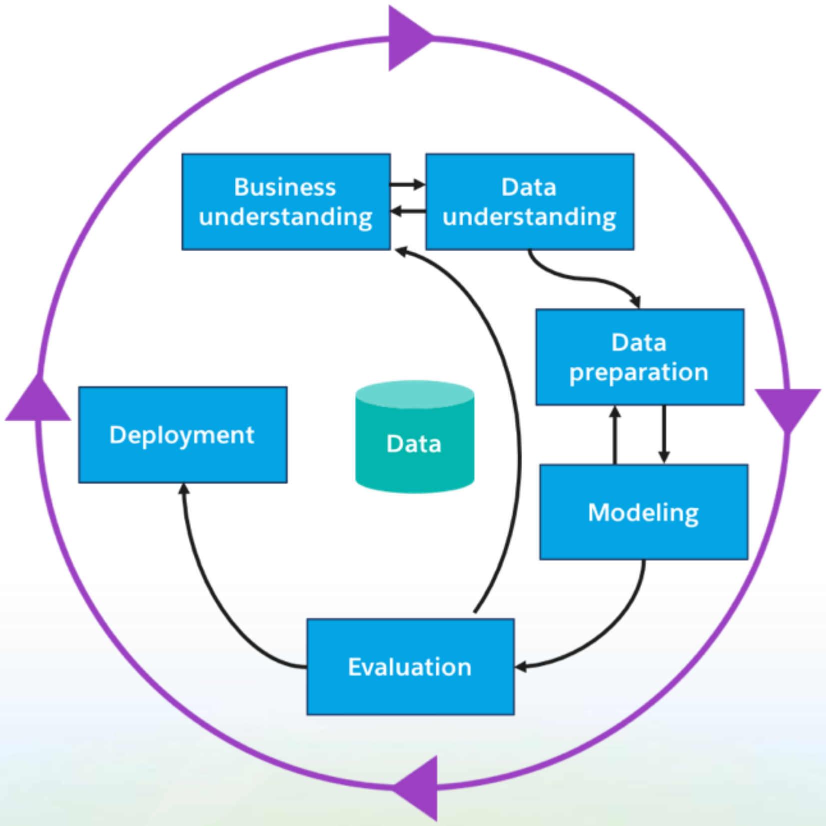


HOW TO BUILD A MINIMUM VIABLE PRODUCT
NOT LIKE THIS
1
2
3
4
LIKE THIS
1
2
3
4
5
image by blog.fastmonkeys.com original idea: spotify product team

Focusing on MVP with Agile Processes

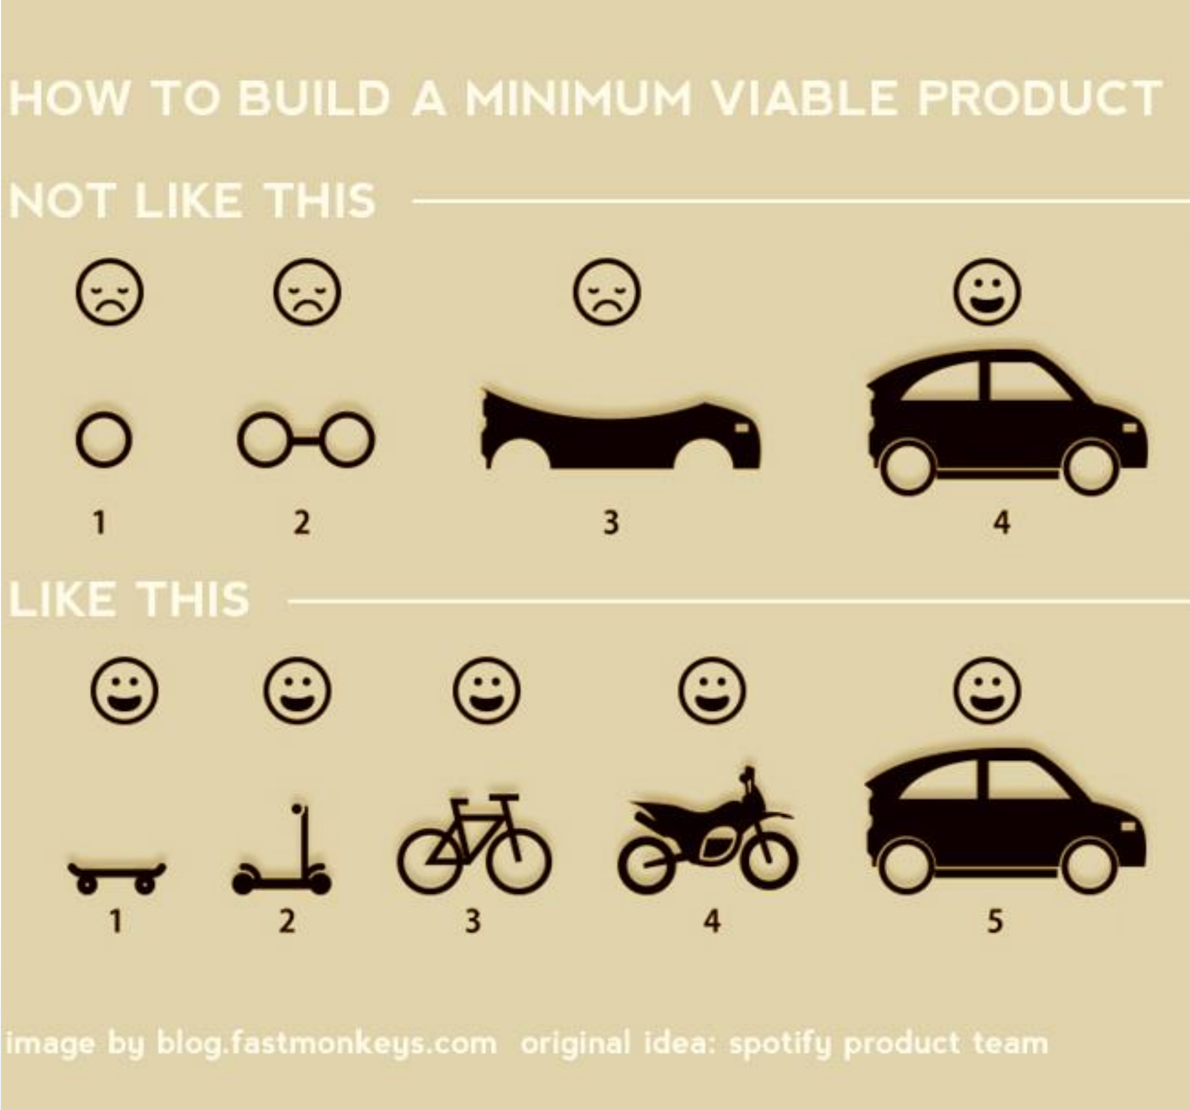

Prioritized Backlog

- ☐ As a rider I want a comfortable seat so I can ride a longer time
- ☐ As a rider I want a self-powered vehicle so I can go further
- ☐ As a rider I want protection because I get hit by bugs
- ☐ As a rider I want shelter because I get wet when it rains
- ☐ As a rider I want to go faster so I can shorten my travel time

Focusing on MVP with Agile Processes

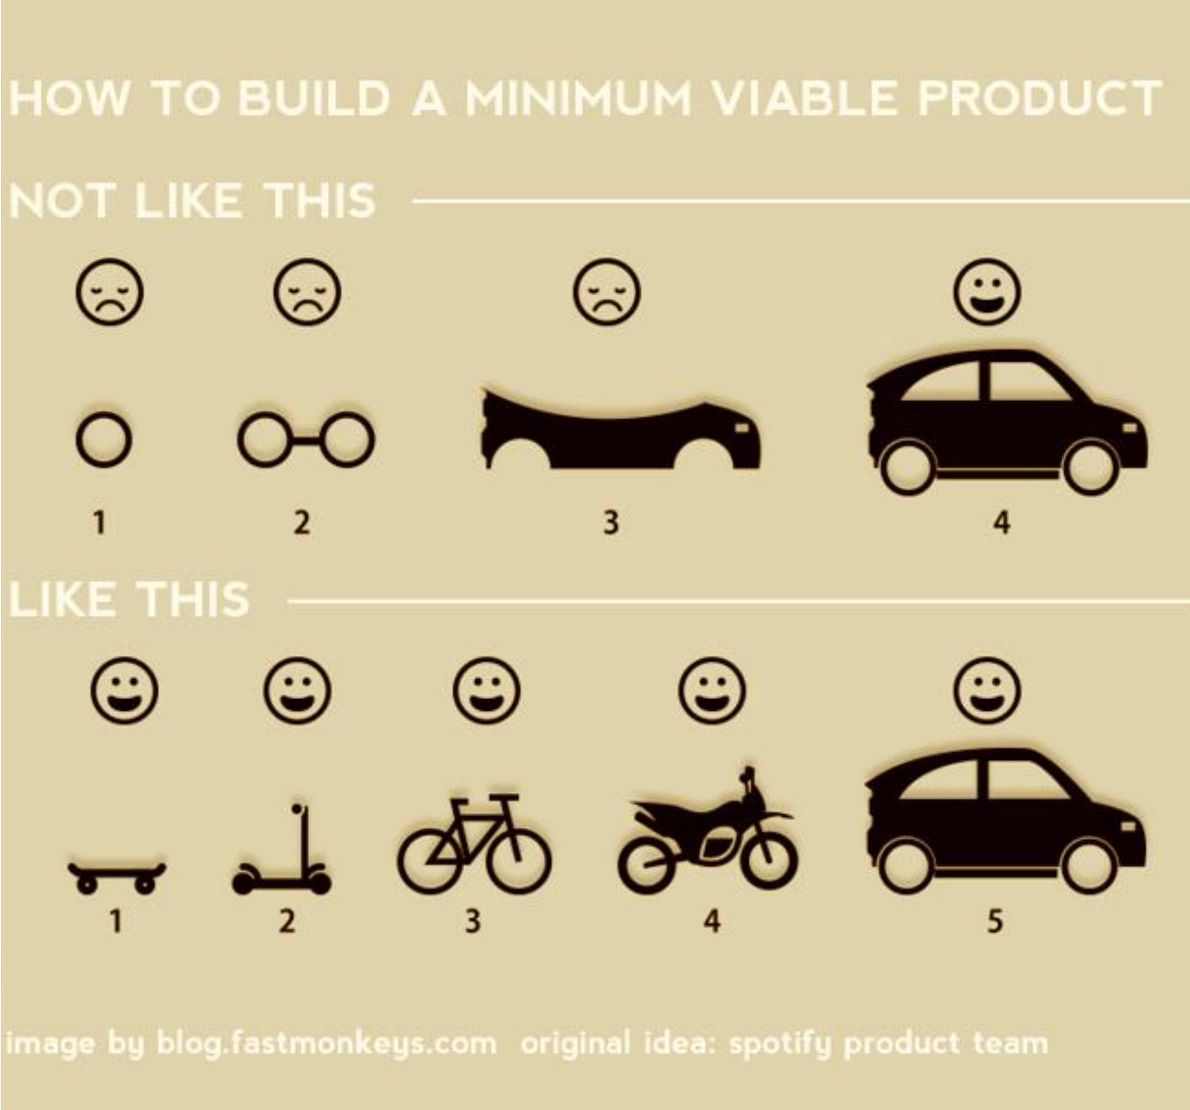

Prioritized Backlog

- ☐ As a rider I want a comfortable seat so I can ride a longer time
- ☐ As a rider I want a self-powered vehicle so I can go further
- ☐ As a rider I want protection because I get hit by bugs
- ☐ As a rider I want shelter because I get wet when it rains
- ☐ As a rider I want to go faster so I can shorten my travel time

Week 1 | Week 2 | Week 3 | Week 4 | Week 5 | Week 6

Sprint 1 | Sprint 2 | Sprint 3

Focusing on MVP with Agile Processes

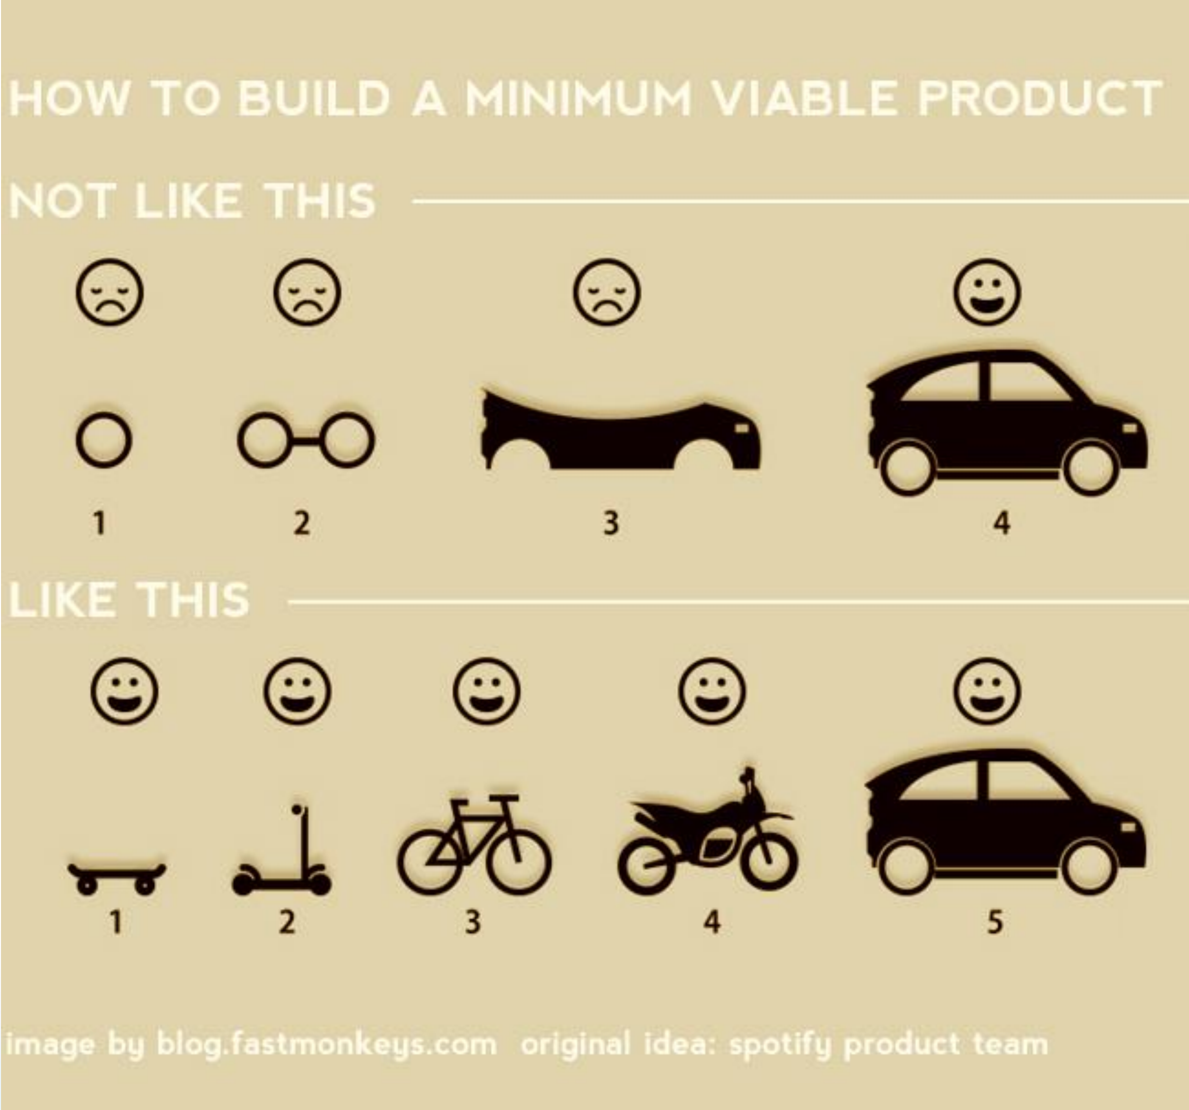

Prioritized Backlog

- ☐ As a rider I want protection because I get hit by bugs
- ☐ As a rider I want shelter because I get wet when it rains
- ☐ As a rider I want to go faster so I can shorten my travel time

Week 1 | Week 2 | Week 3 | Week 4 | Week 5 | Week 6

Sprint 1 | Sprint 2 | Sprint 3

- ☐ As a rider I want a comfortable seat so I can ride longer distances
- ☐ As a rider I want a self-powered vehicle so I can go further

Focusing on MVP with Agile Processes

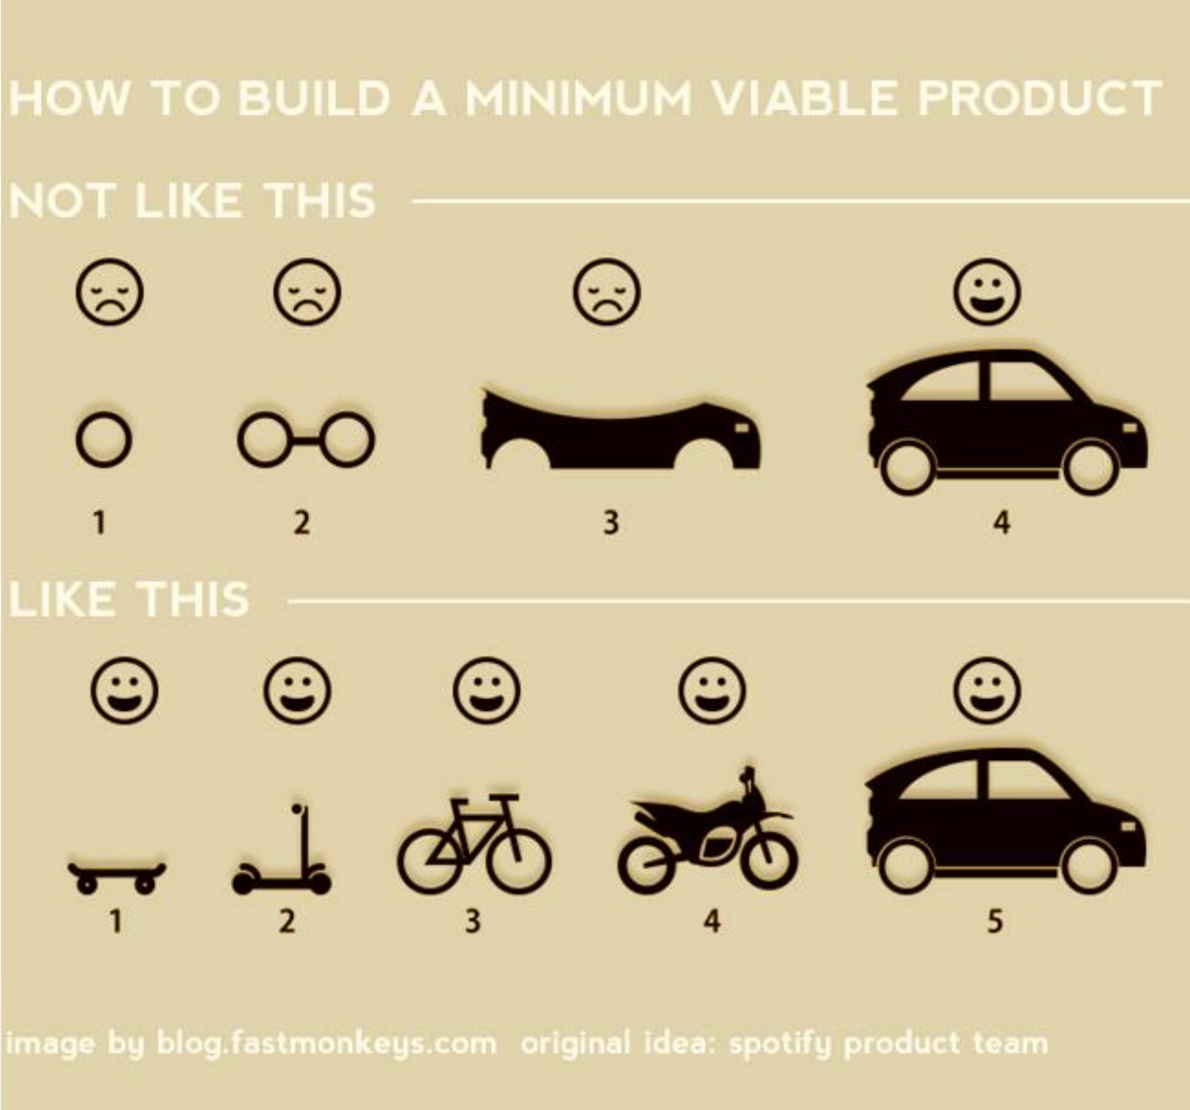

Prioritized Backlog

- ☐ As a rider I want more room because I travel with 3 people
- ☐ As a rider I want protection because I get hit by bugs
- ☐ As a rider I want shelter because I get wet when it rains
- ☐ As a rider I want to go faster so I can shorten my travel time

Week 1 | Week 2 | Week 3 | Week 4 | Week 5 | Week 6

Sprint 1 | Sprint 2 | Sprint 3

- ☐ As a rider I want a comfortable seat so I can ride longer distances
- ☐ As a rider I want a self-powered vehicle so I can go further

Focusing on MVP with Agile Processes

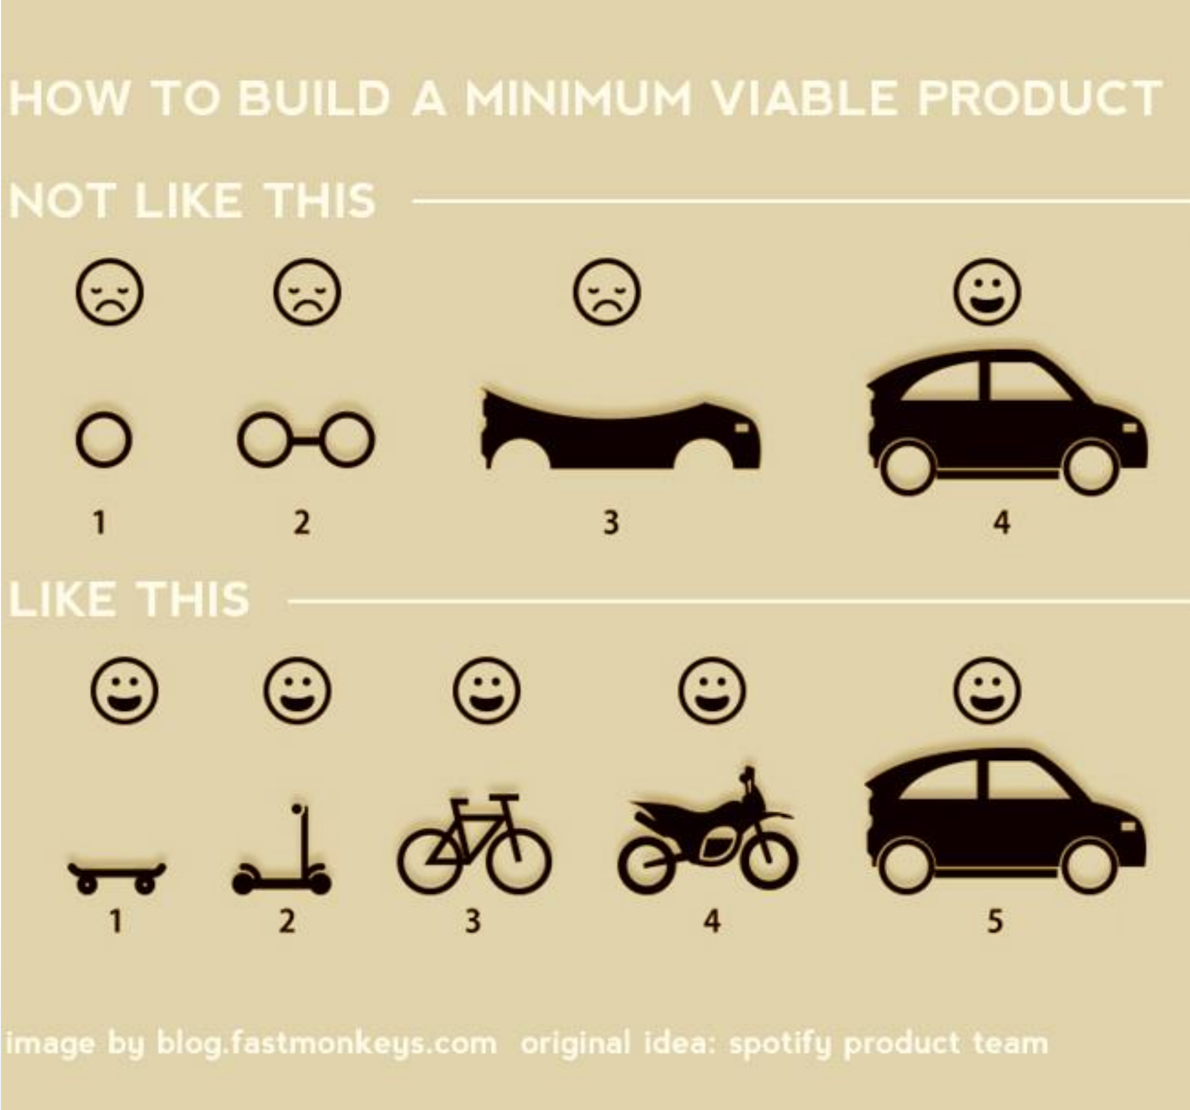

<u>Prioritized Backlog</u>

- ☐ As a rider I want more room because I travel with 3 people
- ☐ As a rider I want to go faster so I can shorten my travel time

Week 1 | Week 2 | Week 3 | Week 4 | Week 5 | Week 6

Sprint 1	Sprint 2	Sprint 3
☐ As a rider I want a comfortable seat so I can ride longer distances	☐ As a rider I want protection because I get hit by bugs	
☐ As a rider I want a self-powered vehicle so I can go further	☐ As a rider I want shelter because I get wet when it rains	

Focusing on MVP with Agile Processes

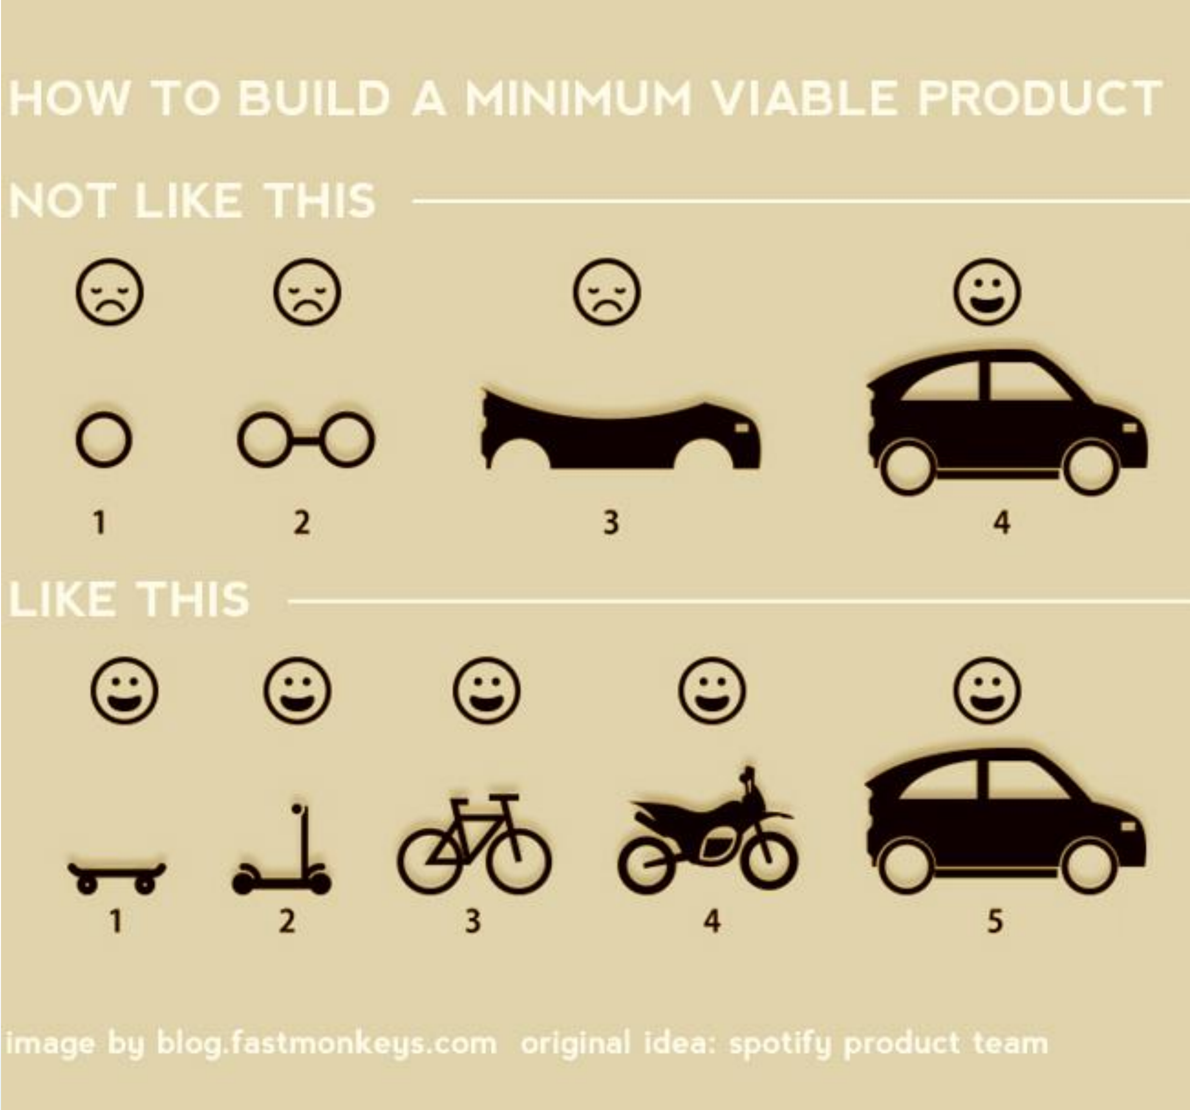

Prioritized Backlog

☐ As a rider I want to go faster so I can shorten my travel time

Week 1 › Week 2 › Week 3 › Week 4 › Week 5 › Week 6

Sprint 1	Sprint 2	Sprint 3
☐ As a rider I want a comfortable seat so I can ride longer distances	☐ As a rider I want protection because I get hit by bugs	☐ As a rider I want more room because I travel with 3 people
☐ As a rider I want a self-powered vehicle so I can go further	☐ As a rider I want shelter because I get wet when it rains	

Focusing on MVP with Agile Processes

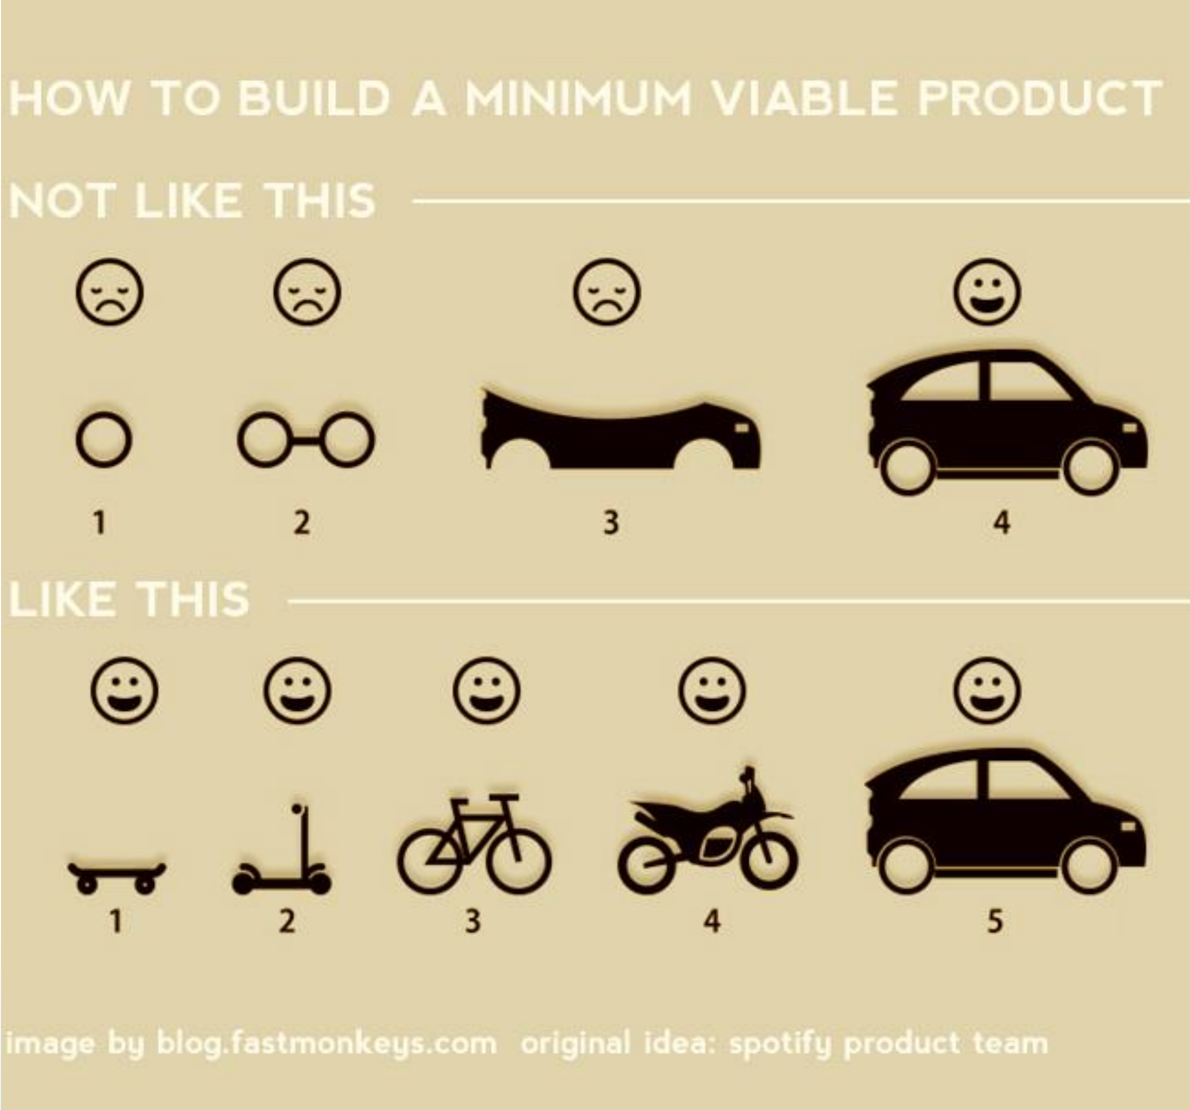

Prioritized Backlog

- [ ] As a rider I want to go faster so I can shorten my travel time

Week 1 | Week 2 | Week 3 | Week 4 | Week 5 | Week 6

Sprint 1	Sprint 2	Sprint 3
☐ As a rider I want a comfortable seat so I can ride longer distances	☐ As a rider I want protection because I get hit by bugs	☐ As a rider I want more room because I travel with 3 people
☐ As a rider I want a self-powered vehicle so I can go further	☐ As a rider I want shelter because I get wet when it rains	

Iterative, Agile, MVP, PSPI, User Stories
WTH does this have to do with Models?

A data scientist's view of the journey to building models

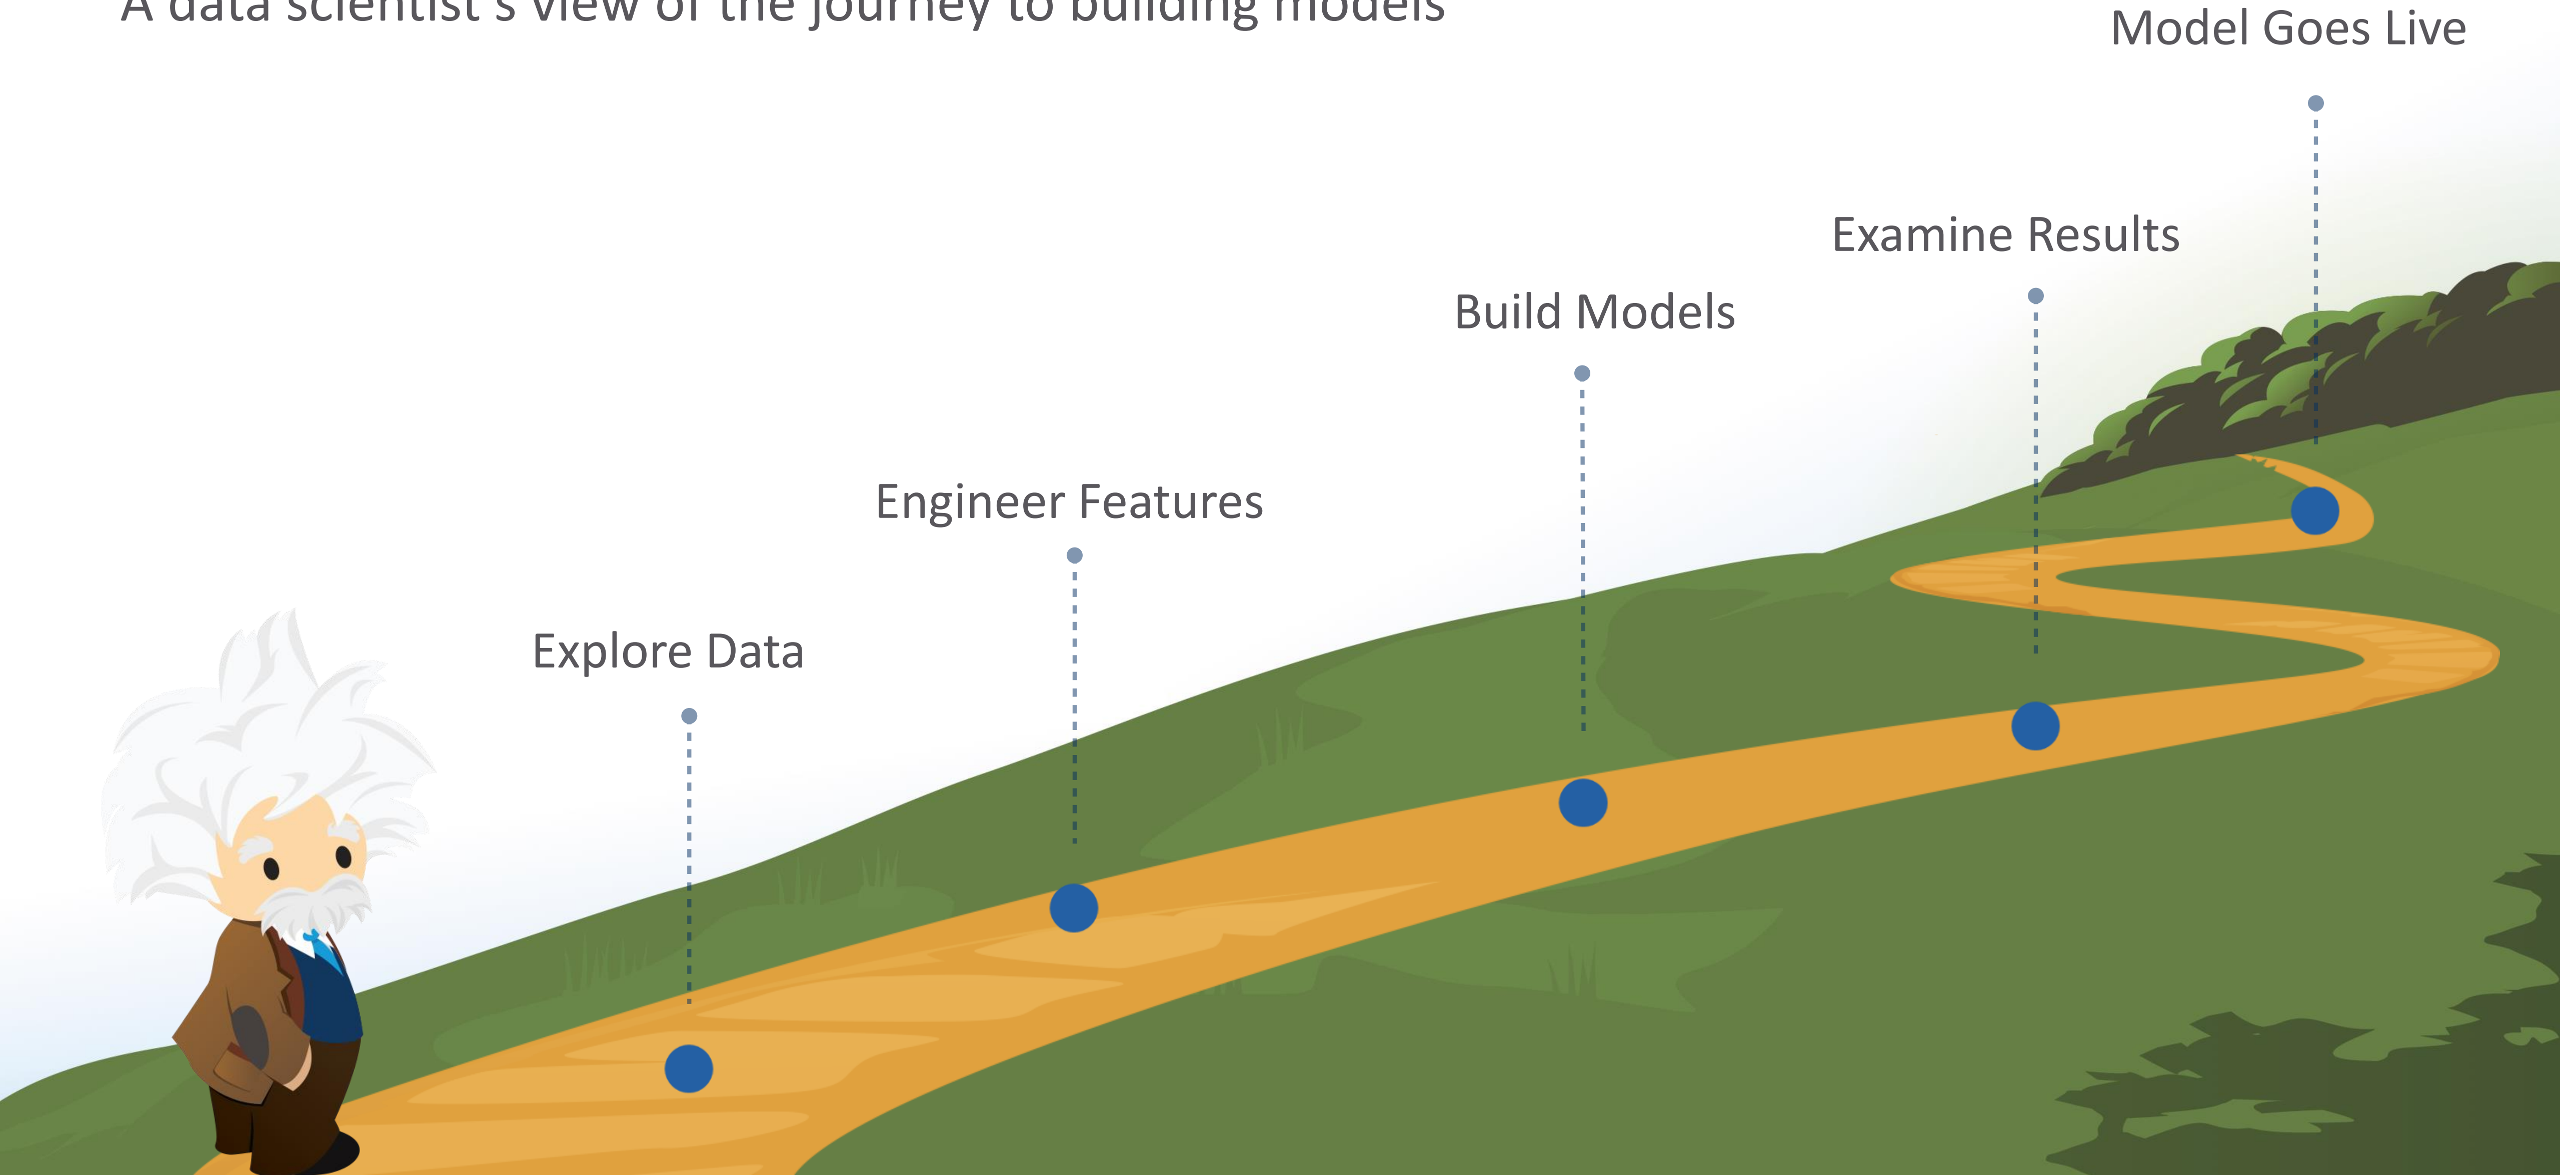

Iterative, Agile, MVP, PSPI, User Stories WTH does this have to do with Models?

Endless choices for ways to improve!

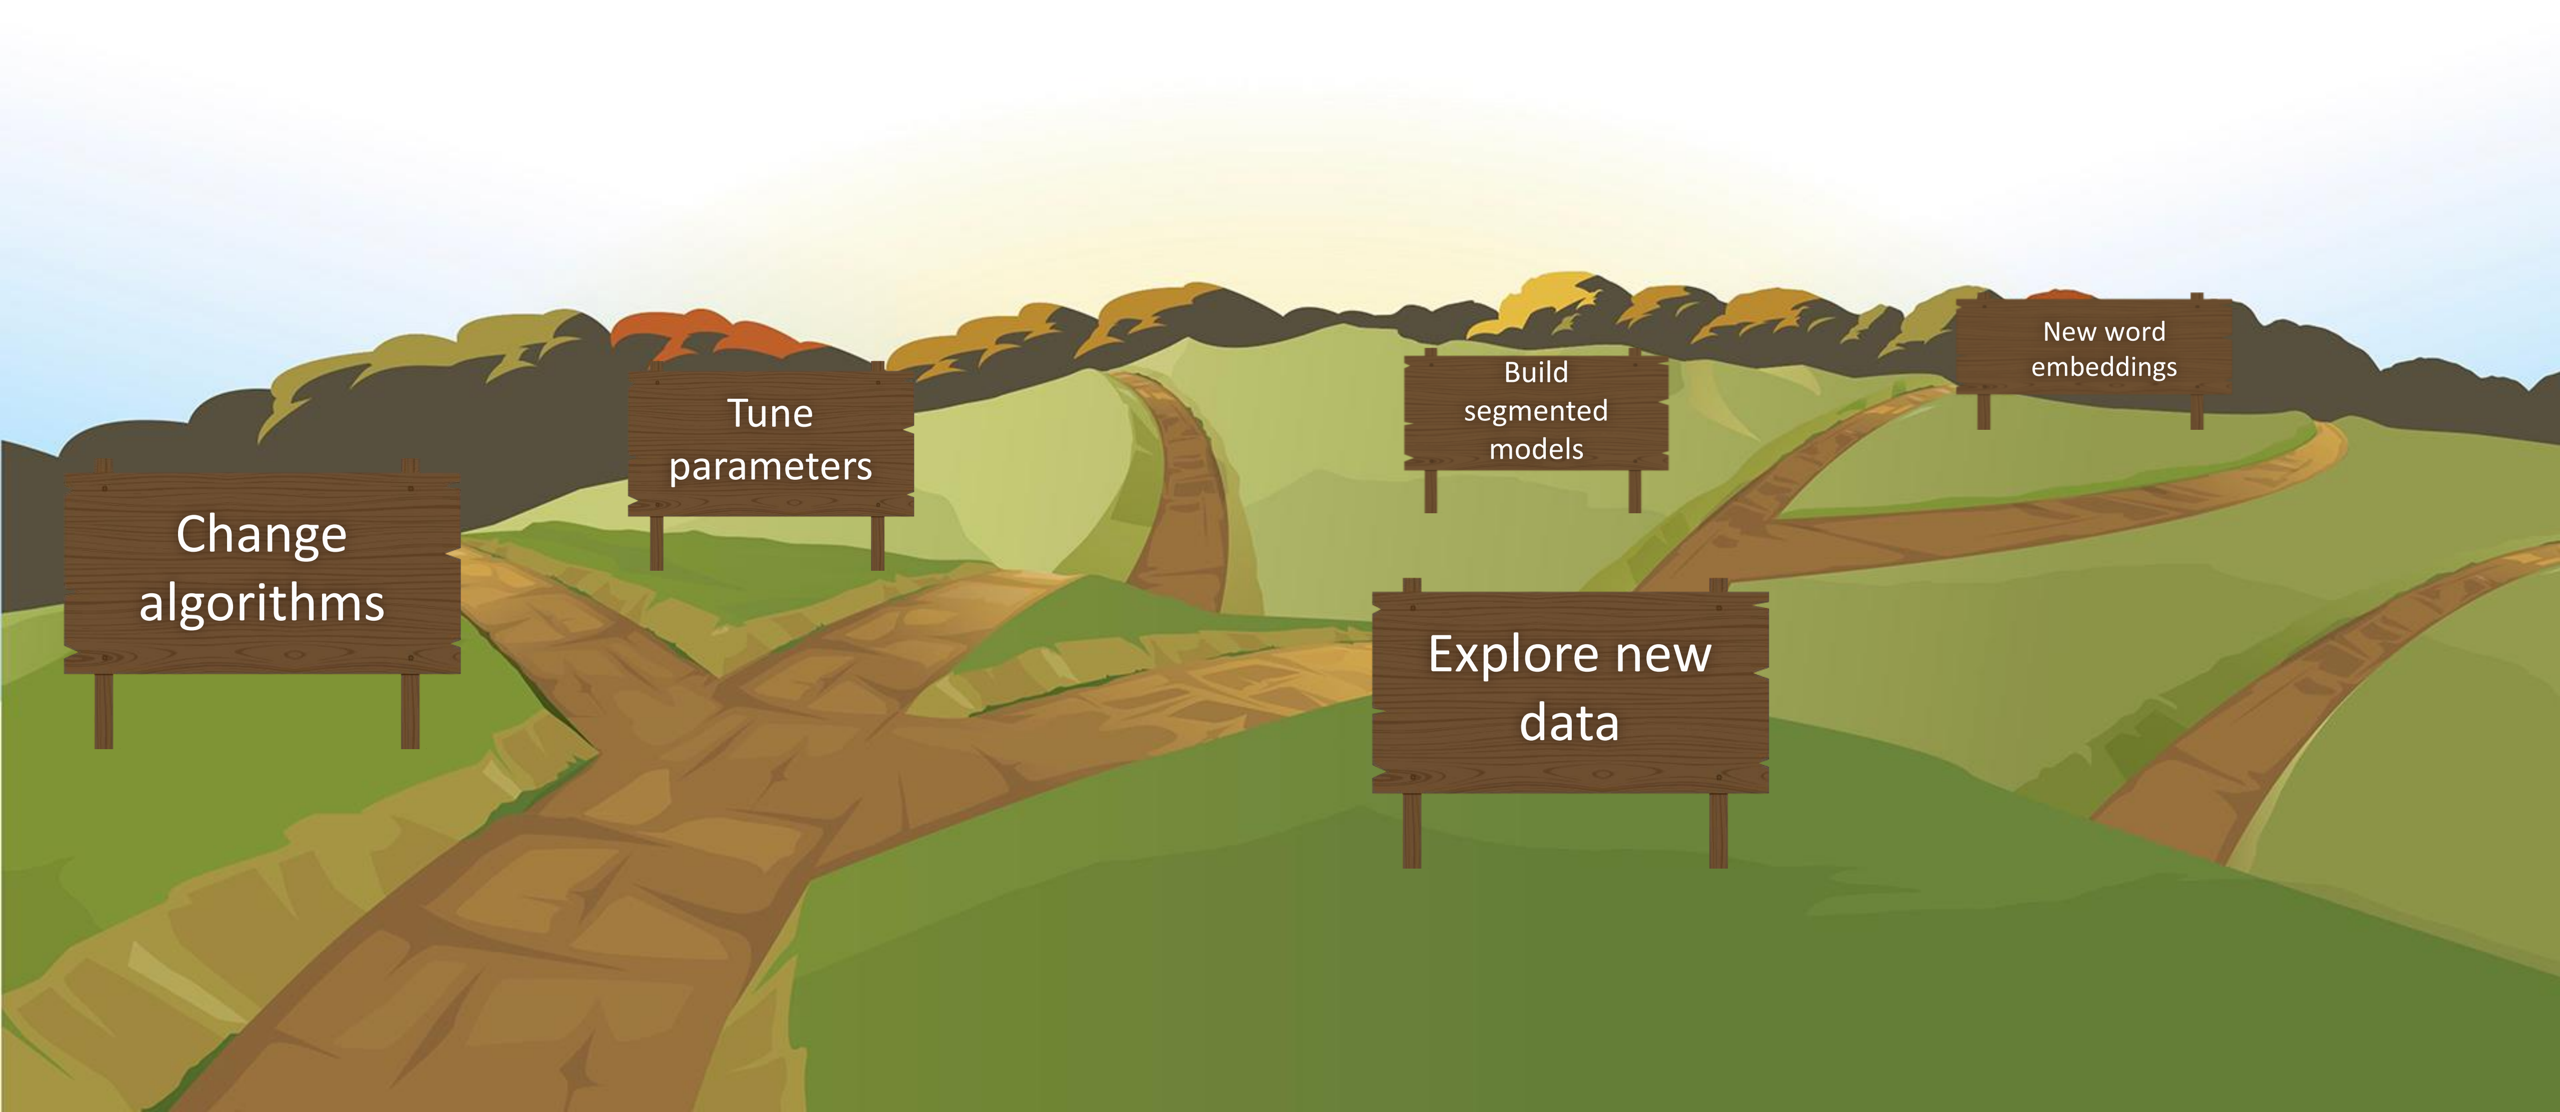

“Le mieux est l'ennemi du bien”
(The best is the enemy of the good)

Voltaire

Invest your time where it is needed!

Pipelines, Model Performance, Scores – Monitors!

105,874
Scores Written Per Hour(1 day moving avg)

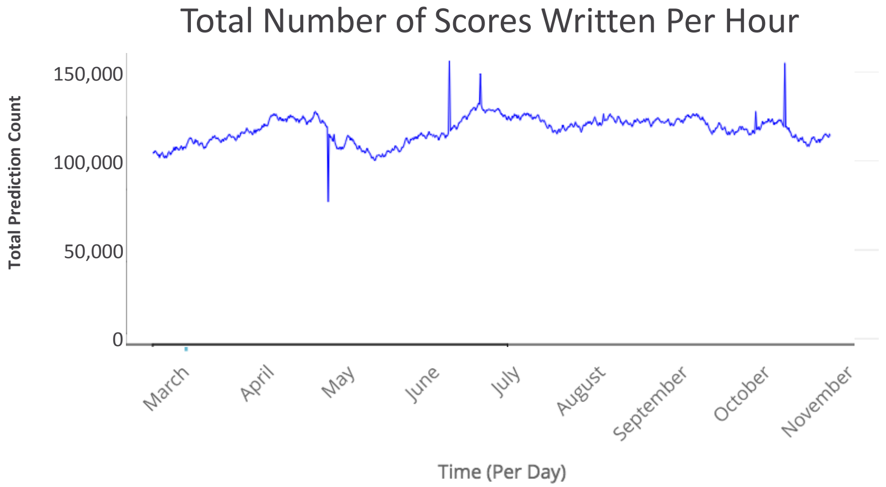

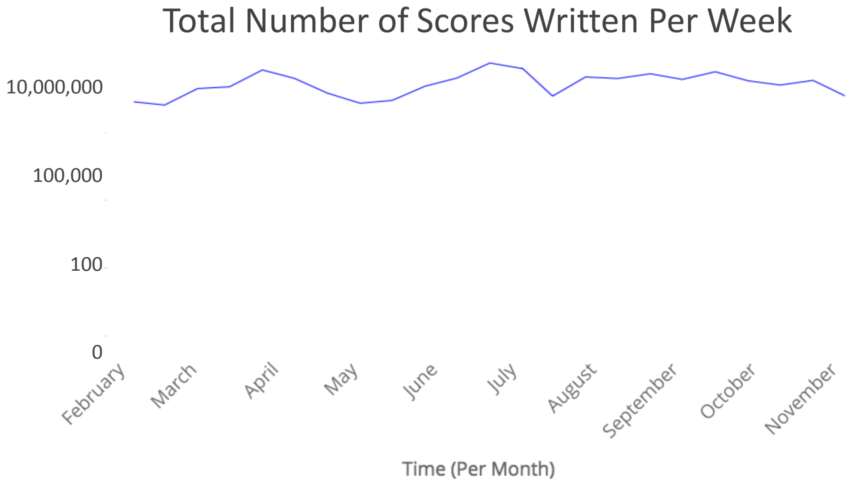

0.86
Evaluation auROC

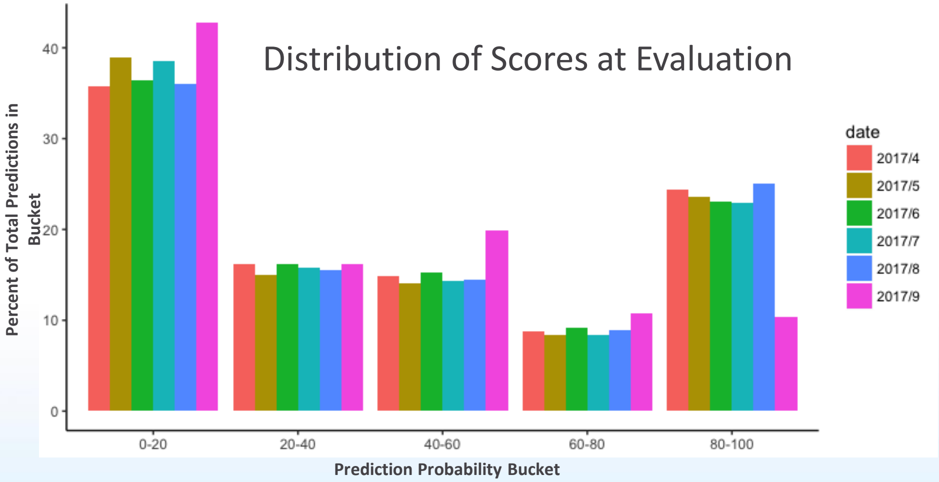

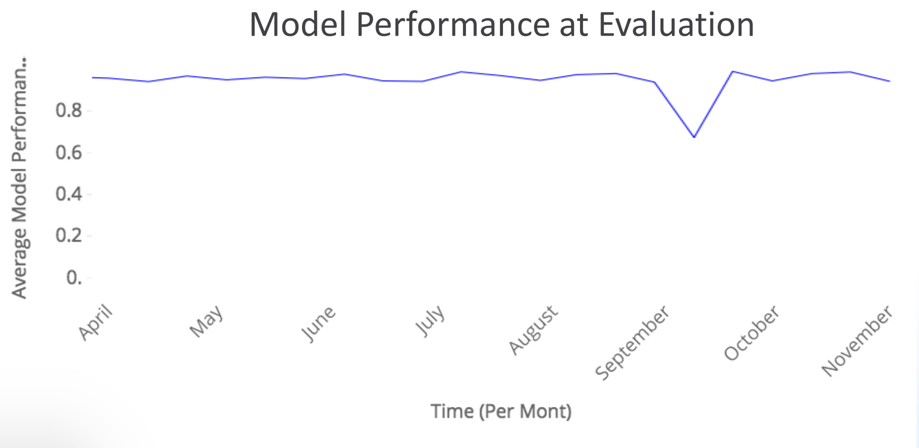

Sample Dashboard on Simulated Data

Remember Our Scale at Salesforce

Customer #1

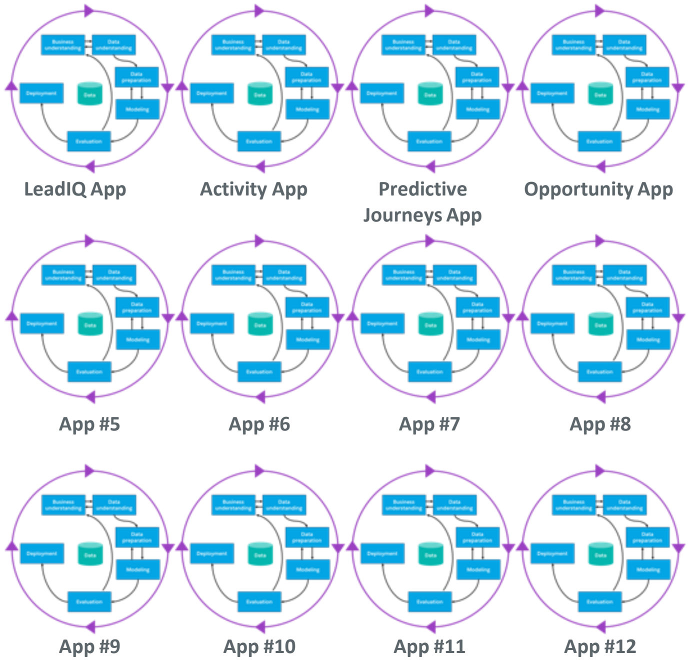

Customer #2

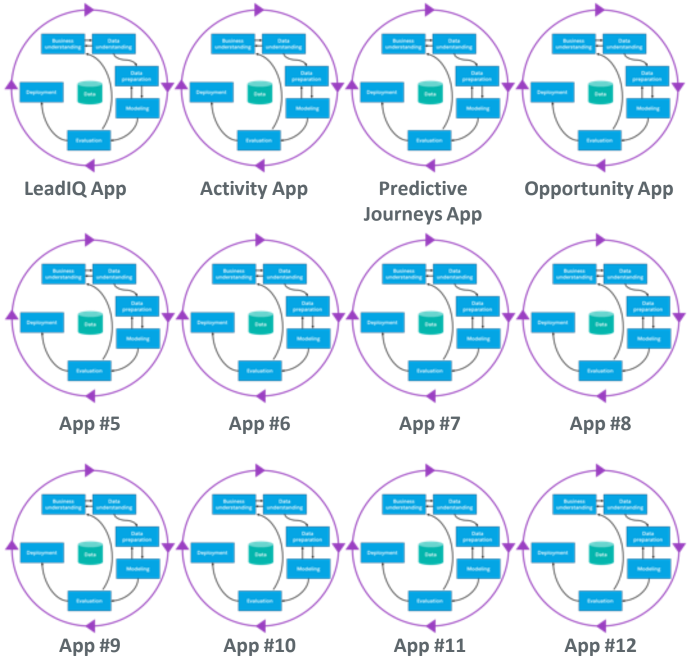

Customer #3

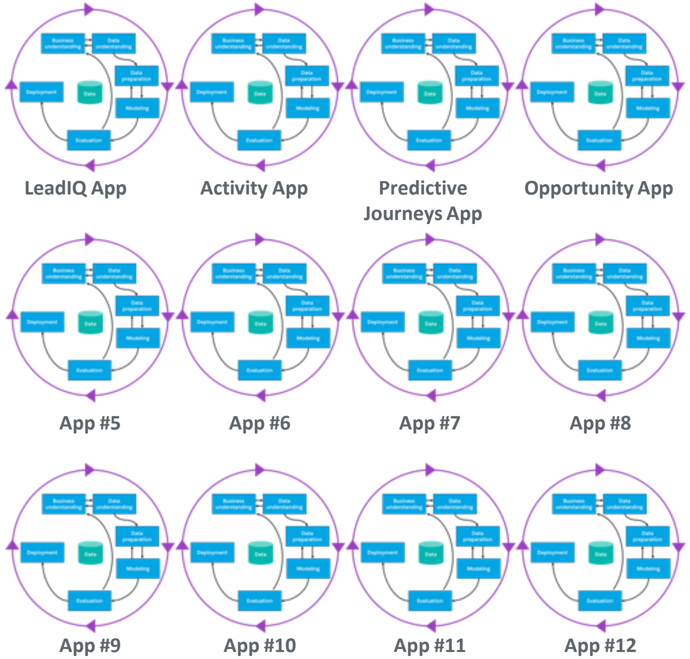

150,000 customers

WHERE DO WE SPEND OUR TIME?

WHERE DO WE SPEND OUR TIME?

Deploy Monitors, Monitor, Repeat!

134 Models in Production	**215** Models Trained (curr.month)	**98.51%** Models with Above Chance Performance

8 Experiments Run this Week	**35,573,664** Predictions Written Per Day (7 day avg)

Sample Dashboard on Simulated Data

What is a sprint? What is a story? What is an investigation? How can agile work in Data Science?

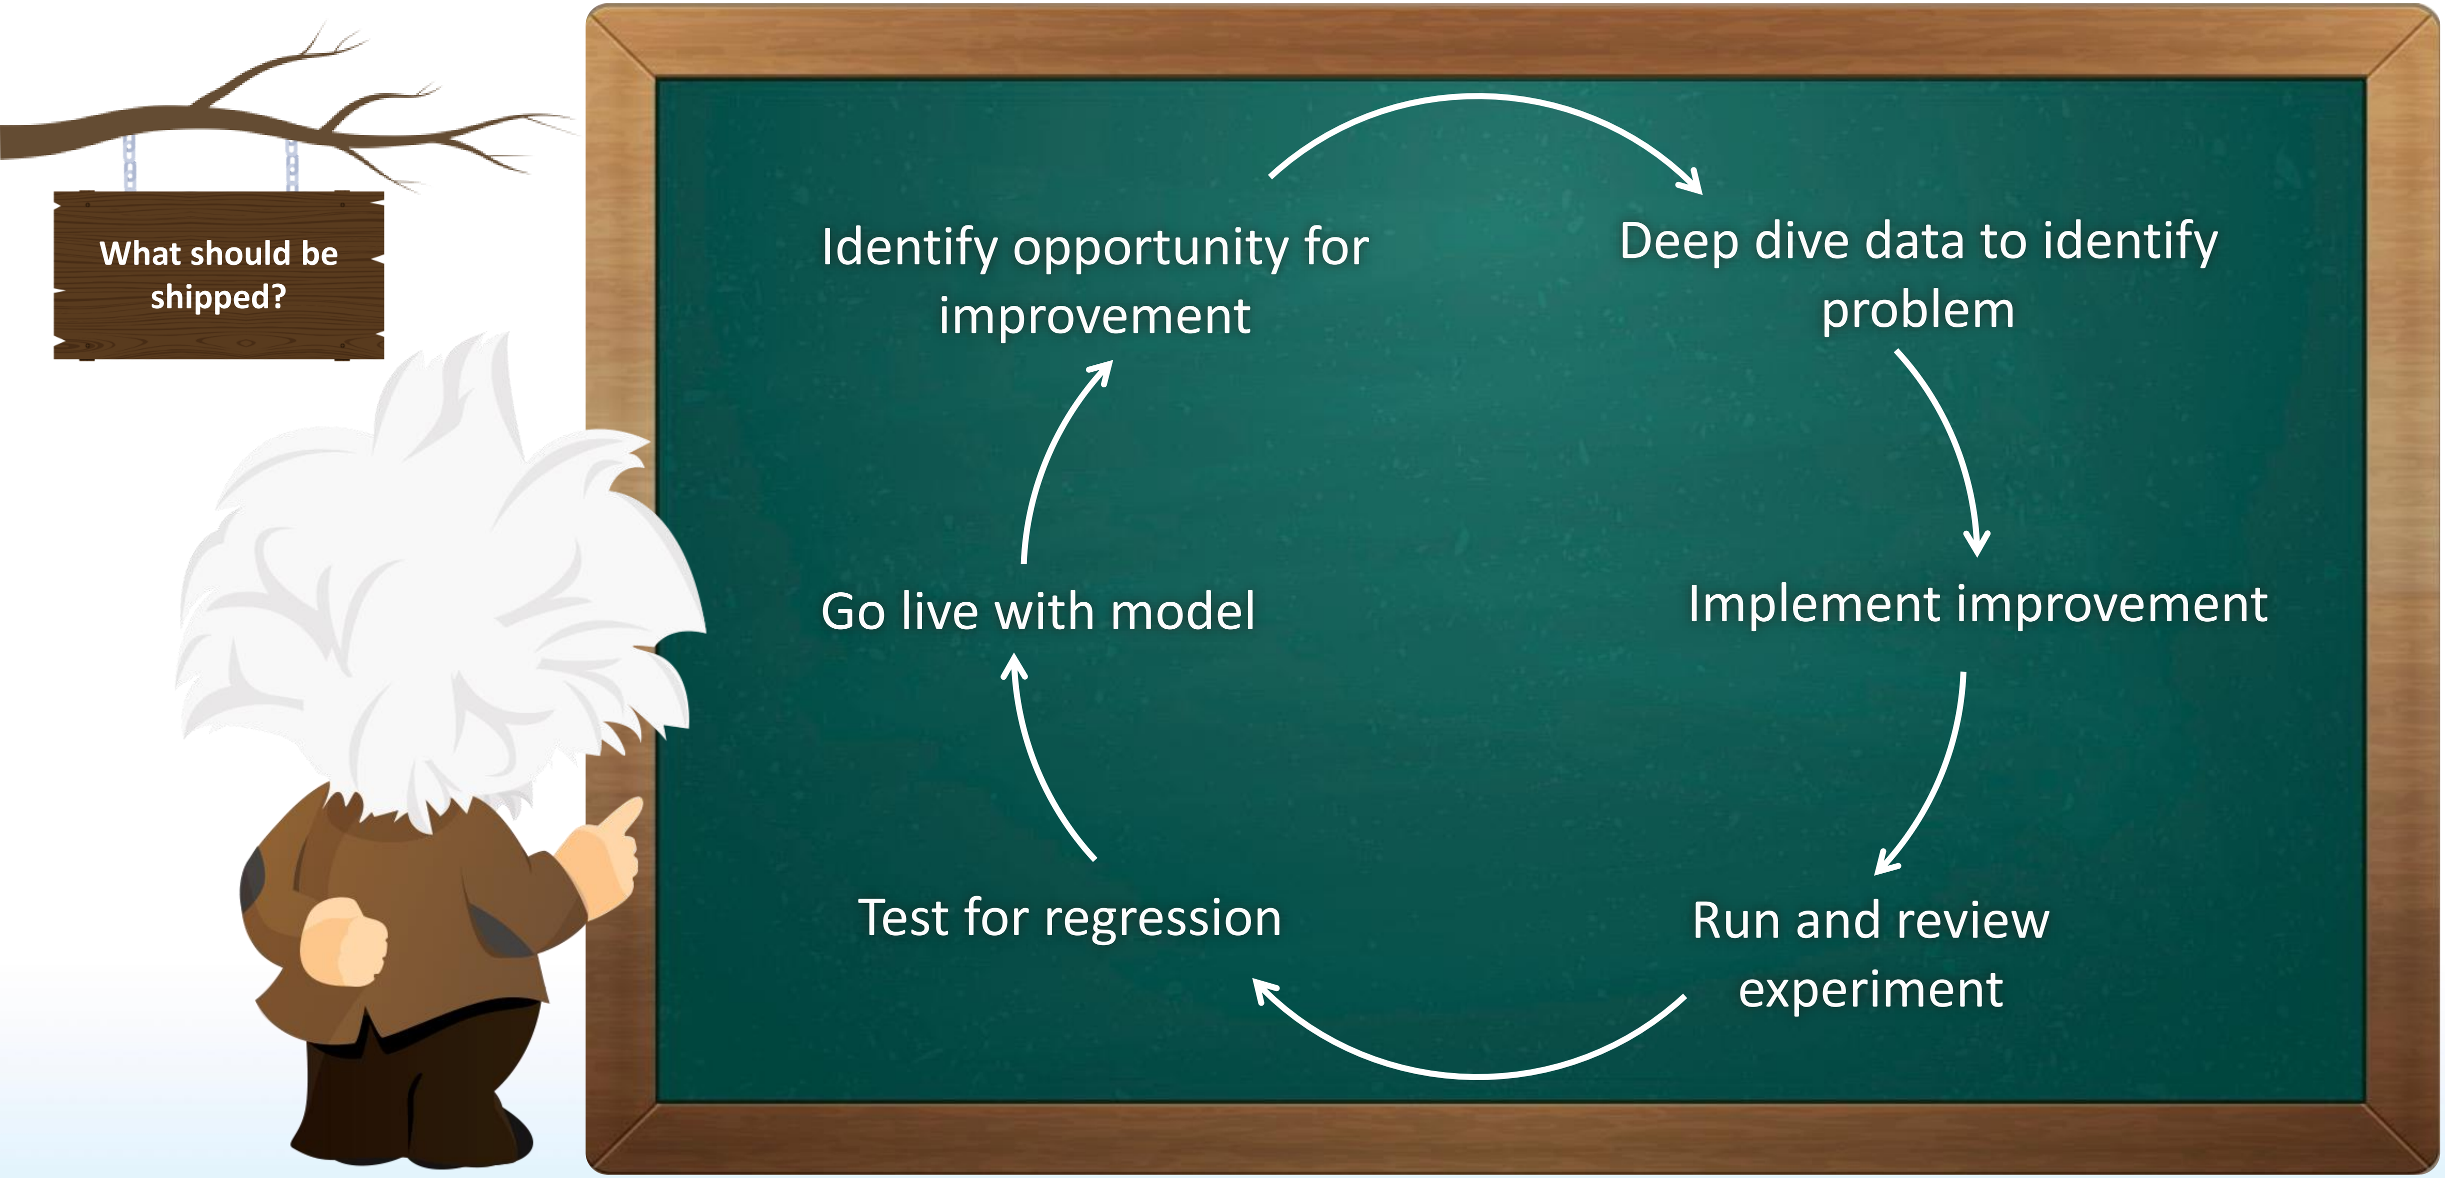

Creating tickets, user stories with clear acceptance criteria

Investigation:

Drop in auROC for model XYZ

AC: Deep dive of model and recent dataset, identify source of issues, create stories for follow-up fixes

User Story:

Add lead data object into model

AC: Engineer features from lead object, add to existing model, provide metrics to assess go-live

SPIKE:

Segmented models

AC: Design doc reviewed and approved by design board, follow-up stories created

User Story:

Identify new hyperparameters settings for random forest

AC: Provide output of experiment with various settings, assess go-live

Bug:

Blank text in field A not being treated as nulls

Creating your prioritized backlog: Value vs Ease of Implementation

Endless choices for ways to improve!

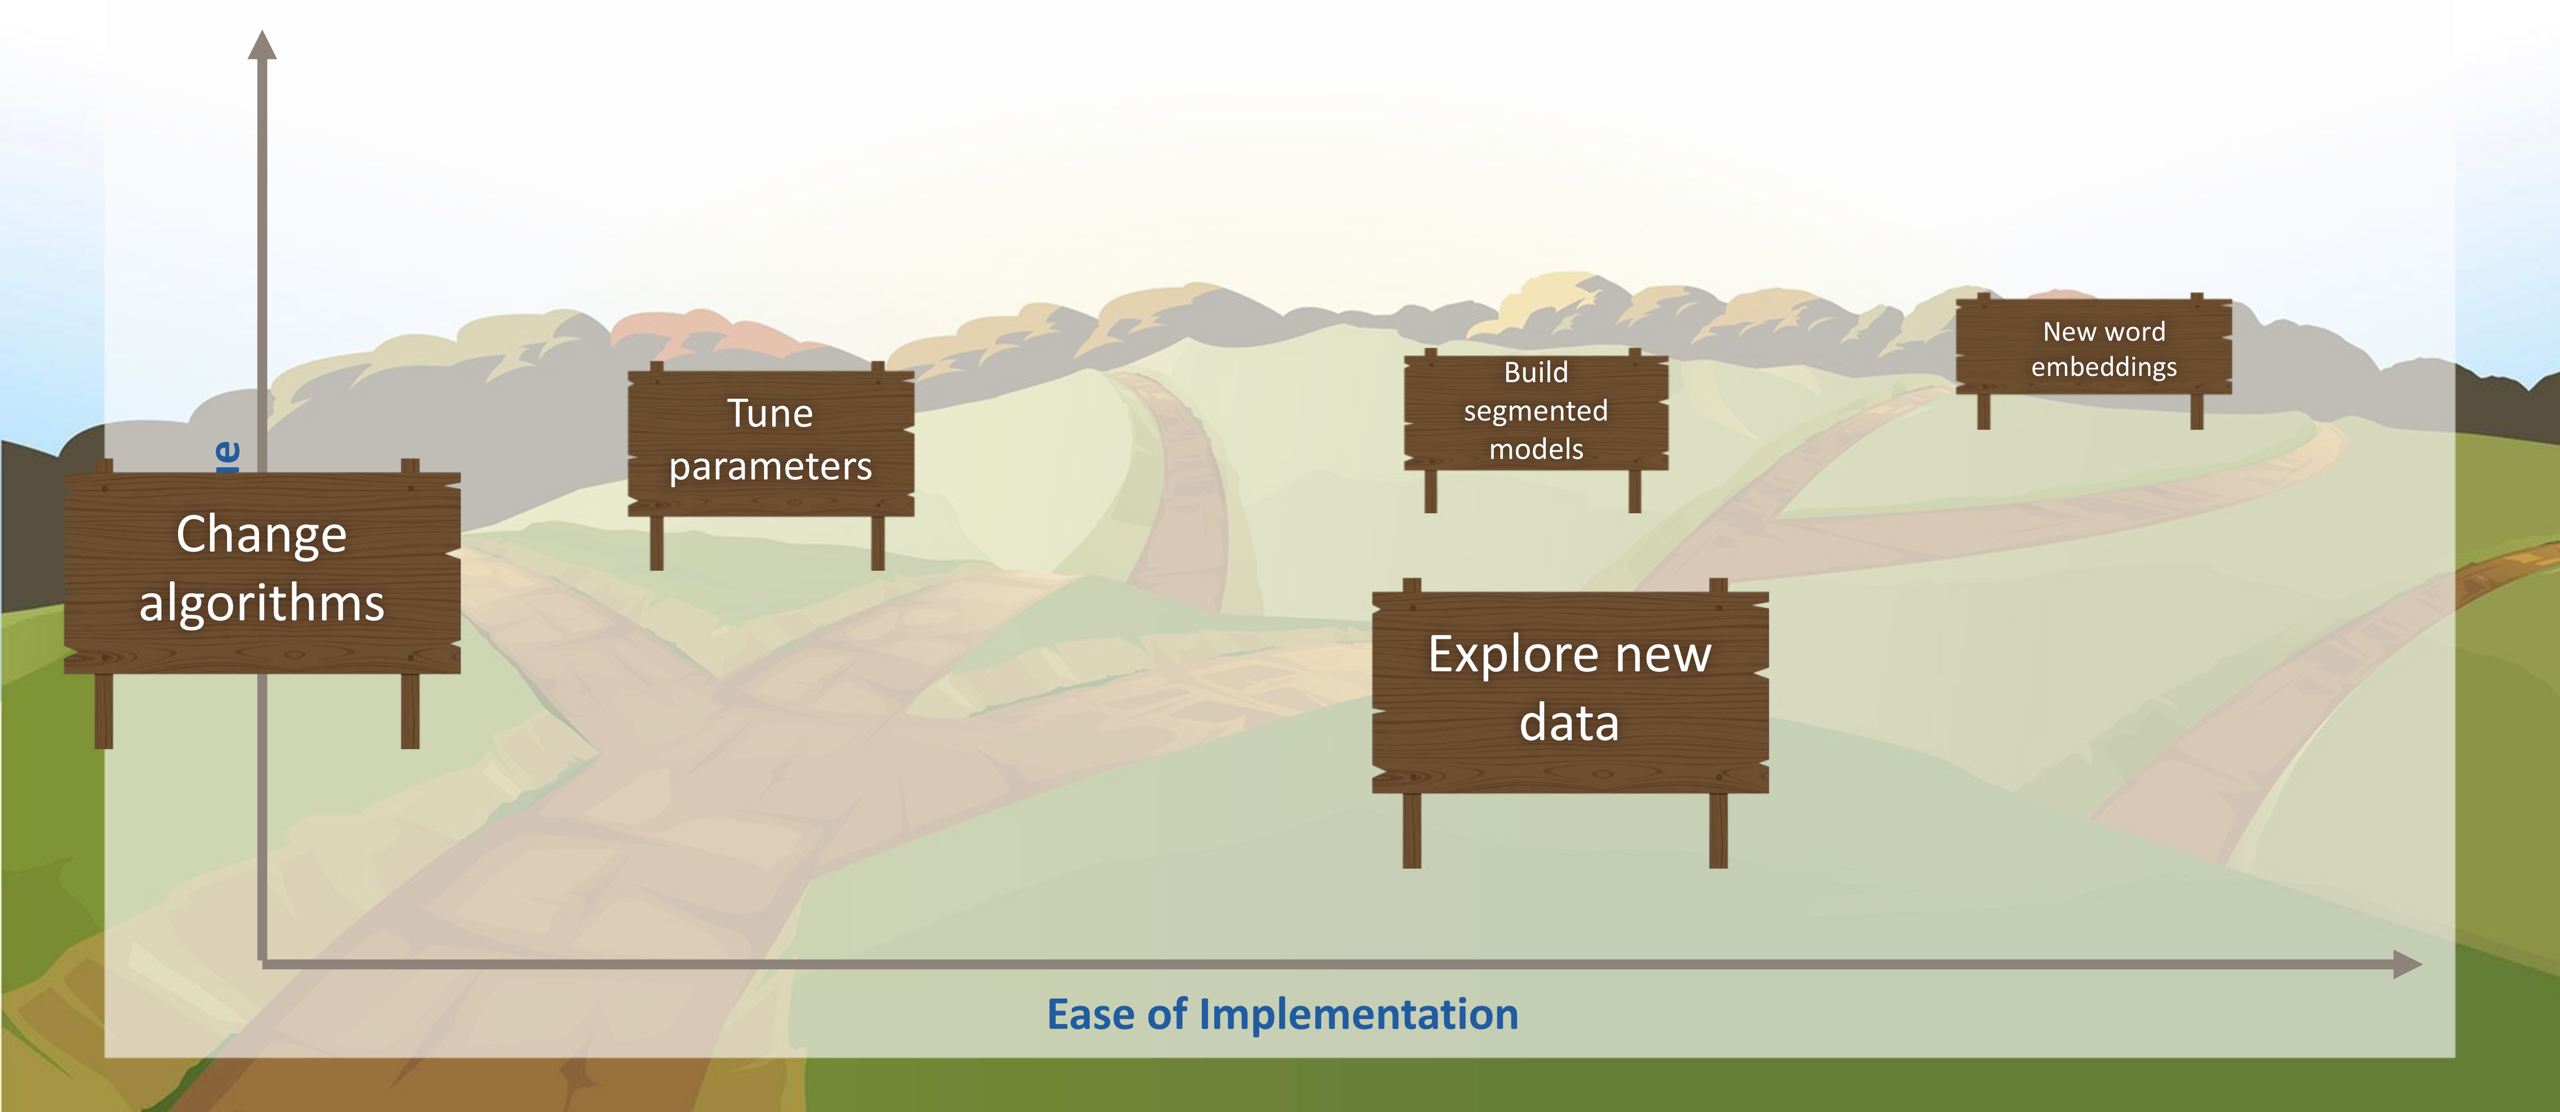

Deploy Monitors, Monitor, Repeat!

134 Models in Production	**216** Models Trained (curr.month)	**99.25%** Models with Above Chance Performance

12 Experiments Run this Week	**35,573,664** Predictions Written Per Day (7 day avg)

Sample Dashboard on Simulated Data

Key Takeaways

- Plan for multiple apps… **always**

  Identify opportunities for reusability in all aspects, even your machine learning pipelines

- Understand your data scientists

  Build a platform to enable their productivity

- Don’t fly blind

  Make sure you can monitor your model health

- Never deploy without a plan for iteration

  How can your data scientist experiment?

  How can your data scientists redeploy?

  What will you do with your old predictions?

Go to github.com/salesforce/TransmogrifAI
to learn more!

Introducing: New Einstein Services for Admins & Developers

Build custom AI you can trust

Einstein Platform Services:

Deploy NLP & Computer Vision inside of Salesforce using APEX Code or Lightning Web Components

Einstein Predictions Service:

Extend predictions outside of Salesforce into any external system like finance, HR, ERP & more

Trusted AI:

Ensure your AI is transparent, responsible & accountable

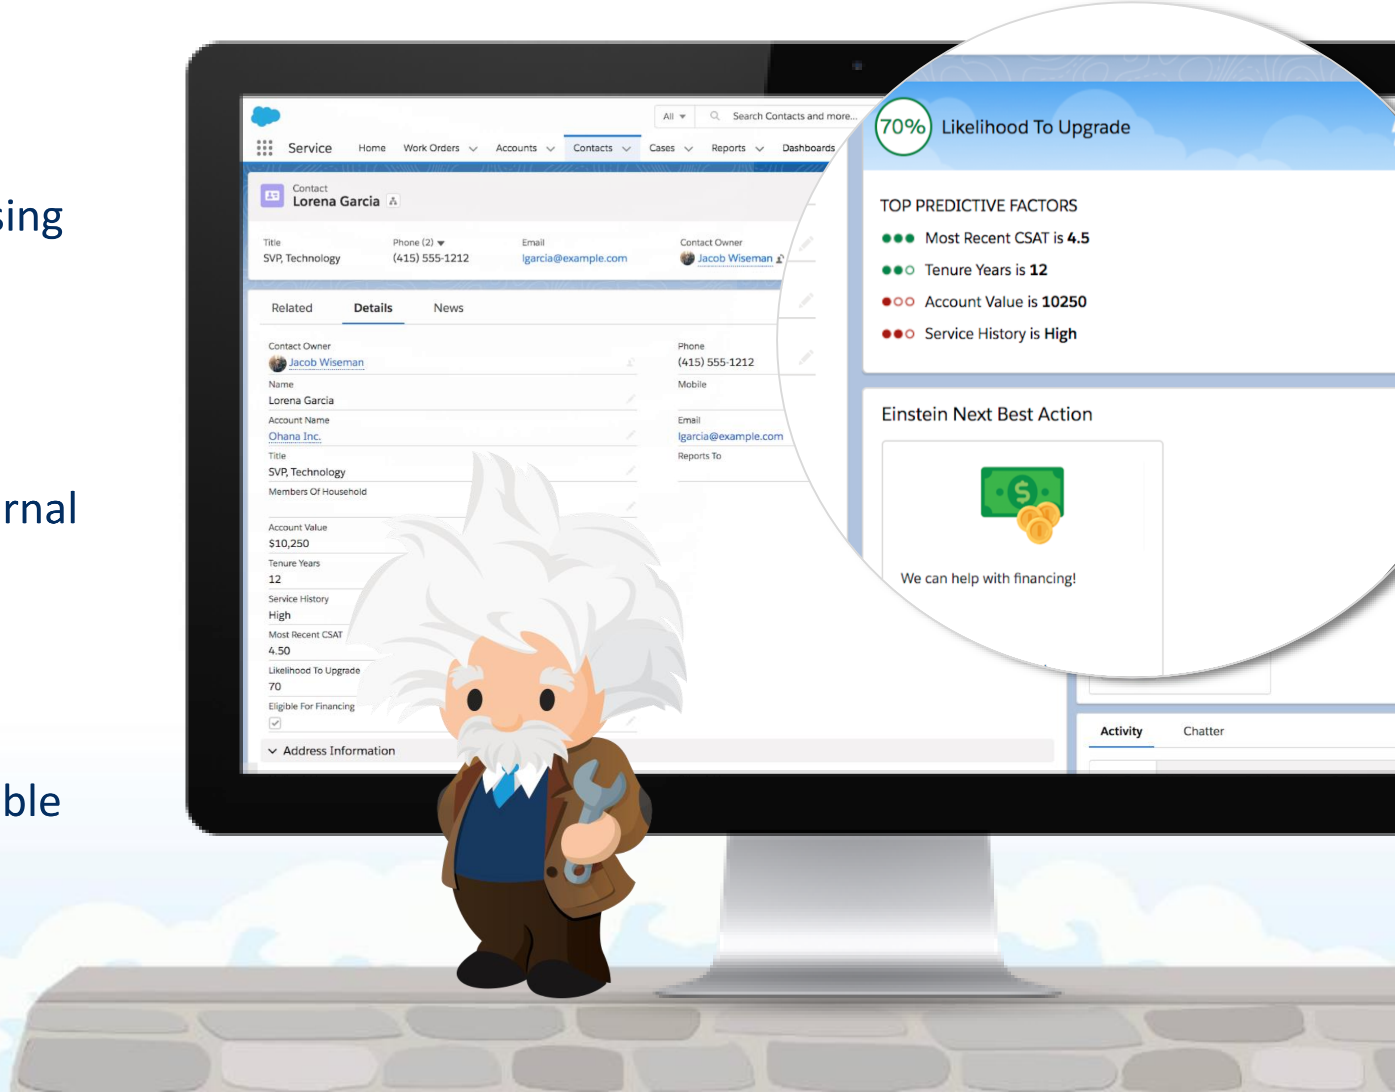


THANK YOU
@itweetsarah

ACM: The Learning Continues…

- Questions/comments about this webcast? learning@acm.org
- ACM Code of Ethics: https://ethics.acm.org
- ACM's Discourse Page: http://on.acm.org
- ACM TechTalks (on-demand archive): https://learning.acm.org/techtalks
- ACM Learning Center: http://learning.acm.org
- ACM SIGKDD: https://www.sigkdd.org/
- ACM Queue: http://queue.acm.org