FAST, RELIABLE, CATASTROPHICALLY FAILING?

SAFELY AVOIDING BEHAVIORAL INCIDENTS IN COMPLEX PRODUCTION ENVIRONMENTS

WHO?

- Software Engineer
- Working on Data Science teams as the engineer
- Exposed to “proper science”
- Put this model/data product into prod

@rmn

fuzzbox

 hyperr

 hieroglyph

 sonar

 supercollider

 errorsmith

 echoplex

@rmn

@rmn

WHAT ARE WE TALKING ABOUT

AND WHAT AREN'T WE TALKING ABOUT

@rmn

COMPLEXITY
EMERGENT BEHAVIOR
UNKNOWNS
DISCOVERED AT SCALE

SUBTLETY
DORMANT BEHAVIOR
FORENSIC INCIDENT DATA
DATA AS A NEW THREAT VECTOR
IS “ACTING WEIRD” AN INCIDENT?

@rmn

LETS TALK ABOUT STUFF
THIS TALK IS NOT ABOUT

@rmn

PREMISE

COMPLEXITY AND FAILURE
GO HAND IN HAND

RELIABILITY AND ROBUSTNESS
COME FROM DIRECT EXPERIENCE
WITH FAILURE

@rmn

COMPLEXITY

SIDE EFFECT OF SUCCESS

@rmn

COMPLEXITY

ESSENTIAL

ACCIDENTAL

@rmn

COMPLEXITY

ACCIDENTAL

@rmn

COMPLEXITY

ESSENTIAL

@rmn

ACCEPT

SOLUTION TO COMPLEXITY
IS NOT SIMPLICITY

@rmn

ACCEPT

COMPLEXITY
HAS TO BE EMBRACED AND MANAGED

@rmn

COMPLEX SYSTEMS

FAILURE

@rmn

FAILURE IN COMPLEX SYSTEMS

HAZARDOUS
LAYERED DEFENSES BUILT OVER TIME
CATASTROPHE INVOLVES MULTIPLE FAILURES
ERROR DETECTION IS HARD

@rmn

COMPLEX SYSTEMS

NO ROOT CAUSE
OPERATORS HAVE DUAL ROLE

@rmn

2 THREATS TO AVAILABILITY

THE SOFTWARE CHANGES
THE ENVIRONMENT CHANGES

@rmn

THE ENVIRONMENT CHANGES

NETWORK LATENCY
RESOURCE CONTENTION / NOISY NEIGHBOR
DISK IS FULL
TIME IS WRONG

@rmn

IN TERMS OF DUAL ROLE

THE SOFTWARE CHANGES

YOU ARE A PRODUCER

@rmn

2 THREATS TO AVAILABILITY

THE ENVIRONMENT CHANGES

DEFENDER
IT IS DOING SOMETHING YOU DIDN'T ANTICIPATE

@rmn

OPERATOR ROLE

ALLOWS YOU TO BE A PARTICIPANT IN SYSTEM DURING AN INCIDENT

AAAAND WE'LL TAKE THIS AWAY LATER TOO 🤓

@rmn

MOSTLY

WHAT AREN'T WE TALKING ABOUT

@rmn

WHY TELL IS NOT THE STUFF RAMIN?!?

MEANS OF DEALING WITH THIS ARE BECOMING WELL UNDERSTOOD

BECOMING A SOLVED PROBLEM

@rmn

LETS REVISIT COMPLEXITY

@rmn

THE AMAZON/NETFLIX KIND

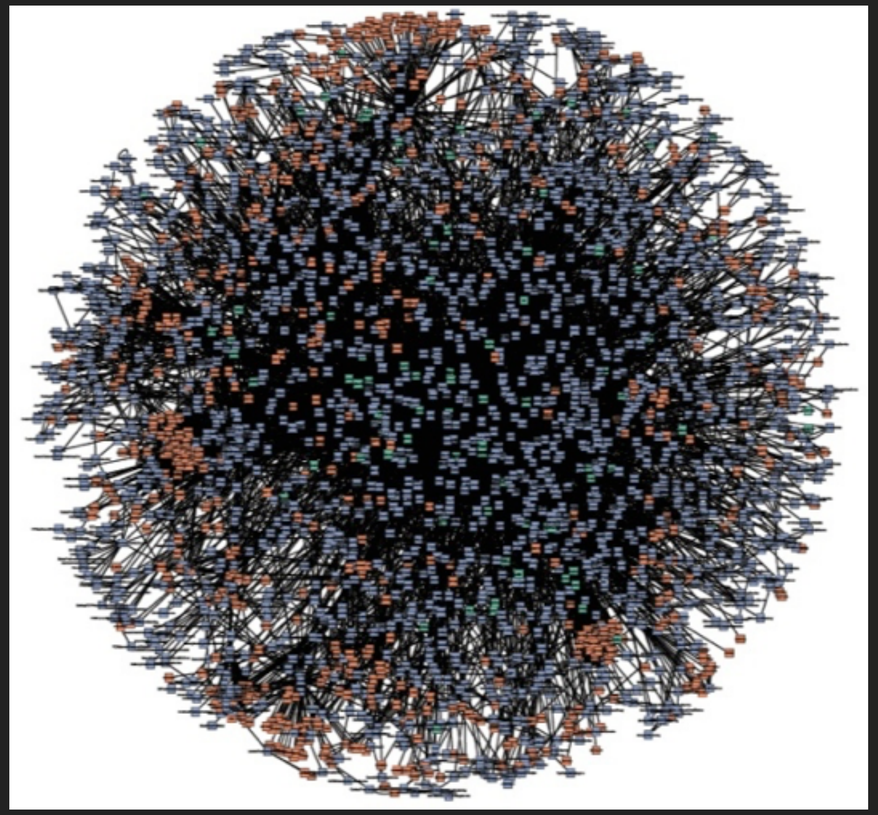

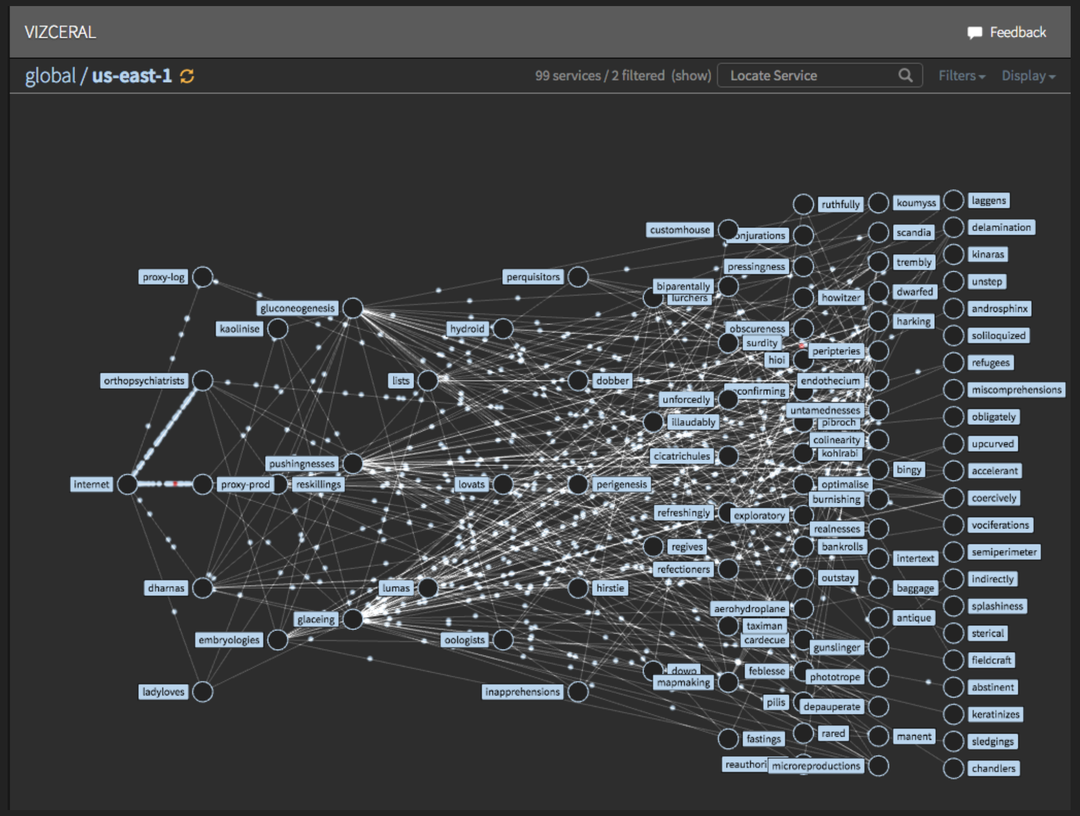

@rmn

FOREGO CORRECTNESS

ADOPT SAFETY

@rmn

RECOGNIZE HAZARDOUS SHARP EDGE

@rmn

SHARP EDGE

SAFETY APPARATUS IS BUILT INTO SYSTEM

HUMANS AND CULTURE

@rmn

3 CONCEPTS

TEST IN PROD
PROGRESSIVE DELIVERY
ERROR BUDGETS

@rmn

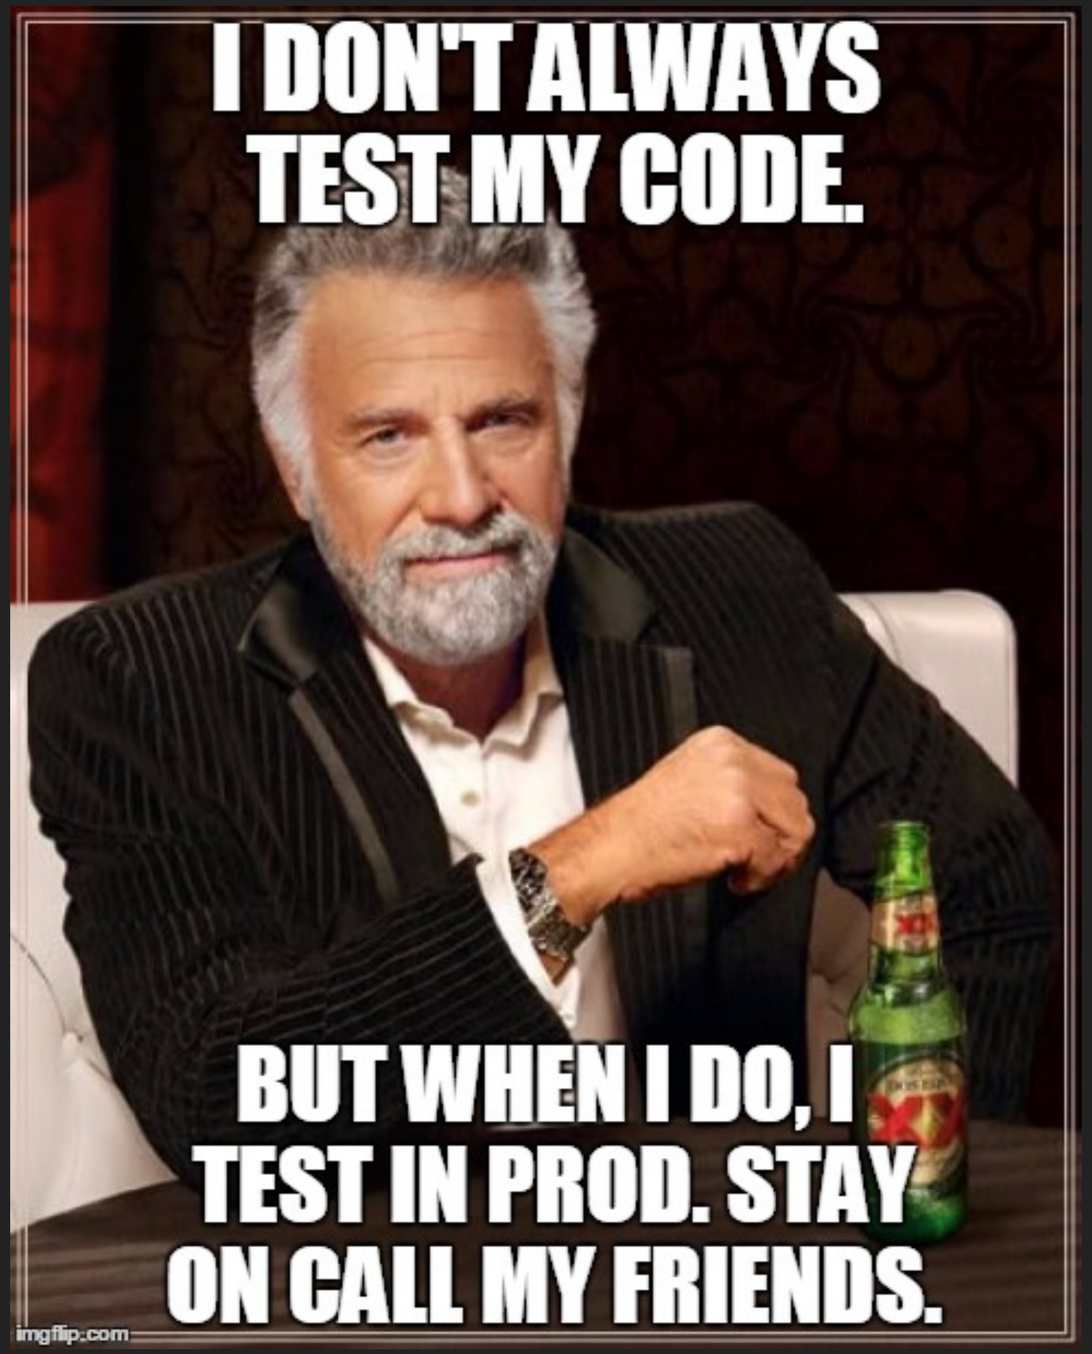
I DON'T ALWAYS
TEST MY CODE.
BUT WHEN I DO, I
TEST IN PROD. STAY
ON CALL MY FRIENDS.
imgflip.com

@rmn

TEST IN PROD DOESN'T MEAN RELEASE WITHOUT TESTING

@rmn

TESTING IN PROD MEANS
EXTENDING THE SOFTWARE
DEVELOPMENT LIFECYCLE
BEYOND RELEASE

@rmn

"REAL USERS, REAL TRAFFIC, REAL SCALE, REAL UNPREDICTABILITIES"

@MIPSYTIPSY

PROGRESSIVE DELIVERY

@rmn

"PROGRESSIVE DELIVERY IS CONTINUOUS DELIVERY WITH FINE-GRAINED CONTROL OVER THE BLAST RADIUS."

James Governor, RedMonk (@monkchips)

SEPARATE DEPLOY FROM RELEASE

WHY?

THINGS ARE DEFINITELY GOING TO GO WRONG IN WAYS YOU DIDN'T ANTICIPATE.
EMBRACE IT

@rmn

FEATURE FLAGS

TARGET SPECIFIC USERS FOR NEW “FEATURES”
ABILITY TO TOGGLE EXPOSURE ON/OFF

@rmn

CANARY

EXPOSE SOME % OF LIVE TRAFFIC TO A NEW SERVICE
MONITOR KEY BUSINESS METRICS FOR THAT POPULATION
A/B TEST OUTCOME OF NEW DEPLOYMENT
WIDER RELEASE WHEN YOU ARE COMFORTABLE

@rmn

ERROR BUDGETS

@rmn

ERROR BUDGETS
OPPORTUNITY FOR LEARNING

@rmn

IT IS NOT YOUR JOB TO CREATE INFINITELY RELIABLE SOFTWARE

WHAT IF YOU COULD CREATE MORE BUSINESS VALUE BY LETTING SOFTWARE BE MORE BROKEN

@rmn

MAYBE THE NATURAL
DISTRIBUTION OF FAILURE
HAS SPARED YOU

@rmn

YOU MIGHT HAVE SOME 9S TO PLAY WITH

@rmn

PERMIT AUDACITY
WHEN AUDACITY
CAPITAL AVAILABLE

@rmn

RECOGNIZE SHARP EDGE

VOCABULARY FOR MANAGING COMPLEXITY SAFELY

TEST IN PROD
PROGRESSIVE DELIVERY
ERROR BUDGETS

@rmn

EXPERIMENT

DELIBERATELY EXPLORE WEIRD BEHAVIOR (CHAOSENG)

TRY NEW THINGS INSIDE YOUR BUDGET

PERMIT AUDACITY WHEN AUDACITY CAPITAL AVAILABLE

ALLOW AN ACCIDENTAL “OVERAGE” OF SLA TO BE YOUR PLAYGROUND

YOU HAVE HEADROOM TO TAKE RISKY CHANGES

@rmn

WHAT WE ARE TALKING ABOUT

@rmn

SUBTLE

OBLIQUE

@rmn

MODELS IN PRODUCTION

LATENT CATASTROPHOIC BEHAVIOR

WHAT ARE THE FAILURE COMPONENTS?

ARE THESE OUTAGES?

@rmn

MODEL ARTIFACT IS COMPLEX

BUT NOT COMPLEX DEPENDENCIES
COMPLEX RESPONSE TO INPUT

DATA NEW VECTOR OF FAILURE

@rmn

~~2~~ 3 THREATS TO AVAILABILITY?

THE SOFTWARE CHANGES
THE ENVIRONMENT CHANGES
THE DATA WAS UNANTICIPATED

@rmn

THE DATA IS HAZARDOUS

DATA ISN'T A TRADITIONAL COMPONENT IN A COMPLEX SYSTEM

@rmn

AN INCIDENT

@rmn

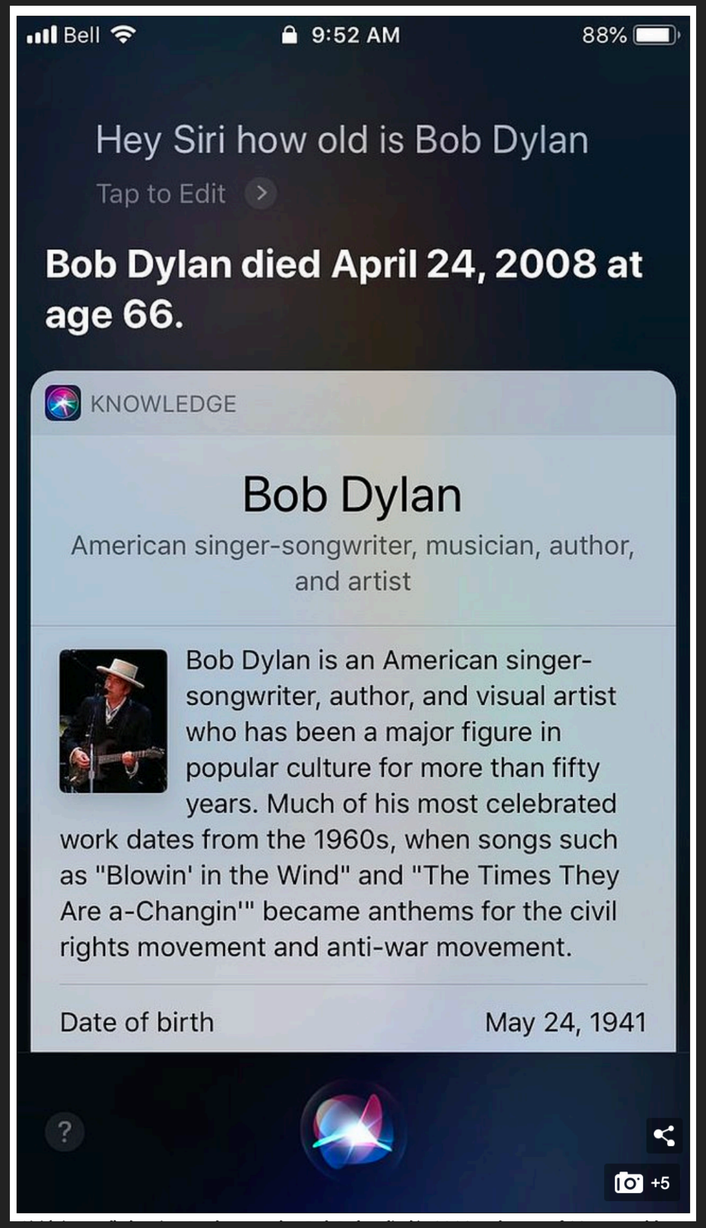
Bell
9:52 AM
88%
Hey Siri how old is Bob Dylan
Tap to Edit
Bob Dylan died April 24, 2008 at age 66.
KNOWLEDGE
Bob Dylan
American singer-songwriter, musician, author, and artist
Bob Dylan is an American singer-songwriter, author, and visual artist who has been a major figure in popular culture for more than fifty years. Much of his most celebrated work dates from the 1960s, when songs such as "Blowin' in the Wind" and "The Times They Are a-Changin'" became anthems for the civil rights movement and anti-war movement.
Date of birth
May 24, 1941
+5

HOW ML WORKS

GET LABELLED DATA
SLICE IT UP
TRAIN ON A SLICE
COMPARE TO OTHER SLICE
TWEAK KNOBS
LOOKS GOOD
DEPLOY

NEW DATA COMES IN
IT INTERPRETS AND RESPONDS

HOW ML WORKS

GET LABELLED DATA
SLICE IT UP
TRAIN ON A SLICE
COMPARE TO OTHER SLICE
TWEAK KNOBS
LOOKS GOOD
DEPLOY

NEW DATA COMES IN
IT INTERPRETS AND RESPONDS

YOUR MODEL KILLS BOB DYLAN

IS IT A BUG?
IT IS AN INCIDENT!
IS IT AN OUTAGE?

@rmn

BEHAVIORAL OUTAGES

DATA DATA DATA DATA DATA DATA DATA DATA DATA DATA DATA DATA DATA DATA REPLACES CODE

@rmn

DATA REPLACES CODE

@rmn

DATA COMPLEXITY
REPLACES CODE COMPLEXITY

@rmn

IS THIS GREY FAILURE?

@rmn

WHAT?

WHAT ON EARTH DO I DO WITH THIS THING

@rmn

TRADITIONAL CHARACTERISTICS OF AN INCIDENT

IT'S SLOW
IT'S DOWN
IT'S INTERMITTENTLY AVAILABLE
IT'S DOING SOMETHING WEIRD
IT'S MAKING SOMETHING ELSE ACT WEIRD

@rmn

CAN BE REASONED ABOUT
WE CAN DEBUG IT WHILE HAPPENING
WE CAN INSTRUMENT CONVENTIONALLY FOR REDUCING MTTD, MTTR (THIS MUST NEVER EVER EVER HAPPEN AGAIN)

(DON'T DRAG ME ON MTTD/MTTR)

@rmn

WE LEARN ABOUT AND
IDENTIFY WHAT IS HAPPENING
WHILE IT IS HAPPENING

@rmn

I'M NOT SAYING:

THE INCIDENT IS EASY TO UNDERSTAND
THE DATA IS SHALLOW
THERE IS A ROOT CAUSE

@rmn

I AM SAYING:

INCIDENT DATA CAN BE GATHERED DURING INCIDENT

IT CAN (LIKELY) BE CIRCUMVENTED AND FIXED DURING THE INCIDENT

PEOPLE WHO ARE PART OF SYSTEM HELP CONTRIBUTE TO DISCOVERY, LEARNING, AND RESOLUTION OF THE INCIDENT DURING THE INCIDENT

@rmn

Hindsight Bias

Before the Accident

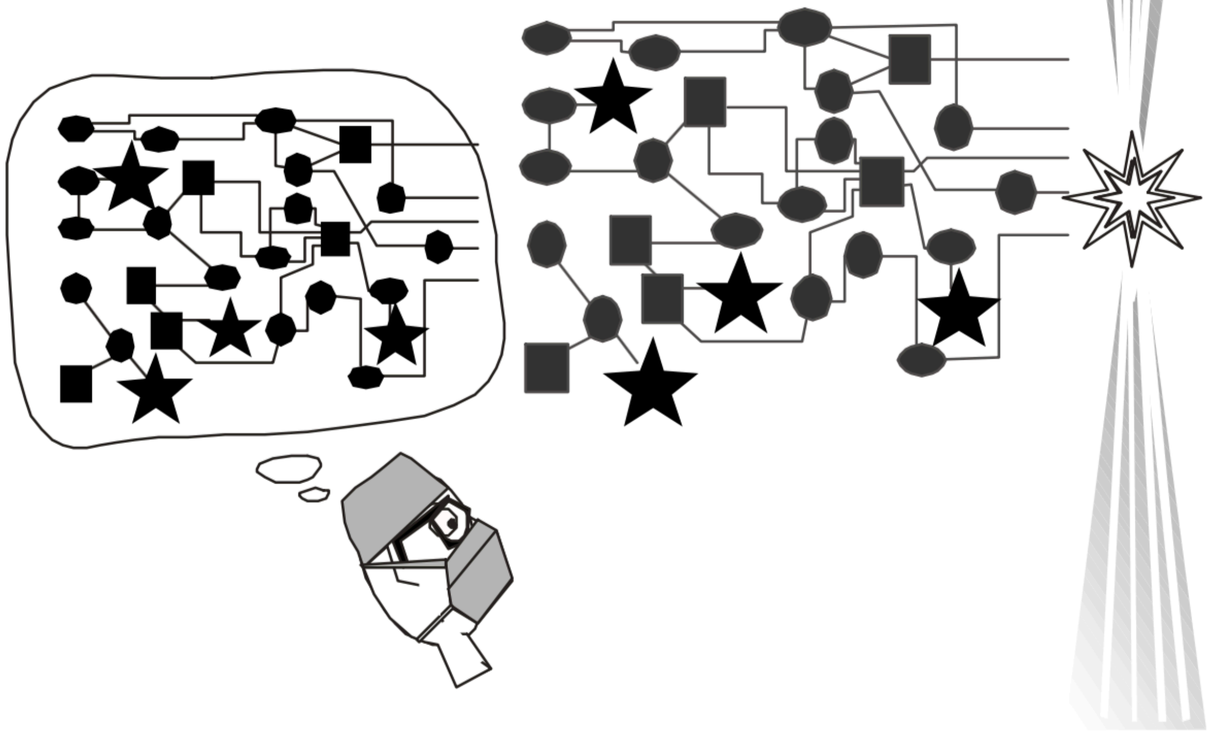

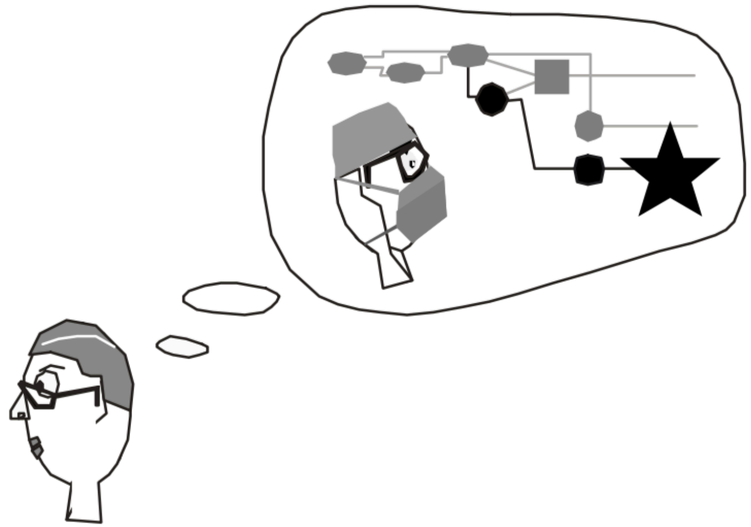

After the Accident

Cook (1999). A Brief Look at the New Look in error, safety, and failure of complex systems. (Chicago: CtL).

CHARACTERISTICS OF SUBTLE/ML INCIDENT

IT'S FASTER
IT'S AVAILABLE
IT'S STABLE
IT'S DOING SOMETHING REALLY WEIRD
IT'S DOING SOMETHING COMPLETELY UNIMAGINABLE
WE CANT UNDERSTAND WHY

@rmn

IT'S FASTER
IT'S MORE AVAILABLE
IT'S STABLE

@rmn

DIFFICULT OR IMPOSSIBLE TO REASON ABOUT

INFORMATION TO LEARN AND DEBUG IS NOT PRESENT DURING INCIDENT

TRADITIONAL APPROACH TO OBSERVABILITY & INSTRUMENTATION DOES NOT APPLY

DETECTION IS HARD AS BEHAVIOR IS WHAT WE PLANNED

PARTICIPANTS ARE EXCLUDED/PROHIBITED FROM LEARNING DURING INCIDENT

@rmn

WE CAN'T ALWAYS IDENTIFY ANYTHING IS HAPPENING
WHILE IT IS HAPPENING

ONCE DETECTED OPTION IS TO STOP / ROLLBACK / UNDO

FORENSICS CAN ONLY HAPPENED AFTER MITIGATION

SOMETIMES IMPOSSIBLE TO REPRODUCE TO LEARN

@rmn

WHAT IS THE METRIC OR DASHBOARD YOU BUILD FOR DETECTING "OUR APPLICATION KILLED BOB DYLAN"

@rmn

SITUATION #1

IT'S STABLE

PIPELINE JUNGLE
STALE DATA WAS USED SO NOTHING CHANGED
SERVING STALE, IRRELEVANT INFERENCES
DIDN'T IMPROVE ANY KPI

@rmn

INCIDENT #1

MODELED AFTER AN OLD ‘VERSION’ OF THE BUSINESS

INCIDENT? OUTAGE?

@rmn

SITUATION #2

IT'S FASTER

TRAINED INCORRECTLY WITH UNSTABLE DATA
DISTRIBUTION OF LABELS CHANGED
MODEL IGNORED NEW INPUTS AT INFERENCE TIME
FASTER RESPONSE TIME
HOORAY

@rmn

INCIDENT #1

INPUTS IGNORED

INCIDENT? OUTAGE?

@rmn

INCIDENT #3

IT'S STABLE

NO AUTOMATION OR REPRODUCIBLE BUILD PIPELINE
PRODUCTION ARTIFACT BUILT ON SCIENTISTS MACHINE
WRONG ARTIFACT BUNDLED
SIMPLE INTERFACE
WRONG ASSIGNMENT IN MARKETPLACE

BONUS INCIDENT: WHAT HAPPENED WHEN SCIENTIST LEFT COMPANY?

@rmn

INCIDENT #3

WRONG CODE

INCIDENT? OUTAGE?

@rmn

INCIDENT #4

IT'S FASTER

EXPERIMENTAL CODE PATH INCORRECTLY IMPLEMENTED
EVERYONE RECEIVED DEFAULT/FALLBACK DATA
DEFAULT RECOMMENDATIONS FOR EVERYONE
YAY!

@rmn

INCIDENT #4

IT DIDN'T DO ANYTHING

INCIDENT? OUTAGE?

@rmn

SAFETY MECHANISMS DO WORK

TEST IN PROD
VERIFICATION
CANARY
BLAST RADIUS

@rmn

BUT WE NEED MATURITY TO SHIFT LEFT IN THE ML SDLC LIFECYCLE

REPRODUCIBILITY
DATA VERSIONING
REPEATABLE PIPELINES
VARIED DATA
BIAS

@rmn

SKYNET IS HERE

IT'S JUST REPEATEDLY BUMPING INTO THE WALL

MODERN SAFETY TECHNIQUES WORK DO WORK FOR PERF REGRESSIONS
AS WE'VE SEEN PERF ISN'T A SIGNAL OF "SOMETHING WRONG"
WHEN THINGS START MURDERING, OUR ONLY HOPE IS TO ROLL BACK AND HOPE WE HAVE REPRODUCIBILITY
THE CATASTROPHE'S IN THIS TYPE OF COMPLEX SYSTEM ARE REALLY WEIRD
LUCKILY FOR NOW ITS PLAYFUL

@rmn

THANK YOU

@rmn