



Monitoring Modern Architectures with Data Science

QCon 2017
Dave Casper, CTO



Abstract

Much has changed since simple distributed client/server architectures and so-too have the technologies and industry practices around monitoring.

Cloud-Native, DevOps, blue/green deployments, server-less, edge/fog, IoT all fit into a world much better handled by the emerging Artificial Intelligence for IT Operations domain more-so than traditional ITIL/SDLC approaches.



Abstract

Software continues to eat the world. Software automates, defines.

The world is "going digital" and it's quite exciting -- but this always-connected-from-everything-to-everywhere world adds complexity to software systems and this talk will dive in to some of that complexity and how modern data science and algorithms are being applied to "fight machines with machines," so to speak.





-25822282



623992118



1343963318





moogsoft





discovery
monitoring
(observing)
analytics



fluid infrastructure

containers dc/os

server-less

software defined/dynamic



data/tx from

**anything
anywhere
anytime**

mobile

IoT

bots/RUM





millions



"if/else"
rules



millions



algorithms
ML

noise filt.

clustering

deja

vu

prc



AIOps

AI for IT Ops

Gartner[®]



customer/ business perspective



COURAGE

INSIGHT

ARE YOU READY

TO GO DIGITAL ?

CONTEXT

VELOCITY

“Silicon Valley is coming.

There are hundreds of startups with a lot of brains and money working on various alternatives to traditional banking. **They are very good at reducing the ‘pain points’ ...”**

JAMIE DIMON

JPMorgan Chase & Co.

Chairman & Chief Executive Officer

April 2015

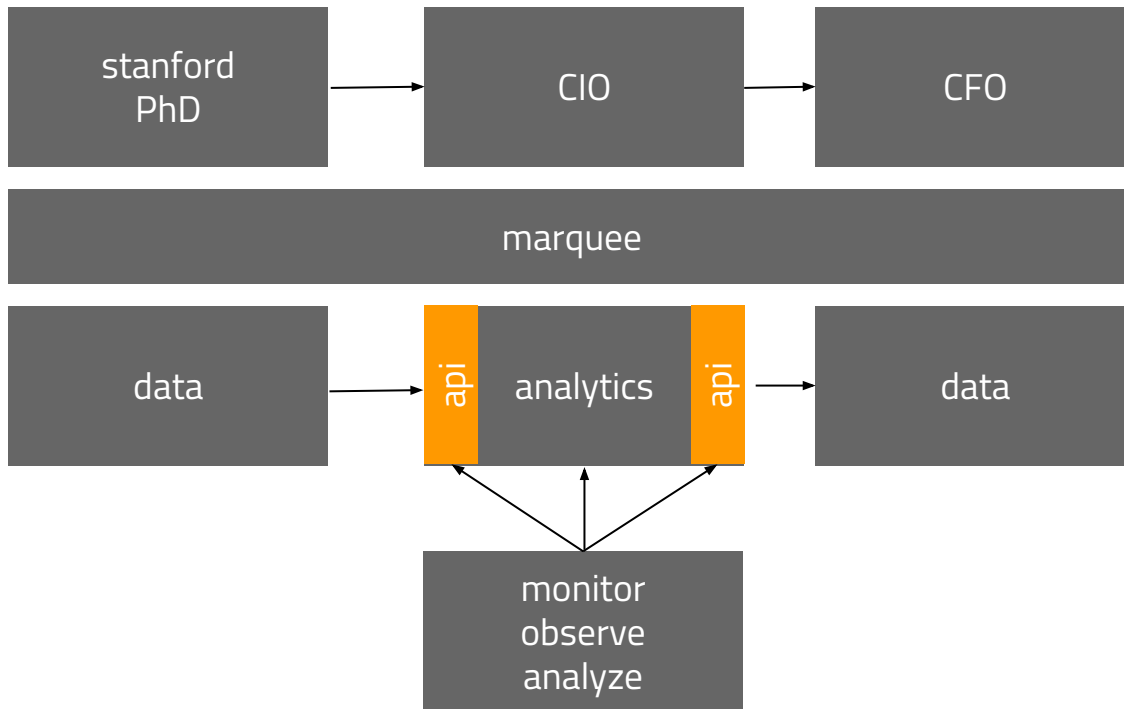


**go digital
or die trying**




**gs wants to
become "google
of wall st."**





THE REALLY BIG PICTURE



In 5 – 10 years, every company will be a **Digital Software Business**

Security, service assurance and consumer centricity become **THE BOARD LEVEL PRIORITY**

Enterprises going **DIGITAL ADOPT HYBRID IT**

traditional

40% Change 60% Run

**Infrastructure Led
Owns Facilities, Data Centers,
Hardware, Networks et al
Has Refresh Cycles caused by
Capital Depreciation
Still using Waterfall for App
Dev**

**Thinking led by Inf
Technologists
(hardware, DB, OS et al)\
Traditional Procurement
Less Agile, Change resistant**

hybrid

60% Change 40% Run

**AppDev starting to lead
Owns less Facilities, Data
Centers, Hardware,
Networks, et al
Still Has Refresh Cycles
caused by
Capital Depreciation
Combination Waterfall &
Agile for App Dev
Thinking led by CIO "Move to
Cloud" Traditional
Procurement weakening
More Agile, Less Change
resistant**


digital

80% Change 20% Run

**AppDev leads decisioning
Doesn't own hardware
Refresh doesn't exist
All Agile for App Dev
Thinking led by CIO "Move to
Cloud"
Cloud Centric "Marketplace"
Procurement
Embraces Change, Very Agile**

2045

?




SNMP / traps

or

Daylight Savings





AIOps



EdgeOps
AMRS

EPS
every ip interface globally



APAC
EdgeOps

EMEA
EdgeOps



algorithms we use



BUZZ WORD BINGO

MACHINE LEARNING

OPERATIONAL INTELLIGENCE

BIG DATA ANALYTICS

STREAMING ANALYTICS

ARTIFICIAL INTELLIGENCE

NEURAL NETWORKS

IT OPERATIONS ANALYTICS

BLENDED ANALYTICS

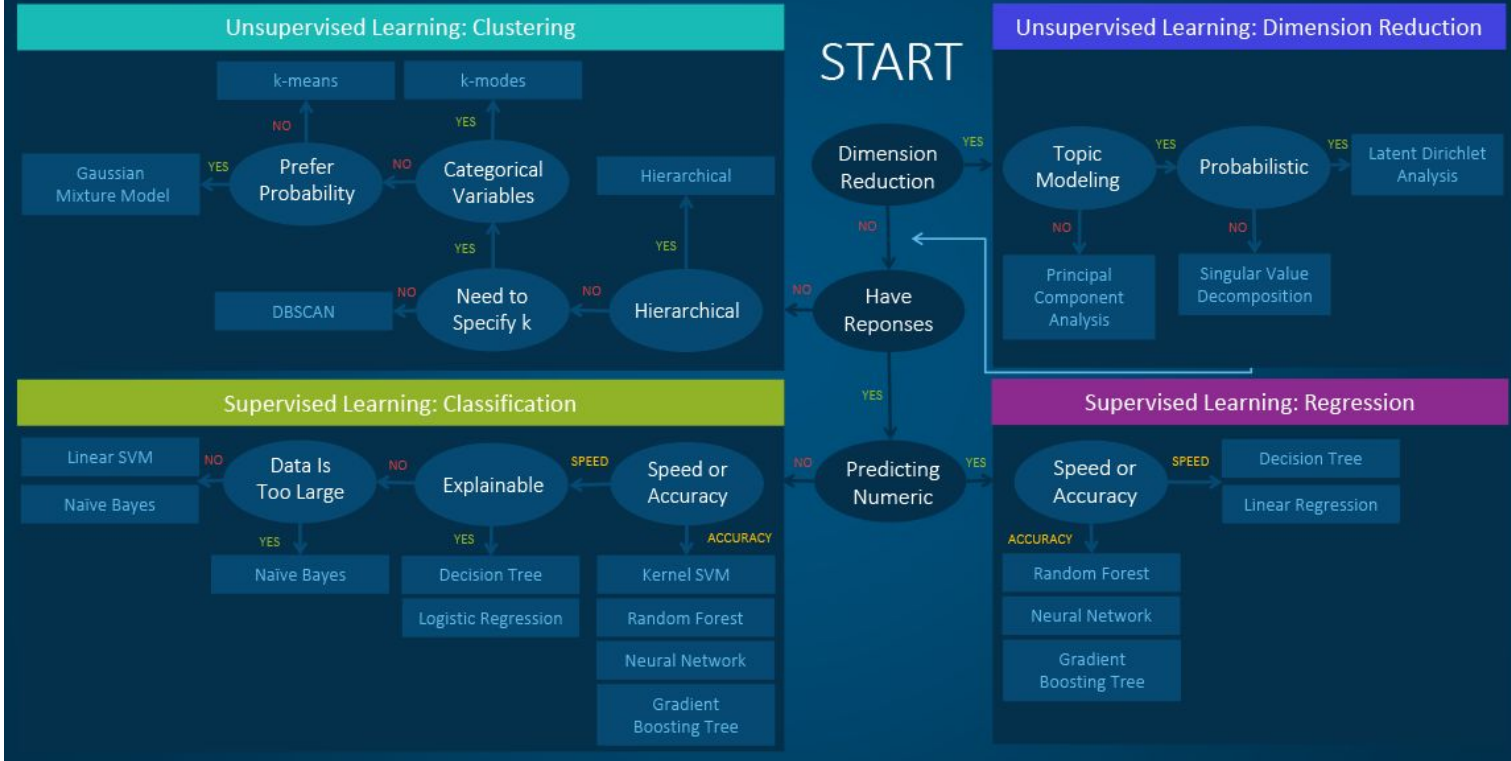
DATA SCIENCE

BIG DATA

This slide courtesy our Chief Scientist Dr. Rob Harper -- Do check out his great 3-part blog on Machine Learning in Moogsoft AIOps: <https://www.moogsoft.com/author/robharper/>




Machine Learning Algorithms Cheat Sheet




By [Hui Li](#) on [Subconscious Musings](#) April 12, 2017



Simple Neural Network




Deep Learning Neural Network



This slide courtesy our Chief Scientist Dr. Rob Harper -- Do check out his great 3-part blog on Machine Learning in Moogsoft AIOps: <https://www.moogsoft.com/author/robharper/>







This slide courtesy our Chief Scientist Dr. Rob Harper -- Do check out his great 3-part blog on Machine Learning in Moogsoft AIOps: <https://www.moogsoft.com/author/robharper/>




regression



classification



clustering



This slide courtesy our Chief Scientist Dr. Rob Harper -- Do check out his great 3-part blog on Machine Learning in Moogsoft AIOps: <https://www.moogsoft.com/author/robharper/>




classification

supervised

✓ Walks

✓ Quacks

✓ Swims



DUCKS



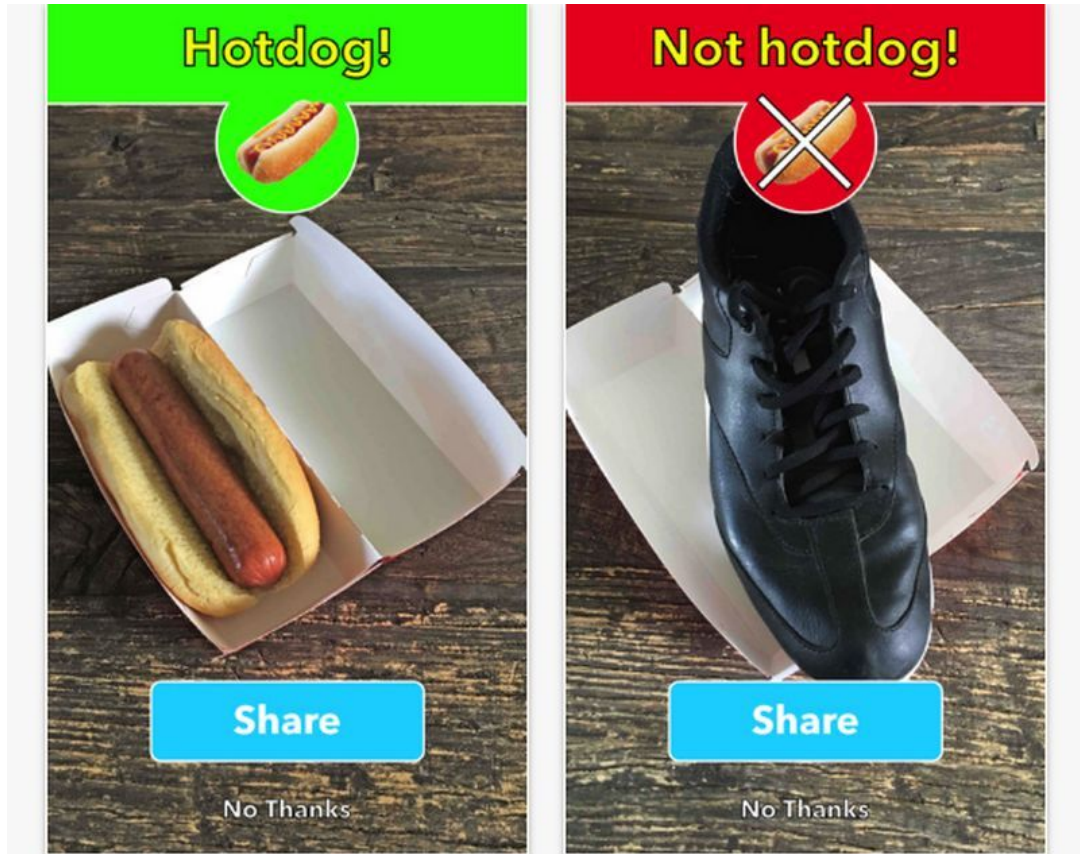
NOT DUCKS

“learn by example” approach. Supervised learning systems need to be given examples of what is “good” and what is “bad”

This slide courtesy our Chief Scientist Dr. Rob Harper -- Do check out his great 3-part blog on Machine Learning in Moogsoft AIOps: <https://www.moogsoft.com/author/robharper/>



classification



clustering

unsupervised



Patterns that you didn't know existed prior. Recommender systems rely heavily on these techniques.

supervised machine learning

"hot dog?" "not hot dog?"



Threshold value Anomalies Average



Untitled graph

1h 3h 12h 1d 3d 1w custom

Line

Actions



20.0k Milliseconds

10.0k

70.5

22:45 4:12 / 6:01

23:45






algorithms we use



SethBling mar i/o

neural nets




<https://www.youtube.com/watch?v=qv6UVOQ0F44>


lua code: <https://pastebin.com/ZZmSNaHX>



k-means clustering



matrix factorization



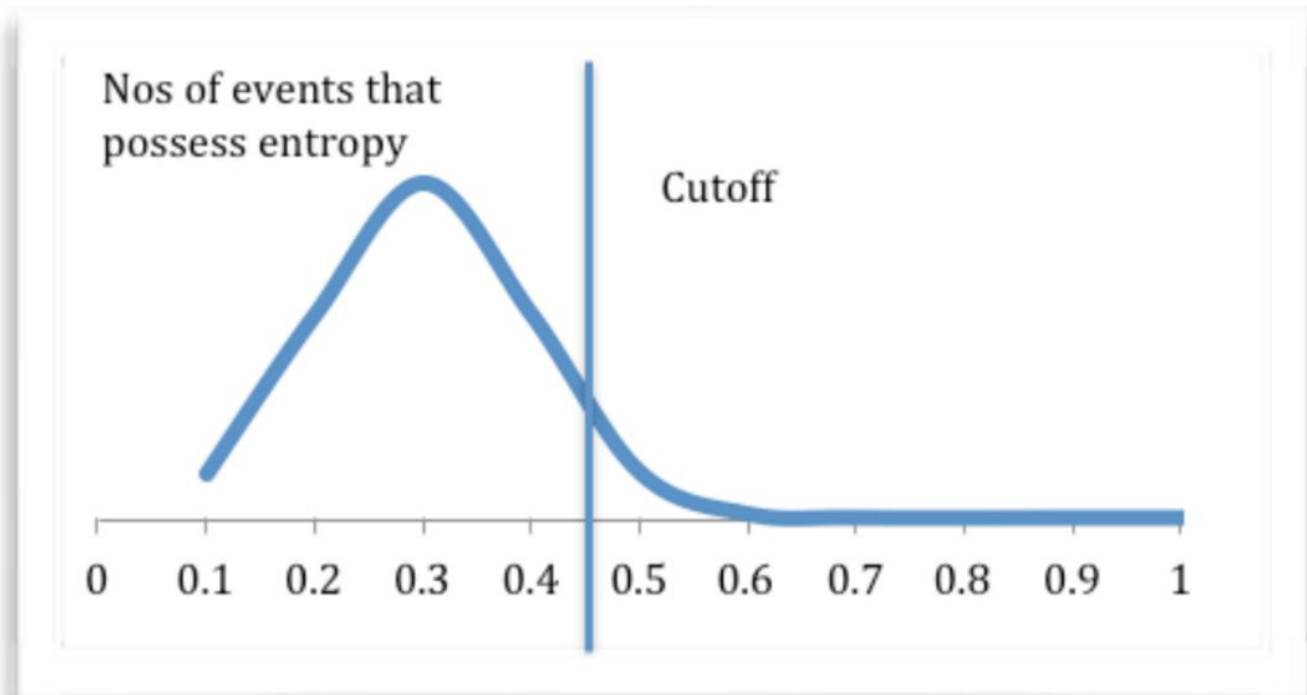
shannon entropy



$$S = \log_2 2^N = N = 2$$



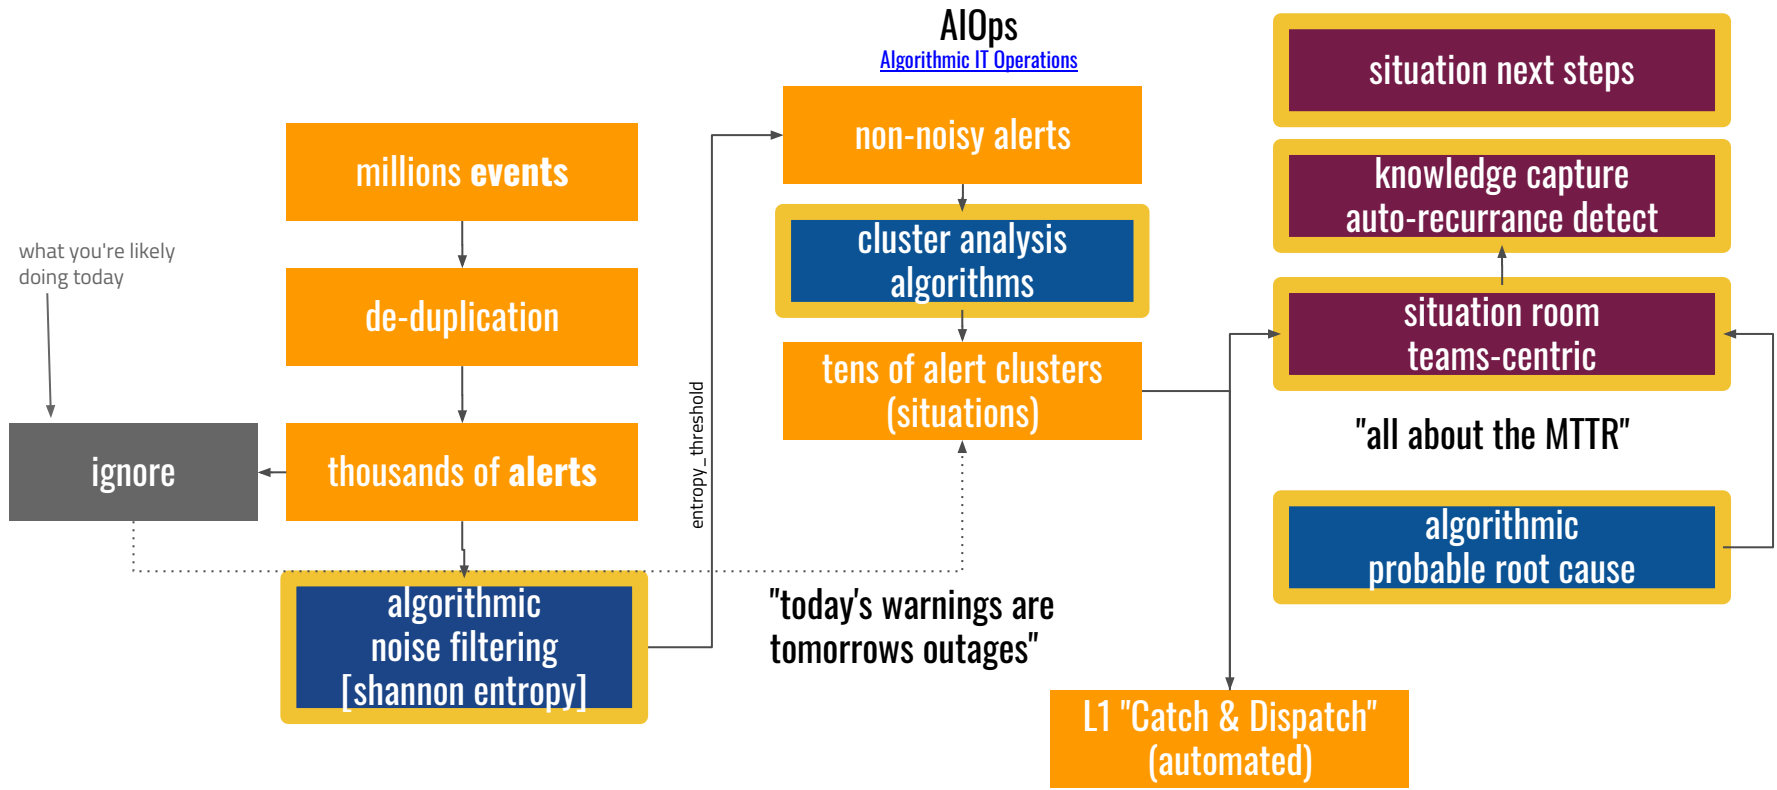
typical entropy distribution





algorithmic workflow





...speaking of classification



Pets

Legacy
Infrastructure

Cattle

Cloud-Friendly
Infrastructure




fault vs audit

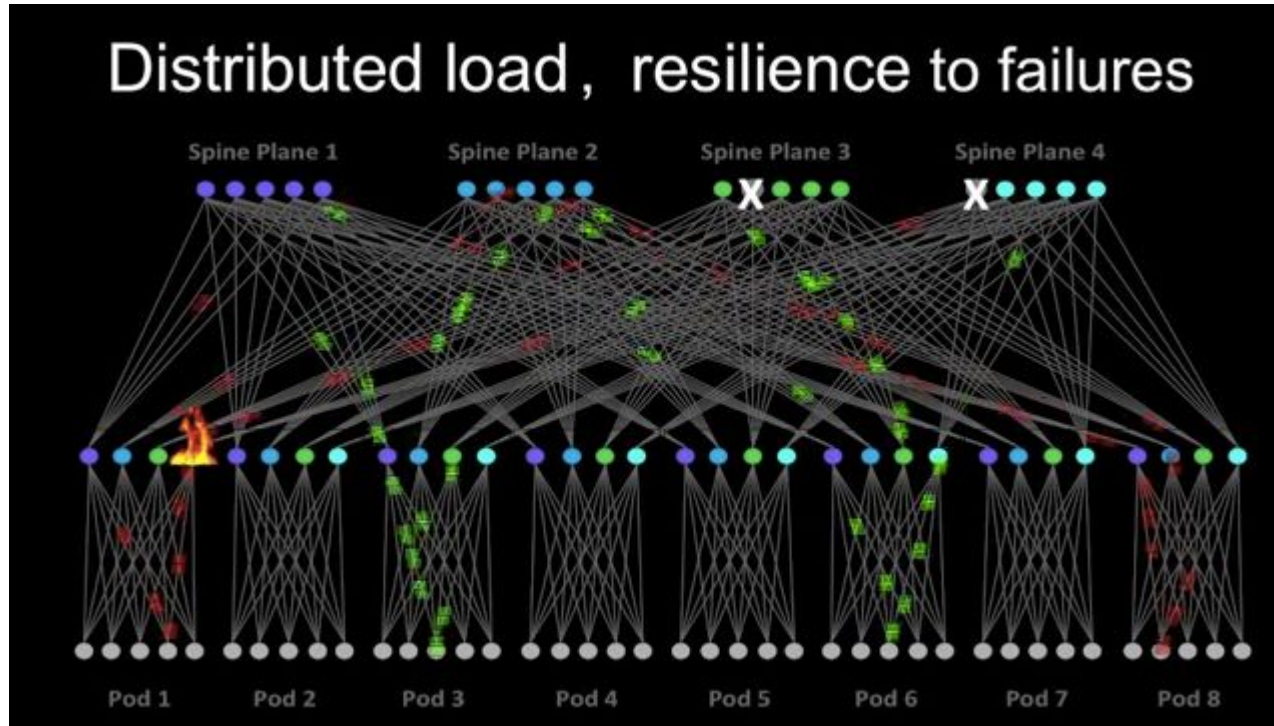
fix → optimize



monitoring fail-around



weld the datacenter doors shut



<lofty_^{non-technical}_tangent>



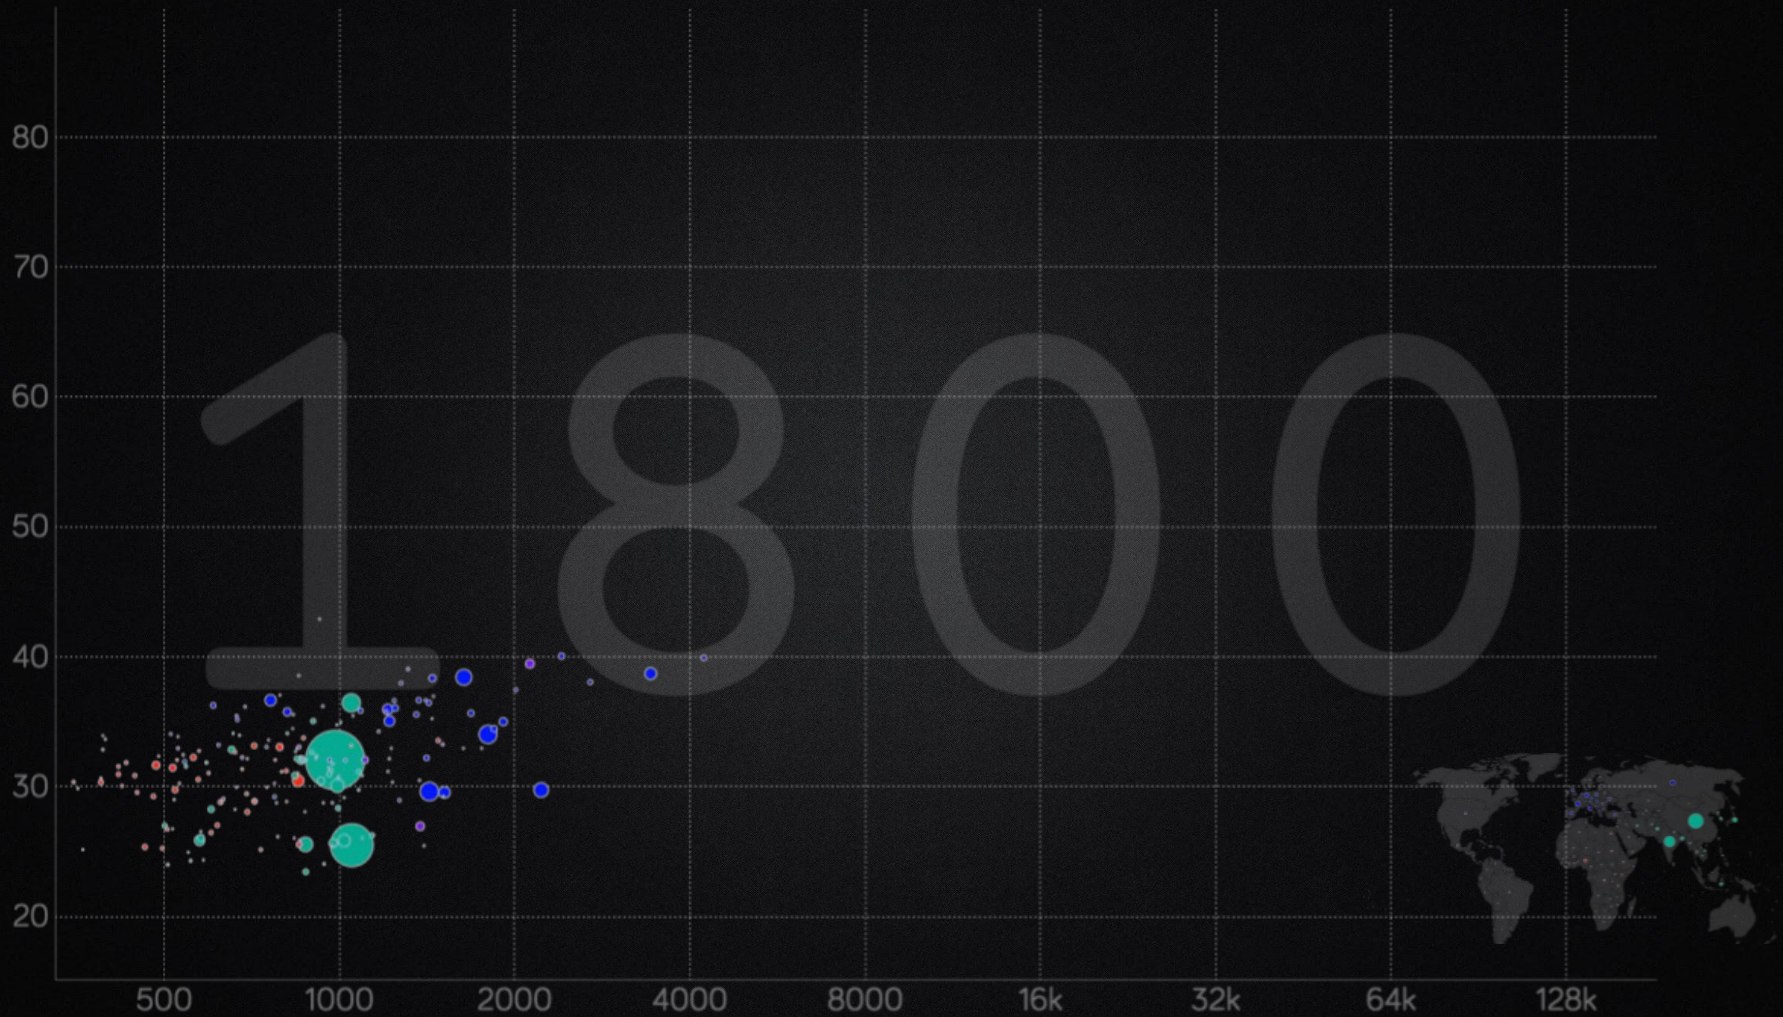


THE

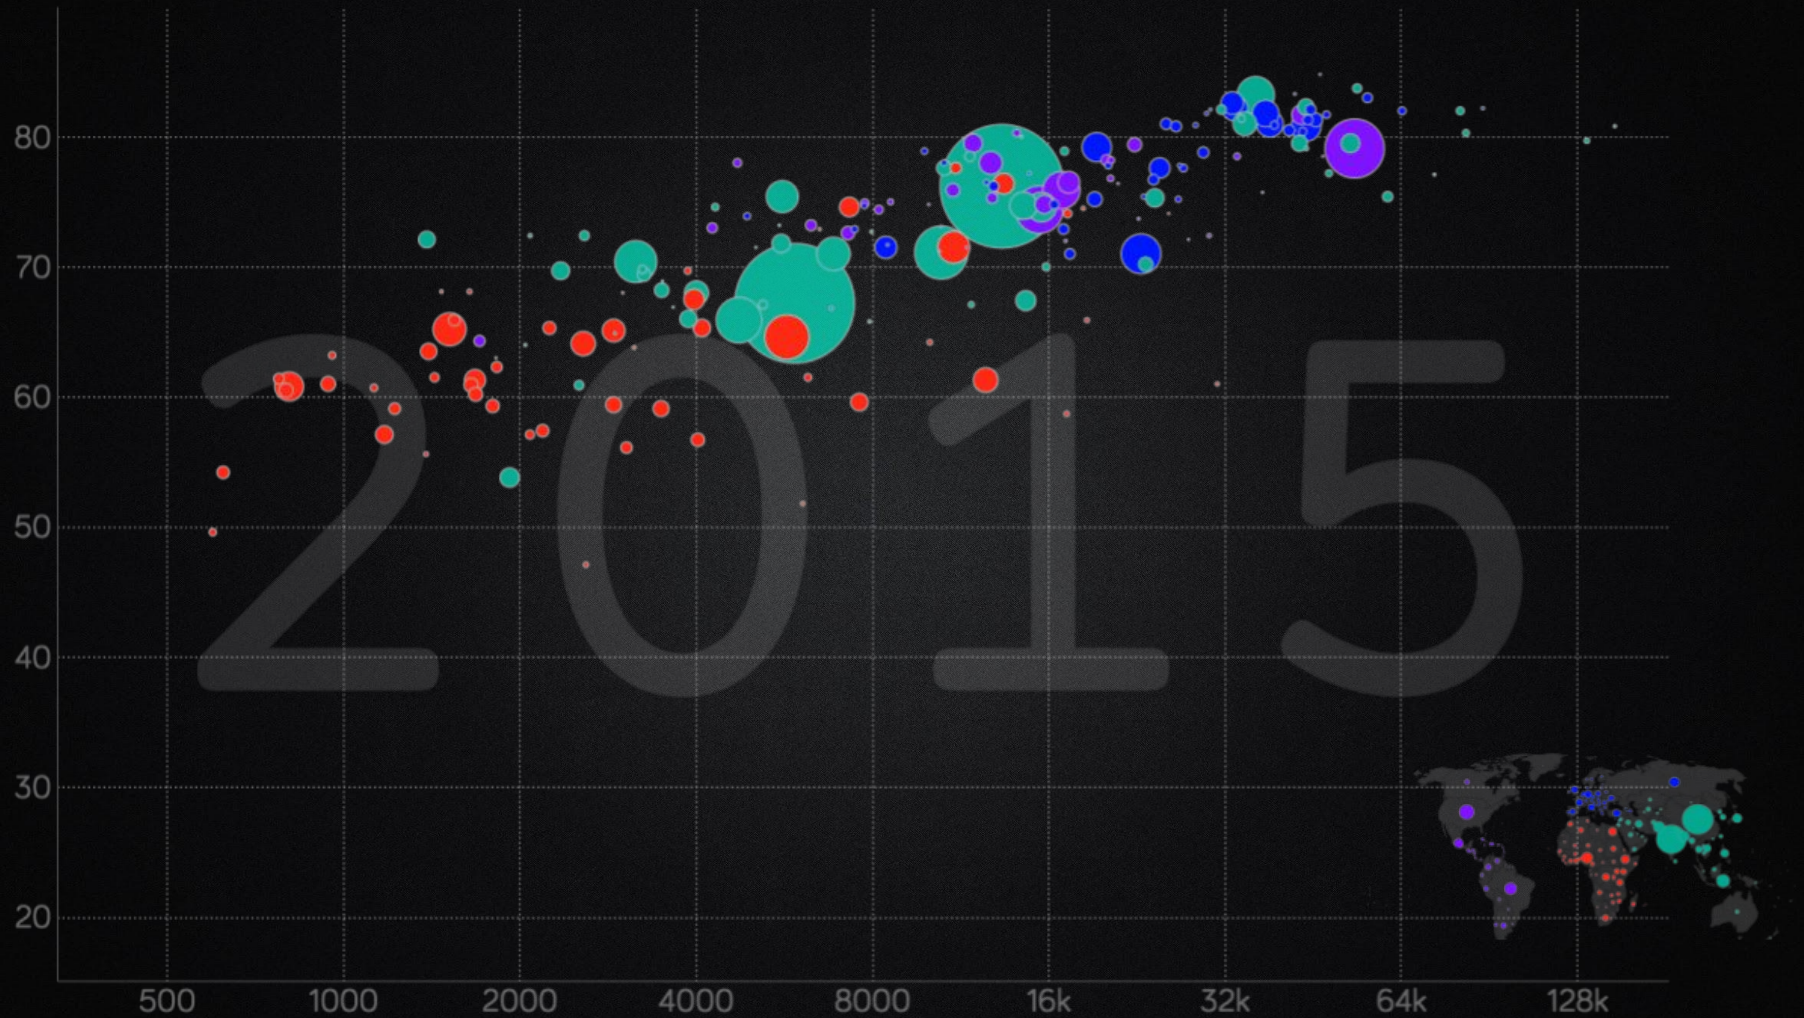
**INDUSTRIAL
REVOLUTION**

1.0

Life expectancy, years



Life expectancy, years



BBC FOUR



An aerial photograph of a city, likely Los Angeles, showing a dense urban landscape with a grid of streets and a prominent highway interchange in the lower-left. The sky is filled with large, white, fluffy clouds, and the overall lighting is bright and somewhat hazy, suggesting a clear day with high contrast.



**WE FACE
EXISTENTIAL
THREATS
TO OUR PROGRESS**







**THE
INDUSTRIAL
REVOLUTION 2.0
CAN HELP SAVE US**

v1.0



v2.0






**COMPLEXITY IS
THE PRINCIPAL
THREAT
TO THE REVOLUTION**


</lofty_{non-technical}_tangent>



Simple Neural Network



Deep Learning Neural Network



```

1  {
2      # The following defines the input neurons in the net
3      # and which activities will prompt a reinforcement learning
4      # step, or a brain to be shot!
5
6      inputs      : [ "source","description","severity","manager" ],
7      learn_queues : [ "manual_create","rated","merge_create","split_create","annotated","diagnosed","refined"],
8      unlearn_queues : [ "rated","split" ],
9      rating_threshold: 3,
10
11
12     # When training a brain, how many negative alerts to
13     # train with positive hits
14
15     alert_pool_size : 100,
16
17
18     # Ok, brain configuration
19
20
21     # The followings are the default values that do not need to be specified,
22     # and in fact not recommended to change by customers
23     neurons      : 7,          # number of hidden neurons is dynamically set
24     layers       : 2,          # number of hidden layer is default to 1
25     learning_rate : 0.15,
26     activation_func : "Hyperbolic",
27     epochs       : 500,
28
29
30     # Match strategy can be "Best" (closest to 1) or "Min_Error" (closest to trained result)
31     # Match strategy can be omitted, in which case it is default to "Min_Error"
32
33     # match_strategy      : "Min_Error",
34
35     precision            : 92.0,      # % precision to train for
36     tolerance            : 0.3,      # % tolerance of error in match
37     exact_match          : false,
38
39 }

```

theory



applied

sharing | giving back



