ECE444: Software Engineering

AI for SE

Shurui Zhou

The Edward S. Rogers Sr. Department
of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

Learning Goals

Understand the AI-enhanced software development processes in practices

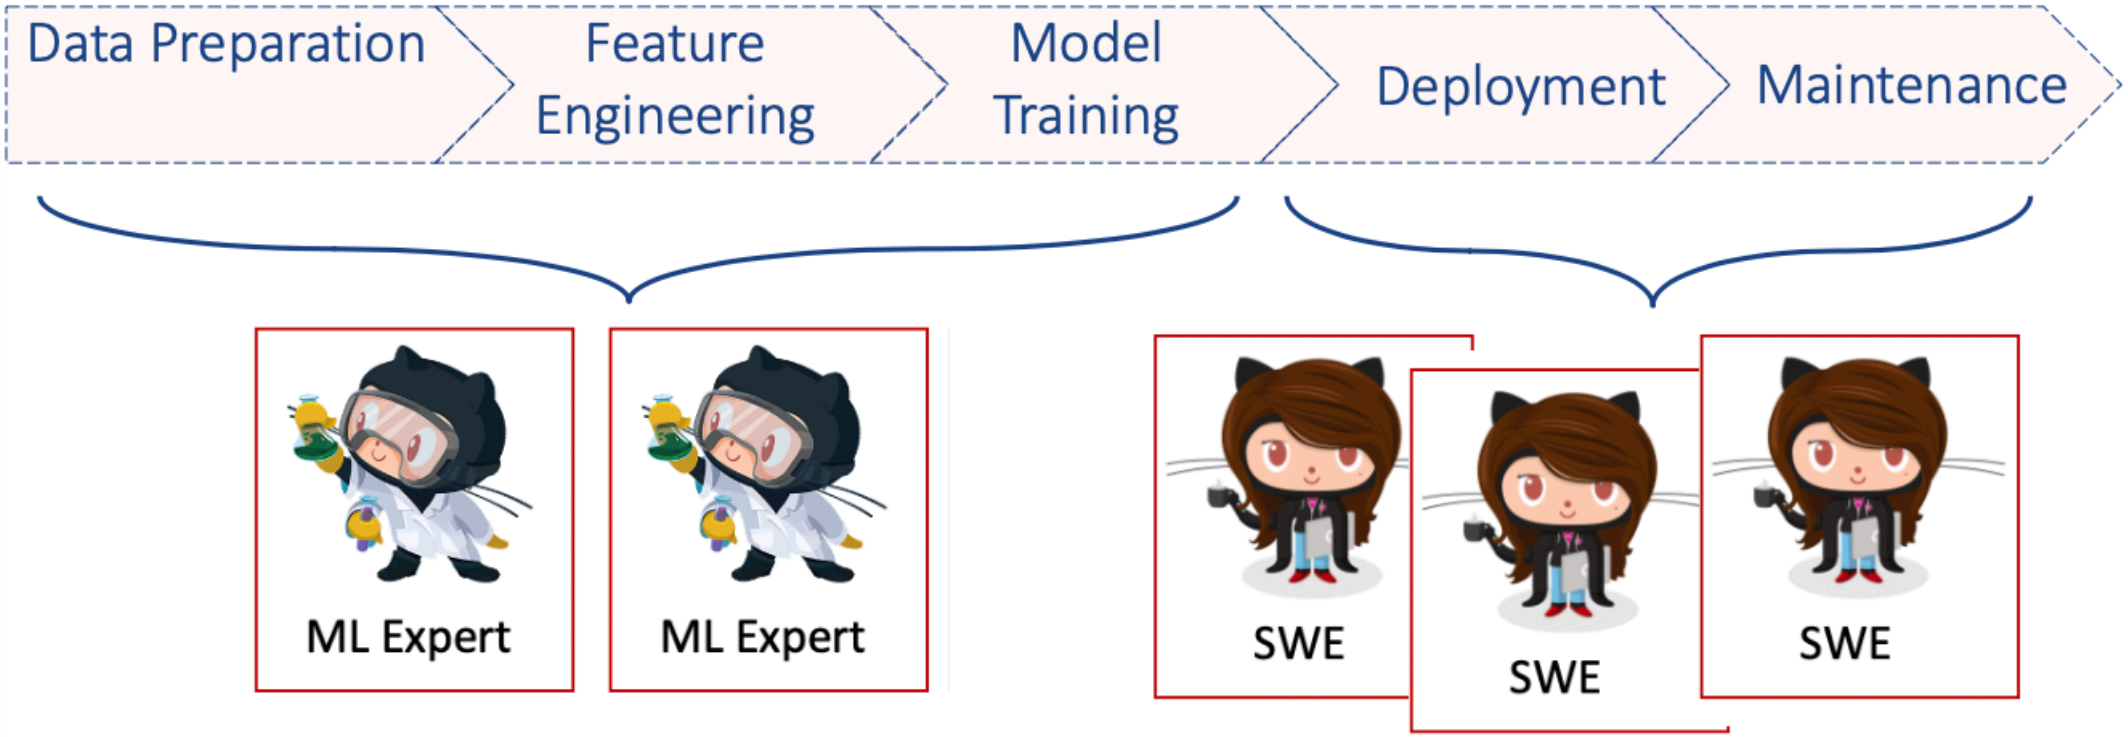
Data Preparation
Feature Engineering
Model Training
Deployment
Maintenance
ML Expert
ML Expert
SWE
SWE
SWE

2016 IEEE/ACM 38th IEEE International Conference on Software Engineering

The Emerging Role of Data Scientists on Software Development Teams

Miryung Kim
UCLA
Los Angeles, CA, USA
miryung@cs.ucla.edu

Thomas Zimmermann   Robert DeLine   Andrew Begel
Microsoft Research
Redmond, WA, USA
{tzimmer, rdeline, andrew.begel}@microsoft.com

ARTWORK: TAMAR COHEN, ANDREW J BUBOLTZ, 2011, SILK SCREEN ON A PAGE FROM A HIGH SCHOOL YEARBOOK, 8.5" X 12"

DATA

Data Scientist: The Sexiest Job of the 21st Century

by Thomas H. Davenport and D.J. Patil

FROM THE OCTOBER 2012 ISSUE

SUMMARY | SAVE | SHARE | COMMENT 5 | TEXT SIZE | PRINT | $8.95 BUY COPIES

When Jonathan Goldman arrived for work in June 2006 at LinkedIn, the business networking site, the place still felt like a start-up. The company had just under 8 million accounts, and the number was growing quickly as existing members invited their friends and colleagues to join. But users weren't seeking out connections with the people who were already on the site at the rate executives had expected. Something was apparently missing in the social experience. As one LinkedIn manager put it, "It was like arriving at a conference reception and

WHAT TO READ NEXT

Big Data: The Management Revolution

Making Advanced Analytics Work for You

Google Flu Trends' Failure Shows Good Data > Big Data

VIEW MORE FROM THE

October 2012 Issue

Typical data science workflow

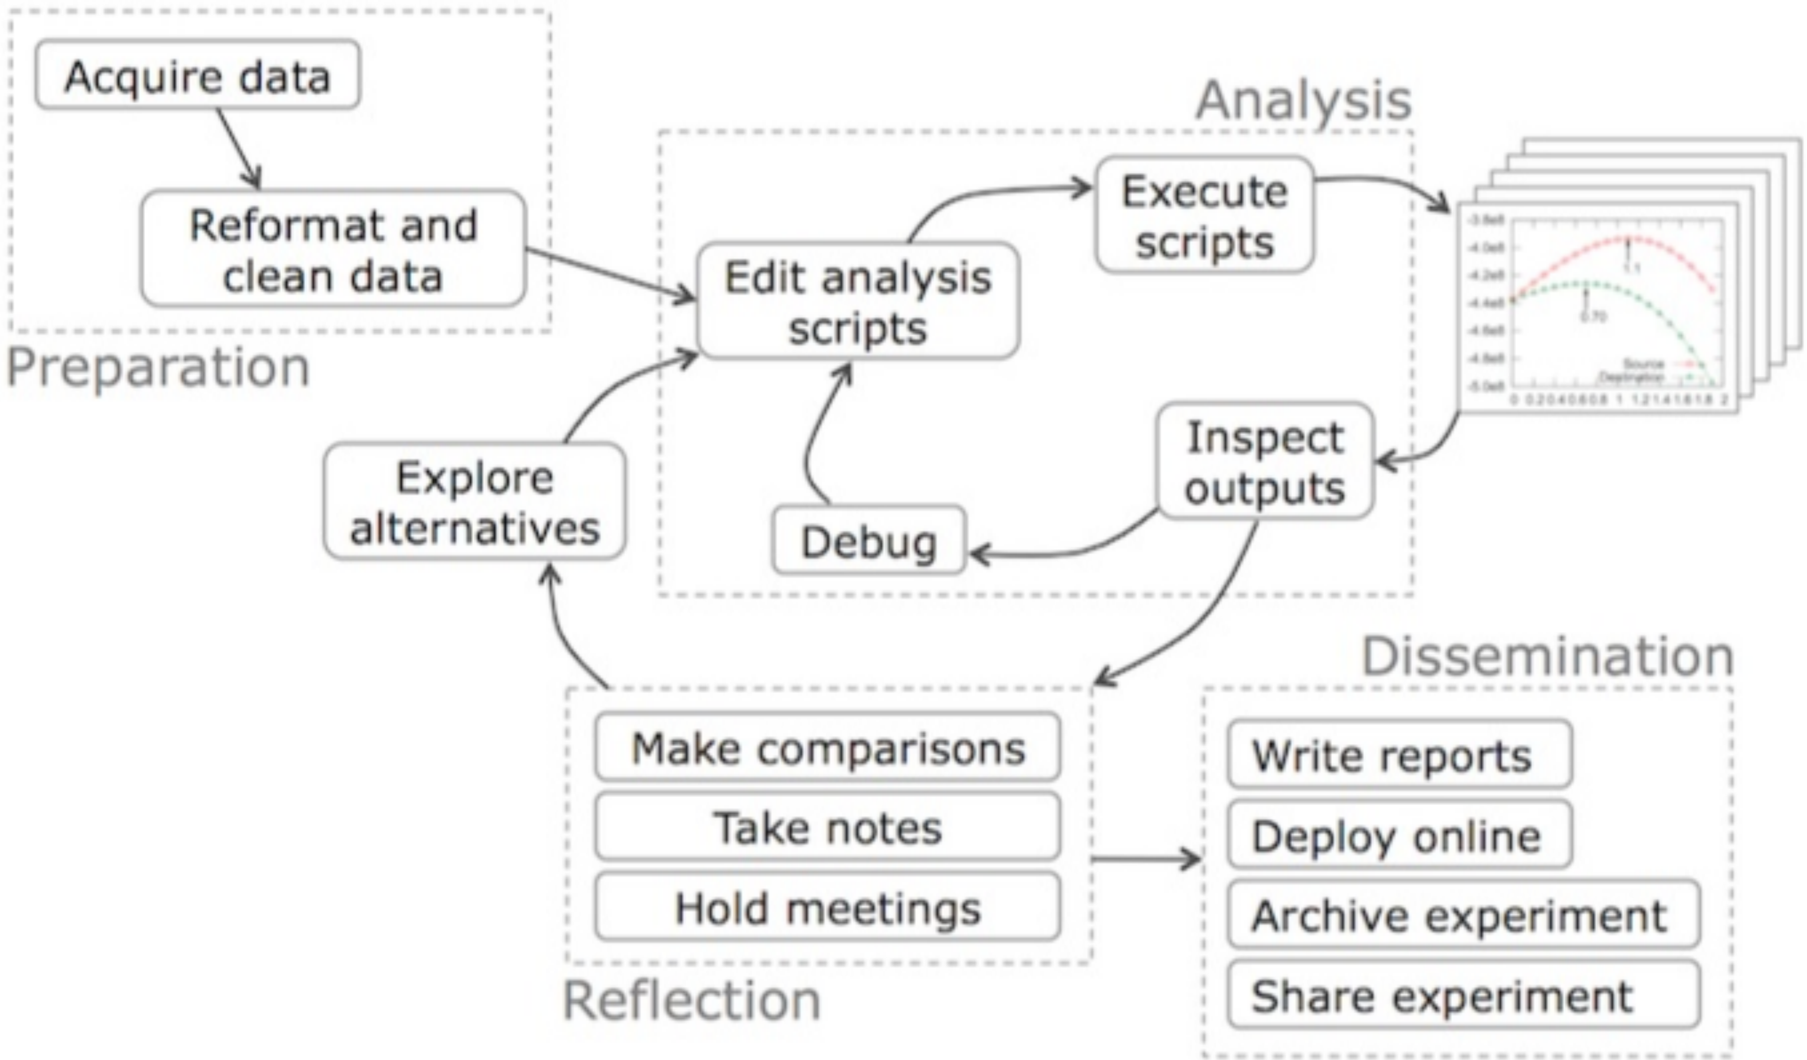

https://cacm.acm.org/blogs/blog-cacm/169199-data-science-workflow-overview-and-challenges/

Why are Data Scientists Needed on SW Teams?

Software companies want to **experiment with real users**, e.g., A/B testing, flighting, games and rewards.

"Instead of having an army of testers to go off and generate a bunch of data, that data's already here. It's more authentic because it's real customers on real machines, real networks. You no longer have to simulate and anticipate what the customer's gonna do." [P10]

Demand for Data Collection Rigor

What about storage, what about speed? What about legal, what about privacy? There is an entire gamut of things that you need to jump through hoops to collect the instrumentation. [P1]

The Emerging Role of Data Scientists on Software Development Teams, Kim et al. ICSE' 16

What Do Data Scientists Work on?

Performance Regression

Are we getting better in terms of crashes or worse? [P3]

Requirements Identification

If you see the repetitive pattern where people don't recognize, the feature is there. [P3]

Root Cause Analysis

What areas of the product are failing and why? [P3]

Bug Prioritization

Oh, cool. Now we know which bugs we should fix first. Then how can we reproduce this error? [P5]

Server Anomaly Detection

Is this application log abnormal w.r.t. the rest of the data? [P12]

Failure Rate Estimation

Is the beta ready to ship? [P8]

Customer Understanding

How long do our users use the app? [P1]

What are the most popular features? [P4]

Cost Benefit Analysis

How many customer service calls can we prevent if we detect this type of anomaly? [P9]

Activities of Data Scientists

Collecting	Building the data collection platform
	Injecting telemetry
	Building the experimentation platform
Analyzing	Data merging and cleaning
	Sampling
	Shaping, feature selection
	Defining sensible metrics
	Building predictive models
	Defining ground truth
	Hypothesis testing
Using and Disseminating	Operationalizing models
	Defining actions and triggers
	Applying insights/models to business

Data Scientist Working Styles

- Insight Providers
- Modeling Specialists
- Platform Builders
- Polymaths
- Team Leaders

Analyze This! 145 Questions for Data Scientists in Software Engineering

Andrew Begel
Microsoft Research
Redmond, WA, USA
Andrew.Begel@microsoft.com

Questions for Data Scientists in Software Engineering: A Replication

Hennie Huijgens
Delft University of Technology
Delft, The Netherlands
h.k.m.huijgens@tudelft.nl

Ayushi Rastogi
Ernst Mulders*
Delft University of Technology
Delft, The Netherlands
a.rastogi@tudelft.nl
ernst@mulde.rs

Georgios Gousios
Arie van Deursen
Delft University of Technology
Delft, The Netherlands
g.gousios@tudelft.nl
arie.vandeursen@tudelft.nl

"Please list up to five questions you would like [a team of data scientists who specialize in studying how software is developed] to answer."

What metrics are the **best predictors of failures**?

If I increase **test coverage**, will that actually increase software quality?

What is the **data quality** level used in empirical studies and how much does it actually matter?

Are there any **metrics that are indicators of failures** in both Open Source and Commercial domains?

I just submitted a **bug report**. Will it be fixed?

Should I be writing **unit tests** in my software project?

How can I tell if a piece of software will have **vulnerabilities**?

Is strong **code ownership** good or bad for software quality?

Do **cross-cutting concerns** cause defects?

Does **Distributed/Global software development** affect quality?

Does **Test Driven Development** (TDD) produce better code in shorter time?

How would you approach these questions with data?

- Where to focus testing effort?
- Is our review practice effective?
- Is the expensive static analysis tool paying off?
- Should we invest in security training?
- What is a good team size?

Evaluate Effectiveness of an Intervention

- Controlled experiments
  - Compare group with intervention against control group without,
  - Randomized controlled trials, AB testing, …
  - Ideally blinded
- Natural experiments, Quasi experiments
  - Compare similar groups that naturally only differ in the intervention
  - No randomized assignment of treatment condition
- Time series analyses
  - Compare measures before and after intervention, preferably across groups with the intervention at different times

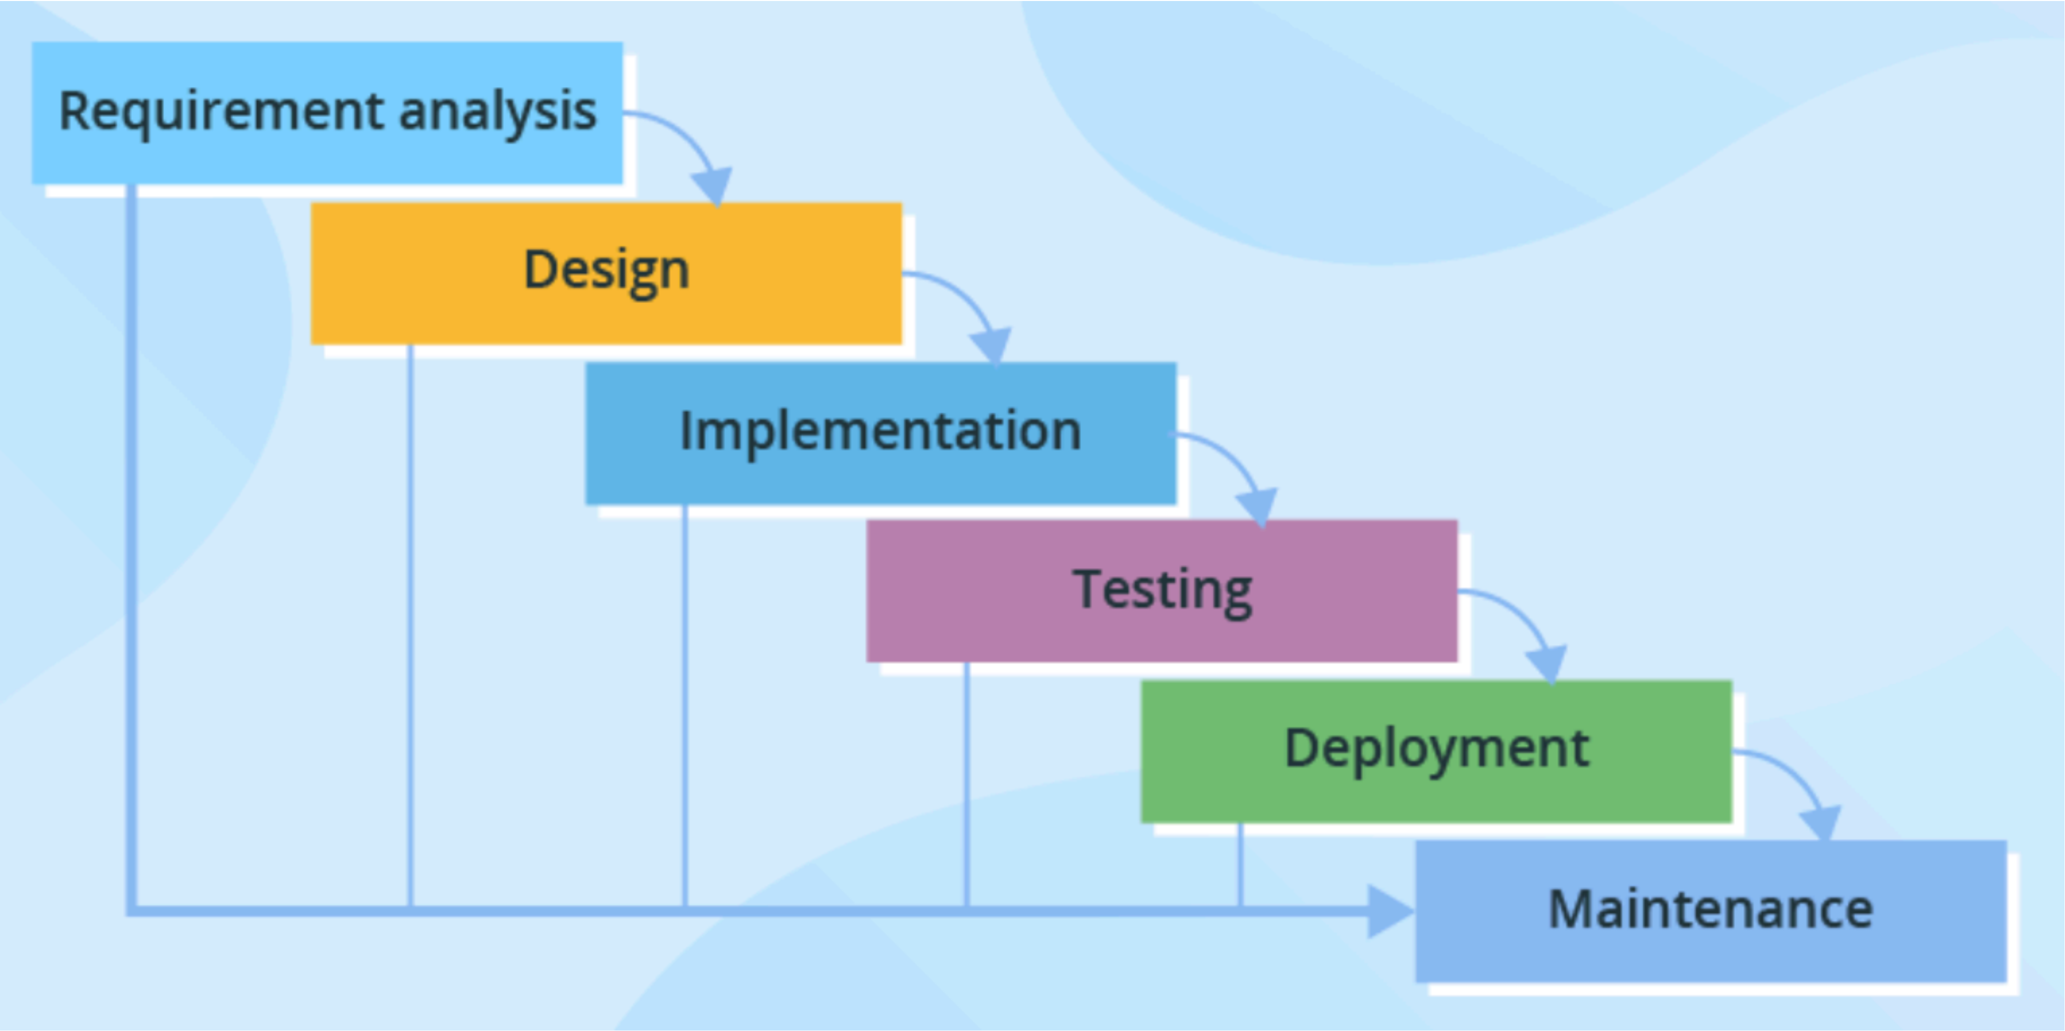

Requirement analysis
Design
Implementation
Testing
Deployment
Maintenance

Requirement

Engineering breakthrough: IBM introduces Watson AI for RQA

By Maggie Mae Armstrong | 2 minute read | February 28, 2019

Watson AI uses natural language processing and understanding to analyze a requirement’s text, suggesting improvements that leverage industry best practices for writing high quality requirements, based on the INCOSE Guidelines for Writing Good Requirements.

https://www.ibm.com/blogs/internet-of-things/iot-ibm-introduces-watson-ai-for-engineering-with-requirements-management-quality-assistant/

visure SOLUTIONS PRODUCTS SERVICES RESOURCES COMPANY BLOG SUPPORT CHOOSE LANGUAGE Download a trial

Requirements Management ALM platform

Visure Requirements is an easy and comprehensive Requirements Management tool. It integrates in

Meet Alice: Your Cognitive Assistant for Business Analysis

Your requirements gathering is about to get easier, better and faster. How? Artificial intelligence (AI) and machine learning. The most tedious part of the requirements process can often be gathering and elicitation. Yet that part of the process is well-suited for AI's capabilities.

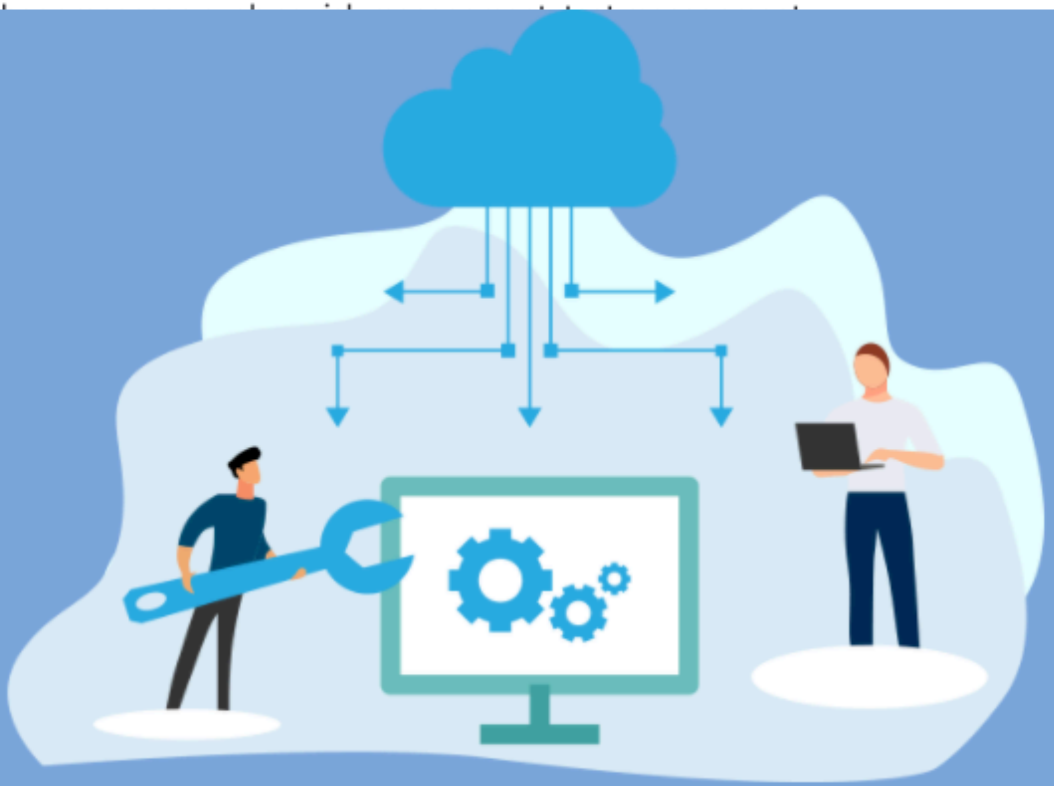

Requirement Analysis

- Detection of Hidden Feature Requests from Massive Chat Messages via Deep Siamese Network . **Shi et al.** *(ICSE), 2020*

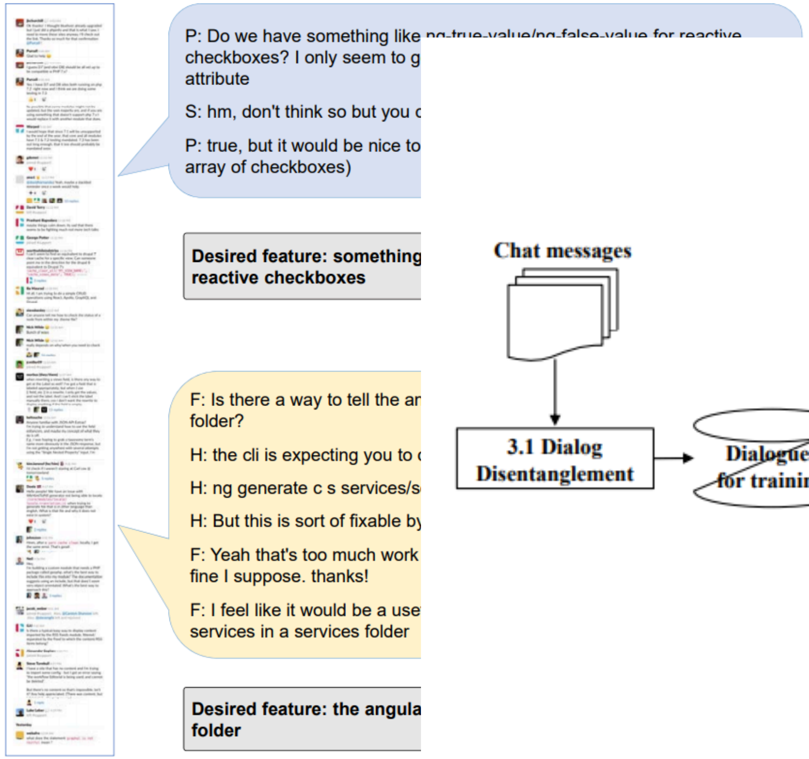

Figure 1: Example chat messa
where requests to desired featu

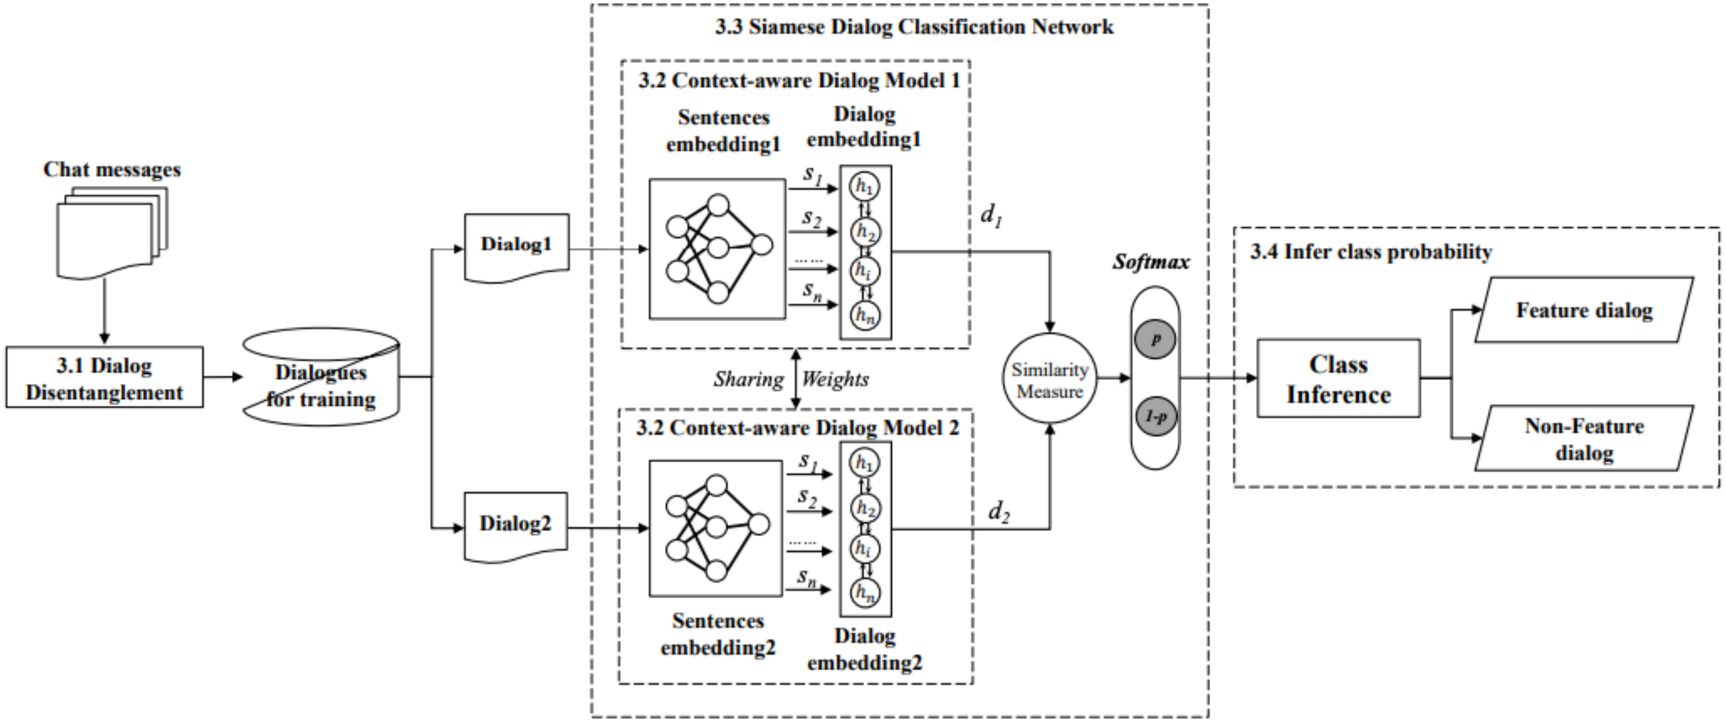

Figure 2: The Overview of FRMiner

Design

Meet AIDA: Your Artificial Intelligence Design Assistant

One of the big topics in design right now is artificial intelligence. Can a computer program actually design a website? Can it help a person speed up or improve the process?

Bookmark is taking the theory to a whole new level with its Artificial Intelligence Design Assistant, or AIDA for short. AIDA learns your needs and desires and uses this knowledge to create the perfect website for you. Today we're taking a look at how it works!

The Ultimate Designer Toolkit: 2 Million+ Assets

BOOKMARK

ON:	19 APR 2017
BY:	CARRIE COUSINS
CATEGORY:	REVIEWS
LENGTH:	4 MIN READ

Coding

Code Completion

codota

Give your development team AI superpowers

Codota automatically learns the patterns and rules in your company's proprietary code and makes sure your developers have the best code insights, whenever and wherever they need them.

https://www.kite.com/

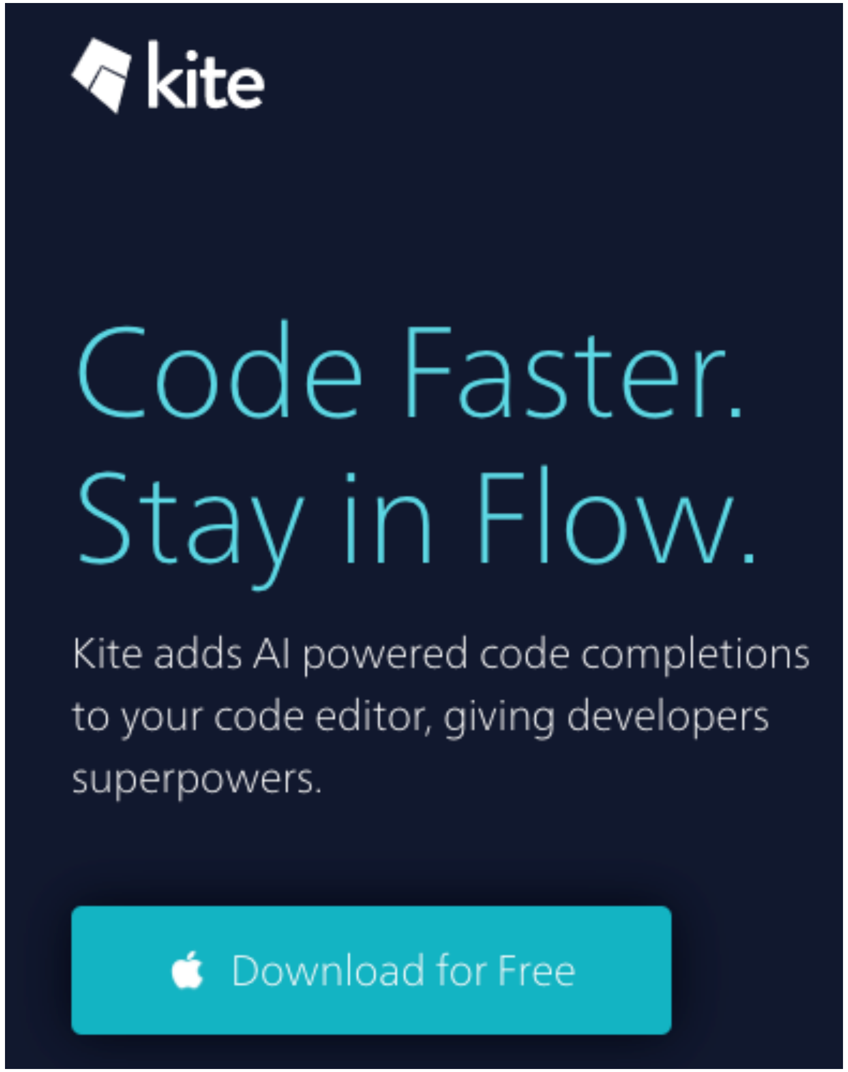


```python
import os
import sys

def count_py_files_in_repos(dirname):
    i
```

kite.com

Use machine-learning-assisted code completion

You can utilize machine learning models to rank most suitable items higher in the suggestions list.

To do this, in the **Settings/Preferences** dialog ⌘, go to **Editor | General | Code Completion** and enable the **Rank completion suggestions based on Machine Learning** option under **Machine Learning-Assisted Completion**.

This feature is experimental, so ranking may not change noticeably.

Visual Studio IntelliCode

Microsoft | 8,246,945 installs | ★★★★☆ (54) | Free

```
loss = tf.reduce_sum(tf.square(linear_model - y))
optimizer = tf.train.GradientDescentOptimizer(0.01)

train = optimizer
```

Deep TabNine: A Powerf Code Autocompleter For Developers

Synced Follow
Jul 18, 2019 · 3 min read

https://medium.com/syncedreview/deep-tabnine-a-powerful-ai-code-autocompleter-for-developers-70454a5953fe

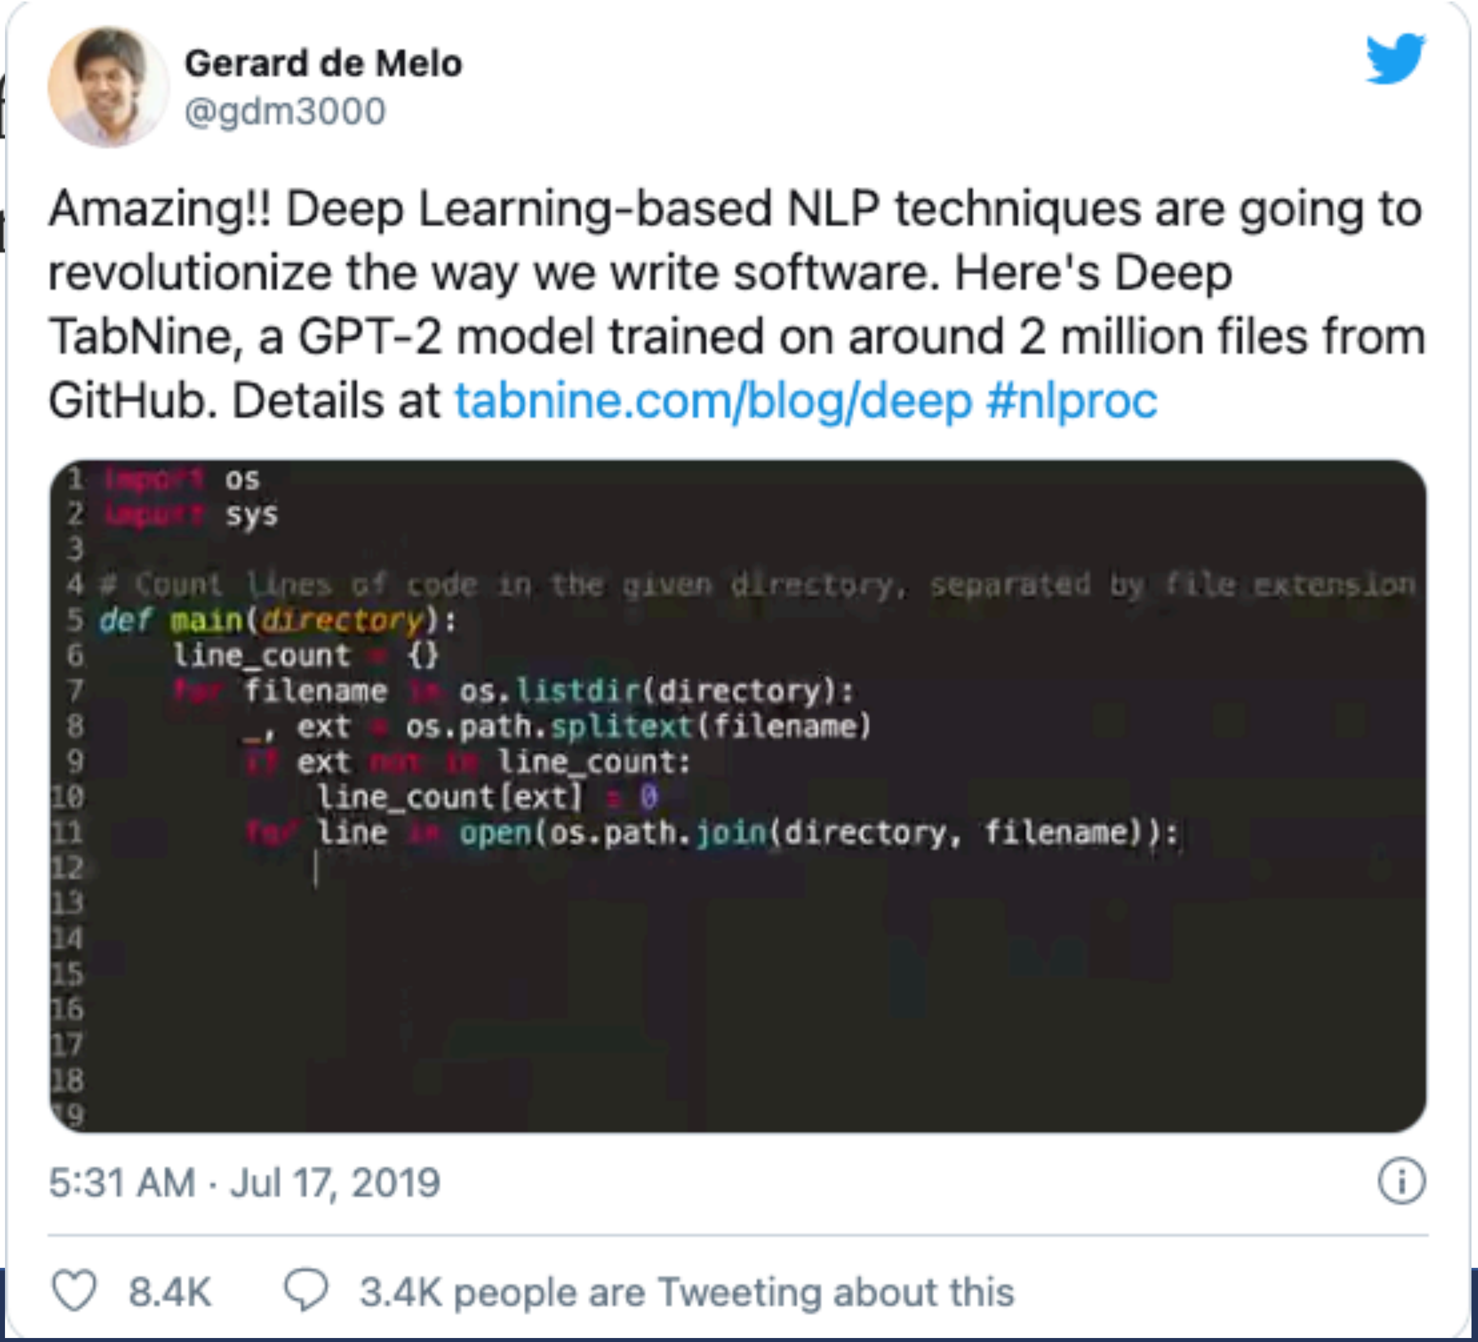

Sharif Shameem @sharifshameem · Jul 13

This is mind blowing.

With GPT-3, I built a layout generator where you just describe any layout you want, and it generates the JSX code for you.

W H A T

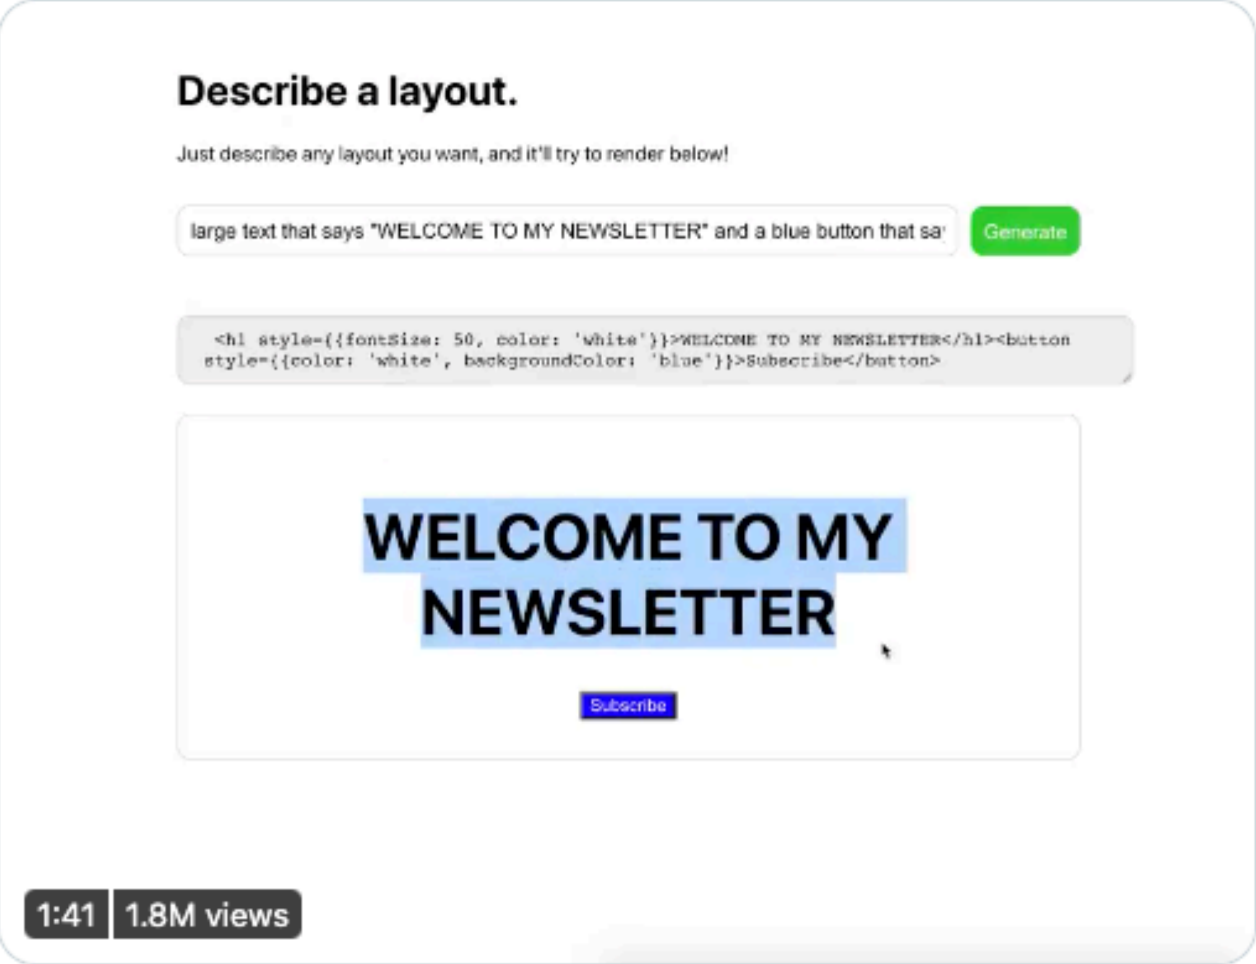

699 14.3K 42.3K

https://twitter.com/sharifshameem/status/1282676454690451457?ref_src=twsrc%5Etfw%7Ctwcamp%5Etweetembed%7Ctwterm%5E1282676454690451457%7Ctwgr%5E%7Ctwcon%5Es1_&ref_url=https%3A%2F%2Fanalyticsindiamag.com%2Fopen-ai-gpt-3-code-generator-app-building%2F

FACEBOOK AI

https://ai.facebook.com/blog/aroma-ml-for-code-recommendation/

ML APPLICATIONS | DEVELOPER TOOLS

Aroma: Using machine learning for code recommendation

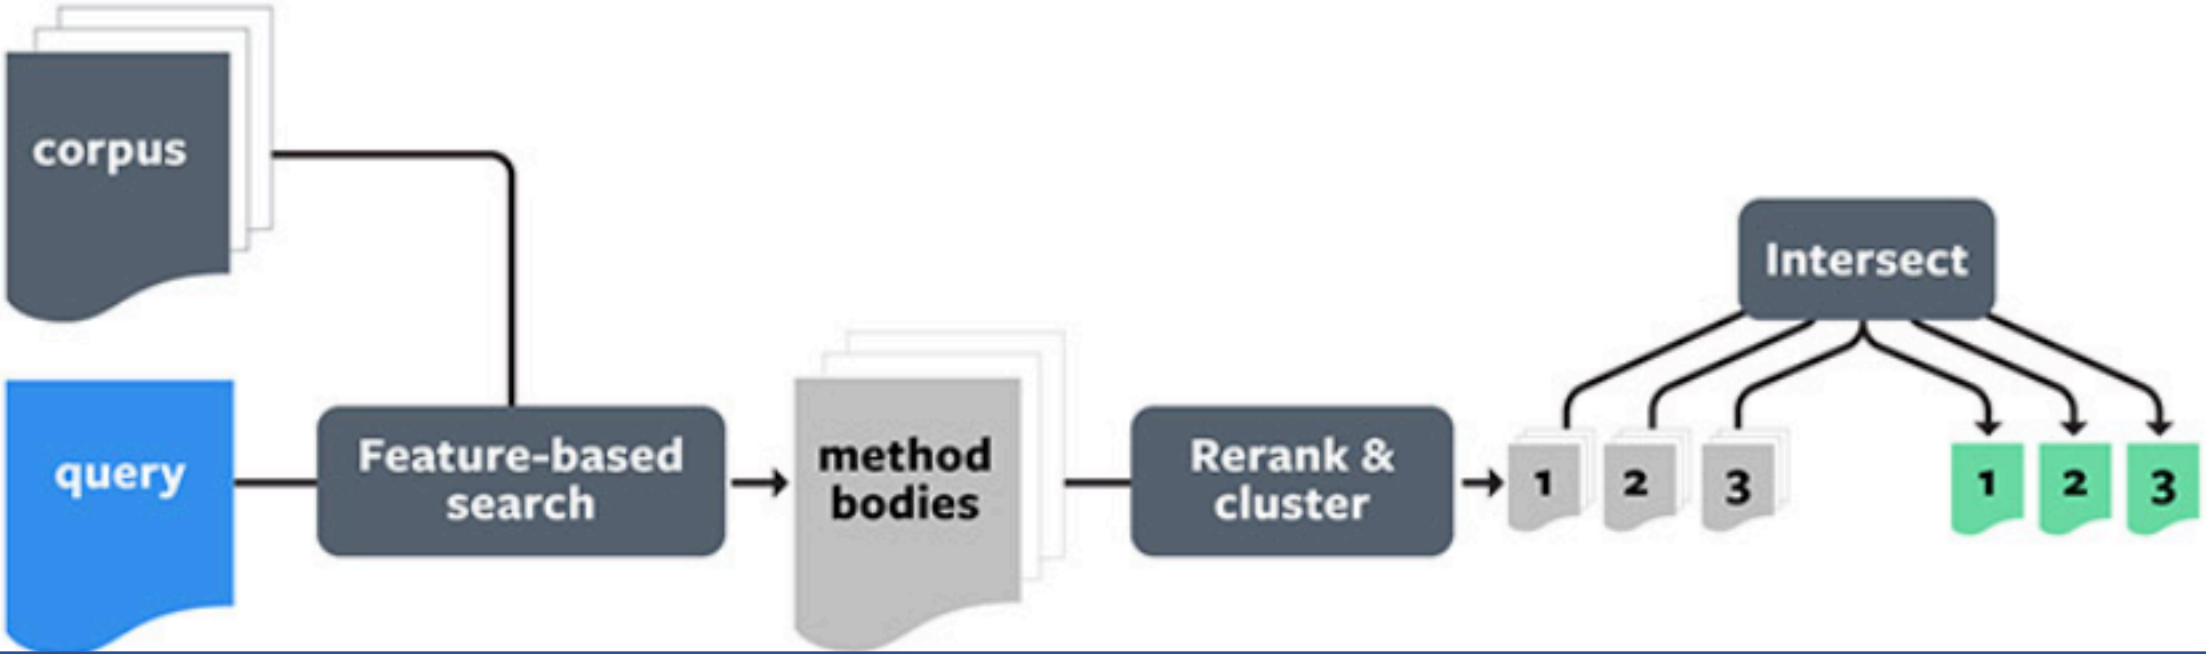

Debugging

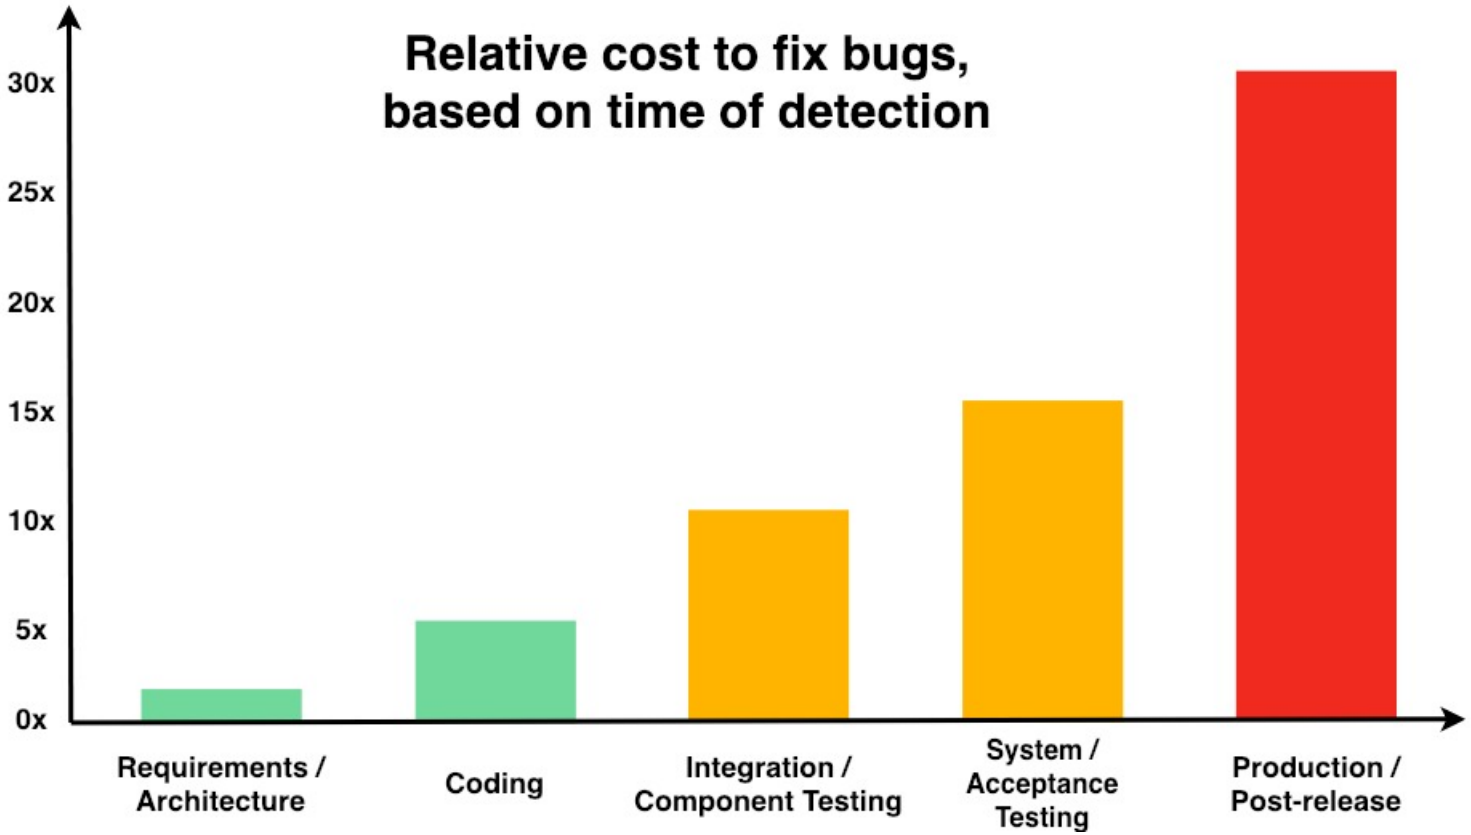

https://deepsource.io/blog/exponential-cost-of-fixing-bugs/

Ubisoft : ML catches 70 % of bugs prior to testing

"The statistical nature of machine learning involves us changing the way we work," he says. Unlike traditional software, in which developers write out rules for the application to follow, machine-learning algorithms use data to guide how the software should act.

-- Yves Jacquier,
executive director, production studio services, Ubisoft Montreal

https://www.pmi.org/learning/library/ai-debug-code-11523

Debugging

https://engineering.fb.com/2018/11/06/developer-tools/getafix-how-facebook-tools-learn-to-fix-bugs-automatically/

Getafix: How Fa
automatically

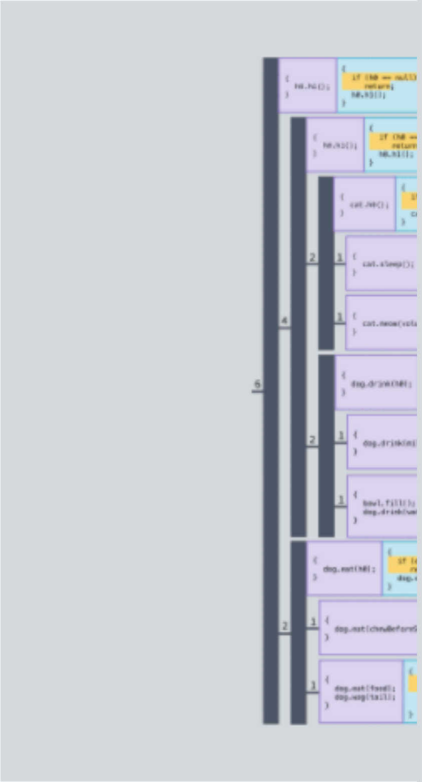

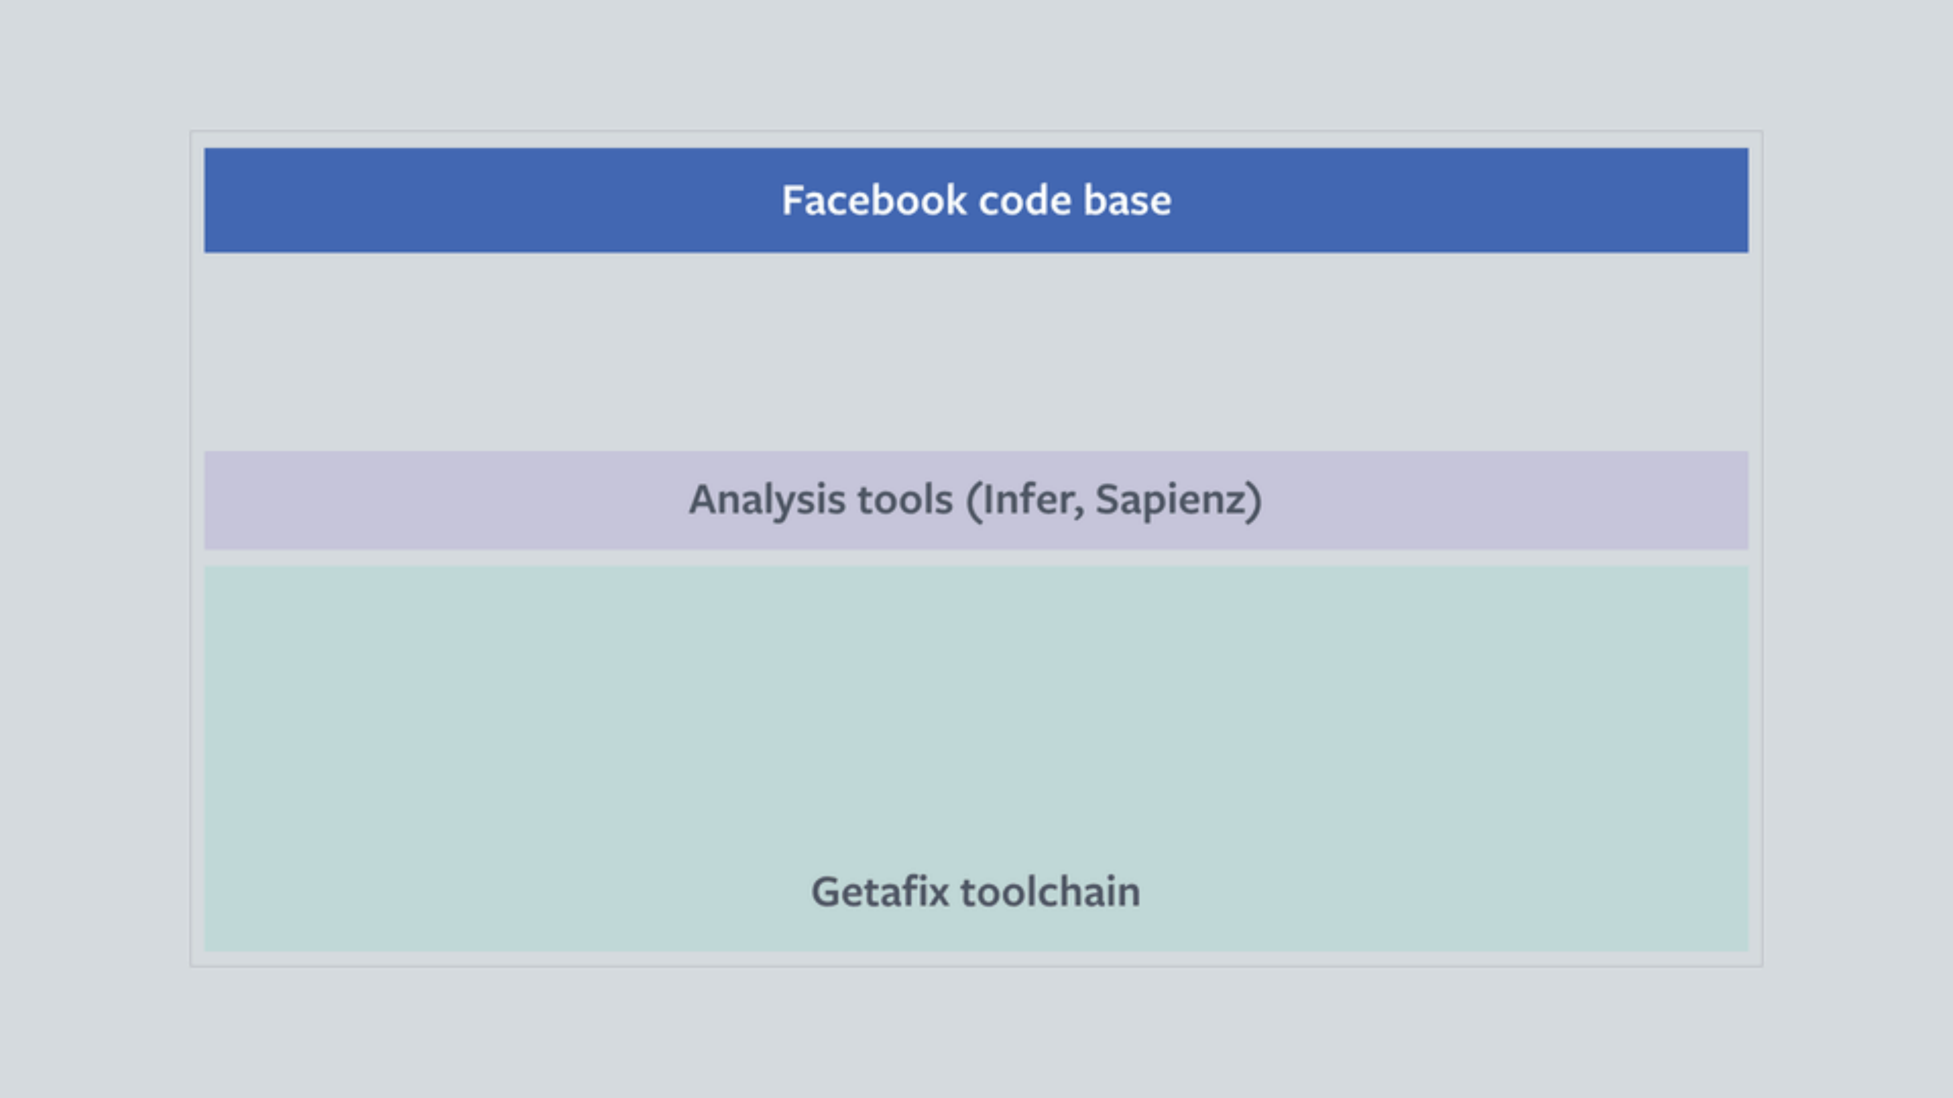

QA

Case study: Functionize Eliminates Agvance Test Maintenance with Machine Learning

September 11, 2018 - Ankur Verma

DeepCode

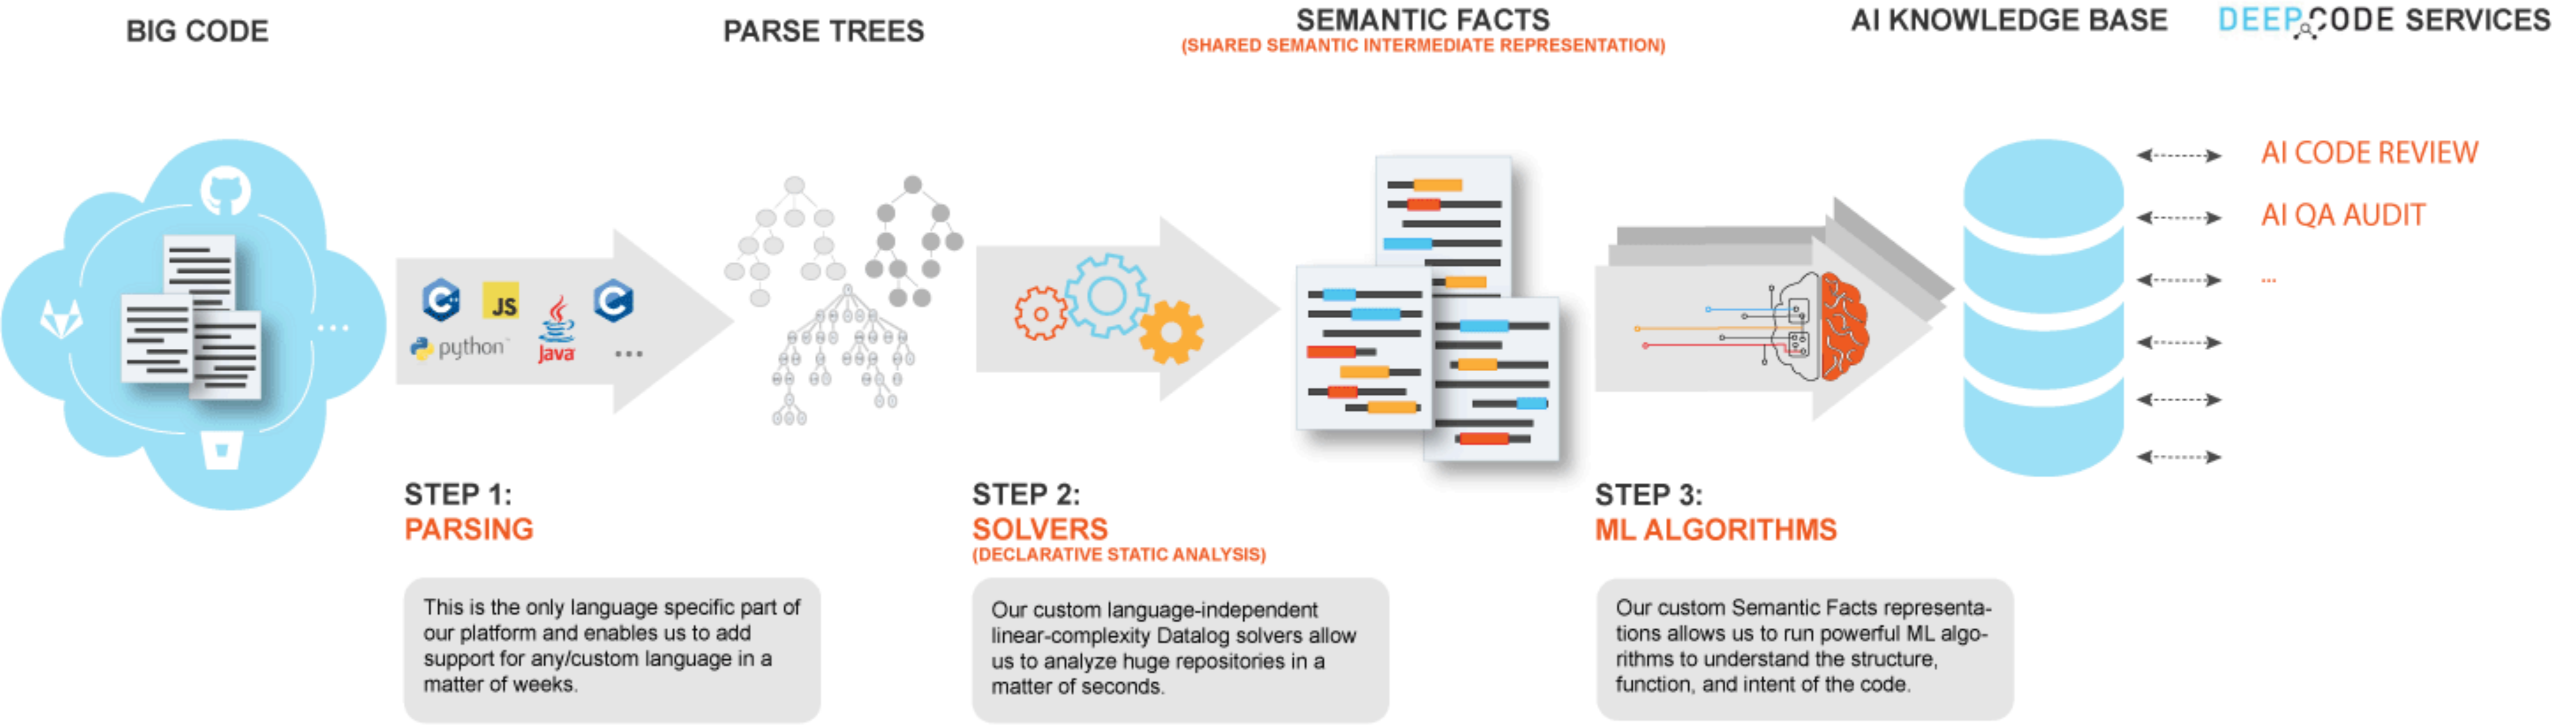
BIG CODE
PARSE TREES
SEMANTIC FACTS
(SHARED SEMANTIC INTERMEDIATE REPRESENTATION)
AI KNOWLEDGE BASE
DEEPCODE SERVICES
python
Java
JS
AI CODE REVIEW
AI QA AUDIT
...
STEP 1:
PARSING
This is the only language specific part of our platform and enables us to add support for any/custom language in a matter of weeks.
STEP 2:
SOLVERS
(DECLARATIVE STATIC ANALYSIS)
Our custom language-independent linear-complexity Datalog solvers allow us to analyze huge repositories in a matter of seconds.
STEP 3:
ML ALGORITHMS
Our custom Semantic Facts representations allows us to run powerful ML algorithms to understand the structure, function, and intent of the code.

CodeQL

CodeQL helps you explore code quickly to find and eradicate all variants of vulnerabilities before they become a problem.

By automating variant analysis, CodeQL enables product security teams to find zero-days and variants of critical vulnerabilities.

Deployment

UPDATED 07:00 EST / APRIL 23 2019

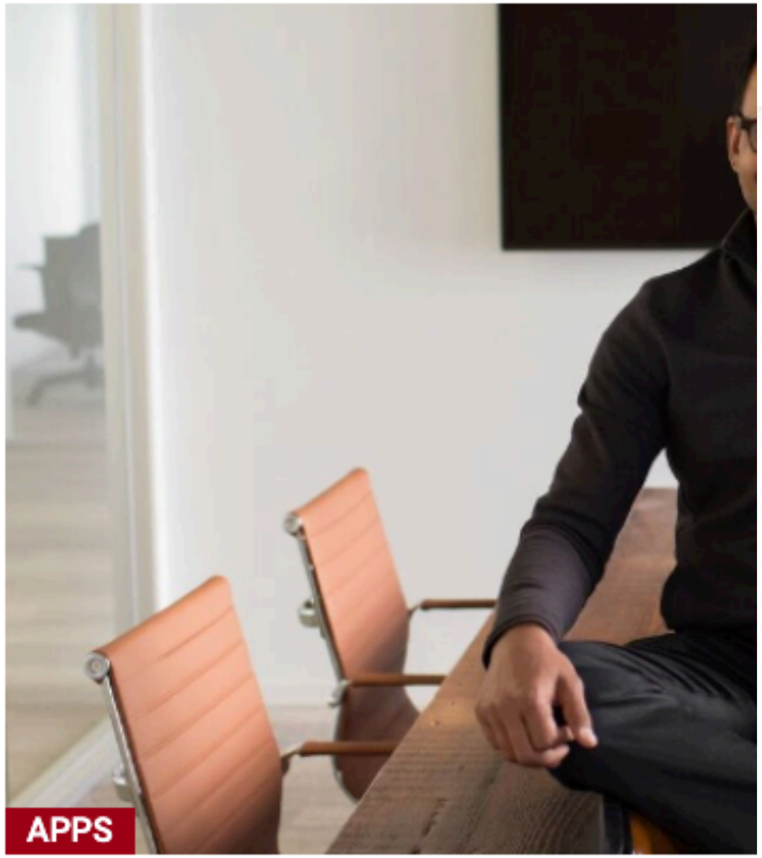

Continuous app delivery firm H

BY MIKE WHEATLEY

https://siliconangle.com/2019/04/23/continuous-app-delivery-firm-harness-raises-60m/

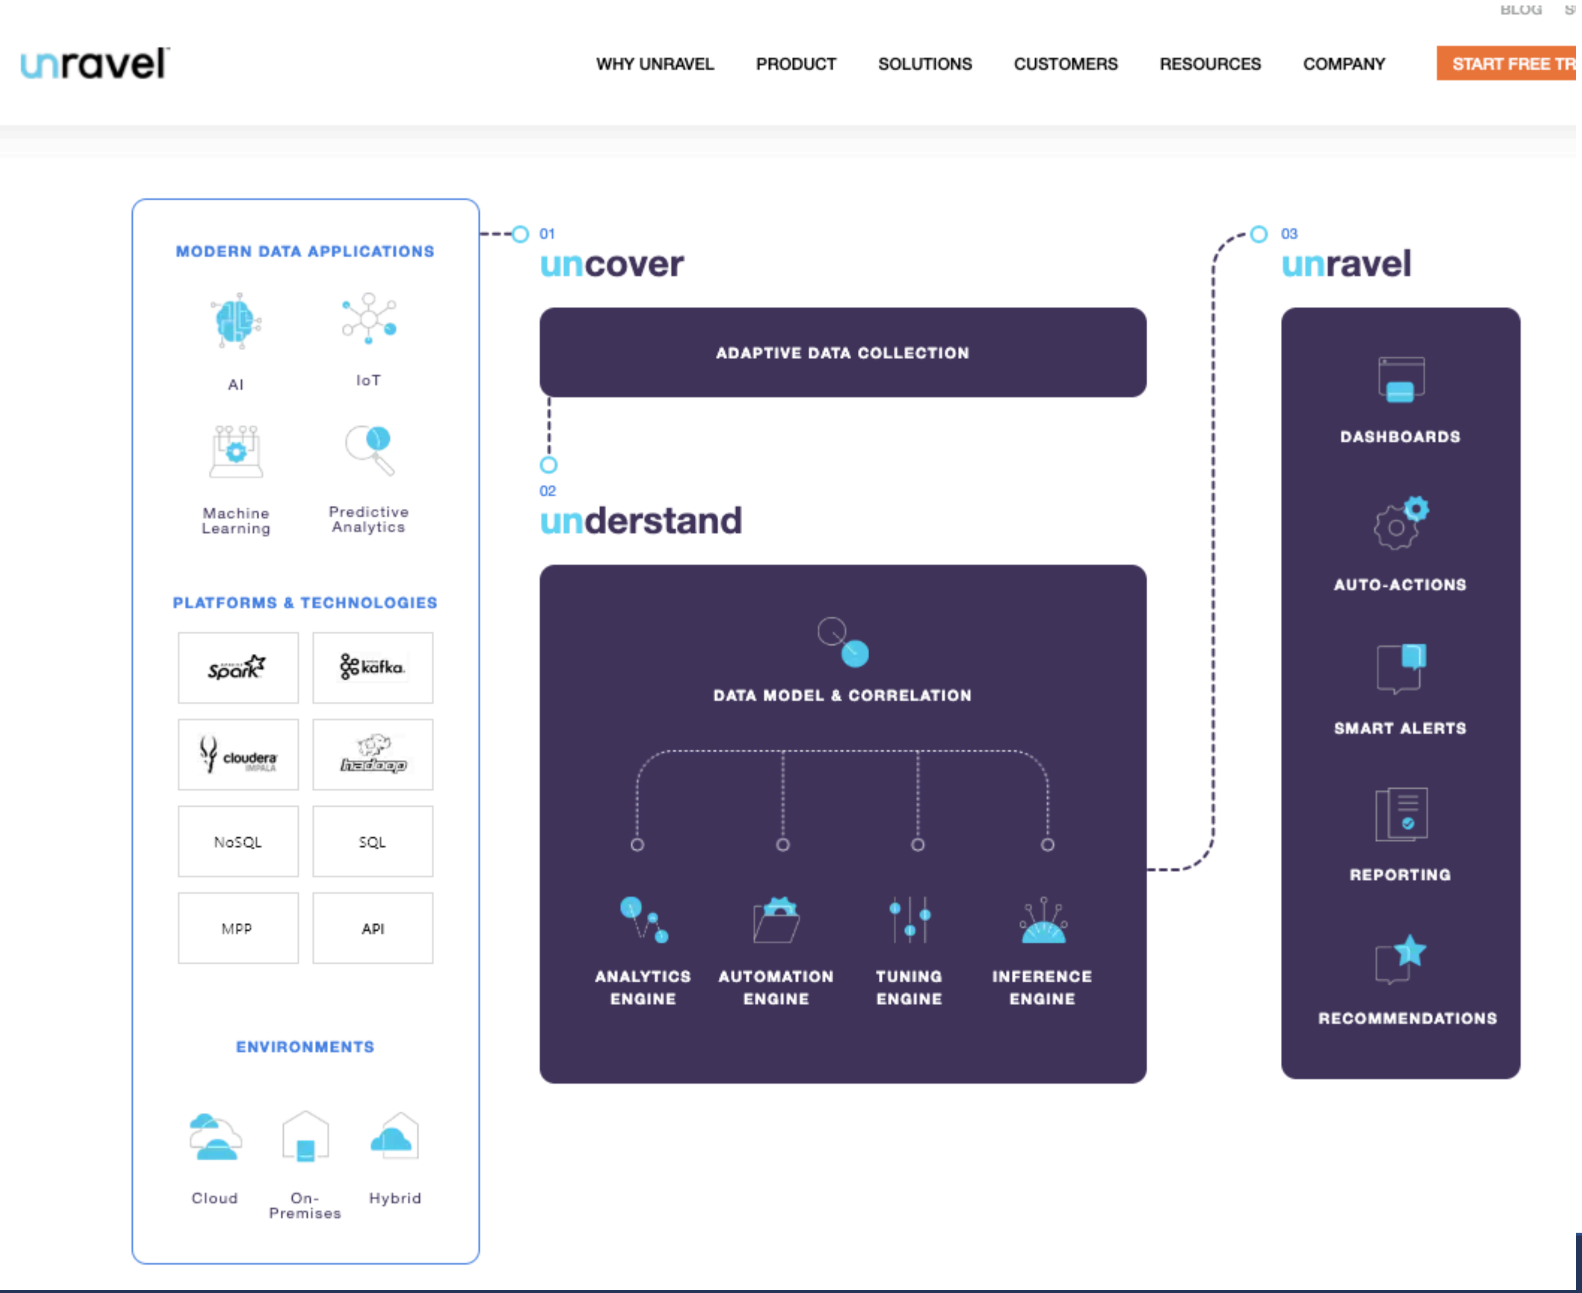

https://opsani.com/product/

OPSANI

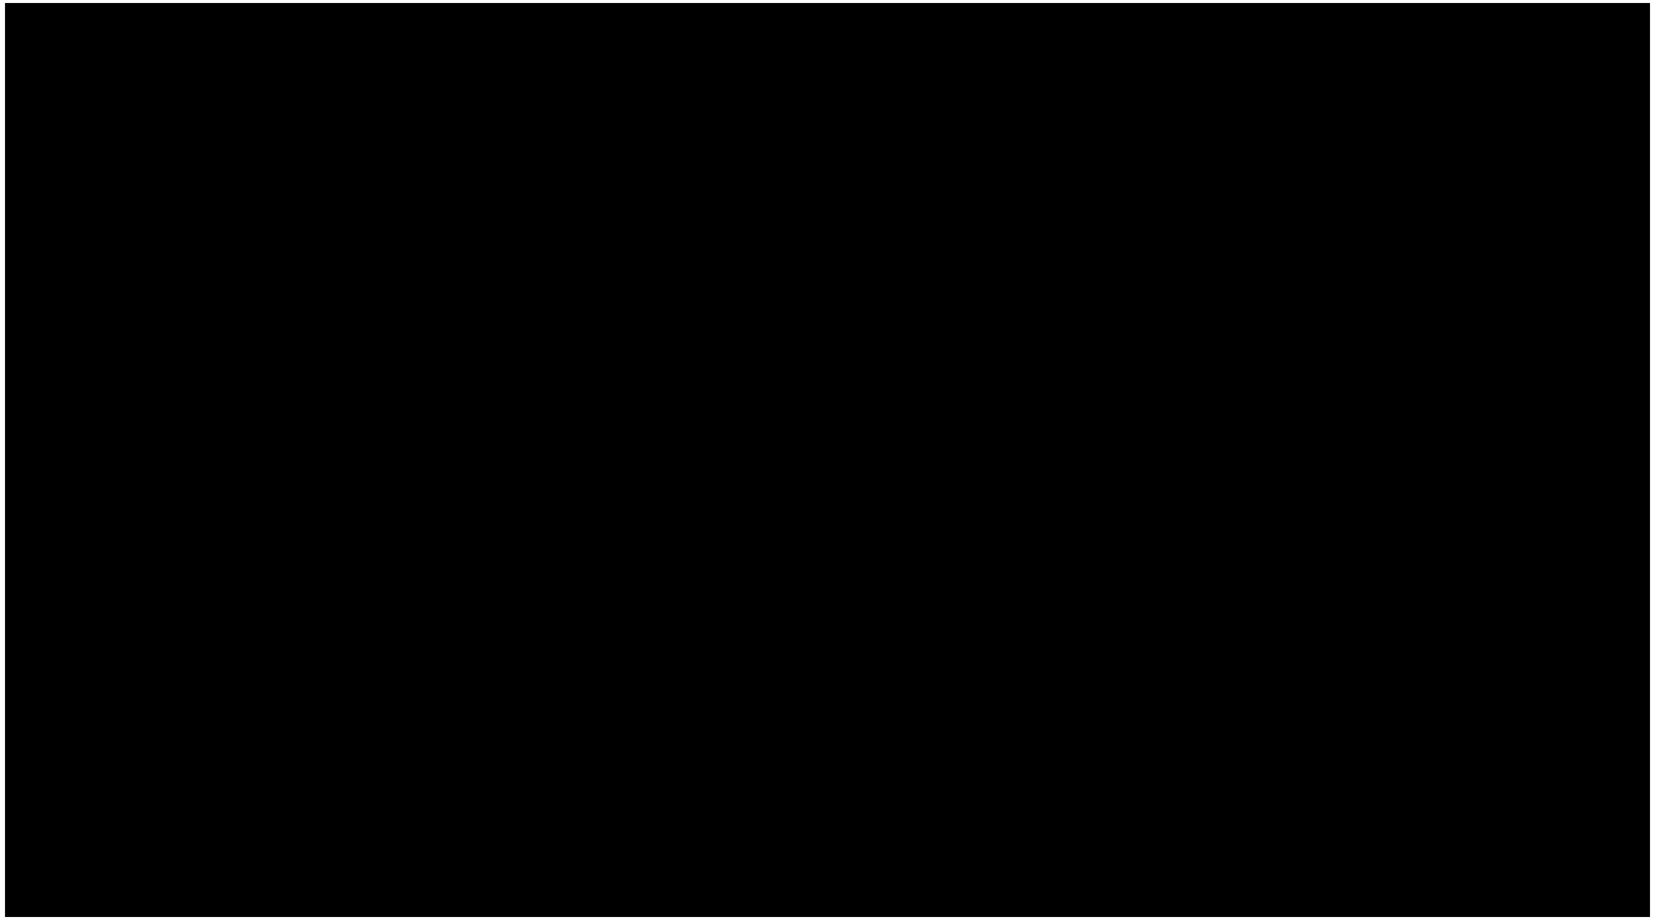

Project Management

How CraneAi uses Artificial Intelligence to help teams build apps faster

Behind the scenes look at how CraneAi's uses artificial intelligence to empower teams

Ryan Hickman Follow

Nov 30, 2018 · 5 min read

TARA

“ Tara’s mission is to help teams develop their plans with visibility and predictability.”

COMPANY NEWS

Ford and Cisco are turning to an AI company to find the best freelance programmers

EDITOR
January 18, 2018

Trade-off?

Software 2.0

Andrej Karpathy Nov 11, 2017 · 8 min read

I sometimes see people refer to neural networks as just "another tool in your machine learning toolbox". They have some pros and cons, they work here or there, and sometimes you can use them to win Kaggle competitions. Unfortunately, this interpretation completely misses the forest for the trees. Neural networks are not just another classifier, they represent the beginning of a fundamental shift in how we write software. They are Software 2.0.

https://medium.com/@karpathy/software-2-0-a64152b37c35

I got myself a cool AI T-Shirt - then the sticker began to peel off