The Life-Changing Magic of Data Governance

Paula Griffin
Director of Product @ ZipRecruiter

paulag@ziprecruiter.com

Our Mission: Actively Connecting People to Their Next Great Opportunity

Marketplace of jobs and job seekers

110M+	2.8M+	#1
job seekers	businesses	rated job search app on Android and iOS

12 years of company history...

Now known as an **AI-driven marketplace**, acting as a personal recruiter for job seekers

Earlier, gained traction by posting jobs to many job boards

Began as a way to add screening questions to job applications

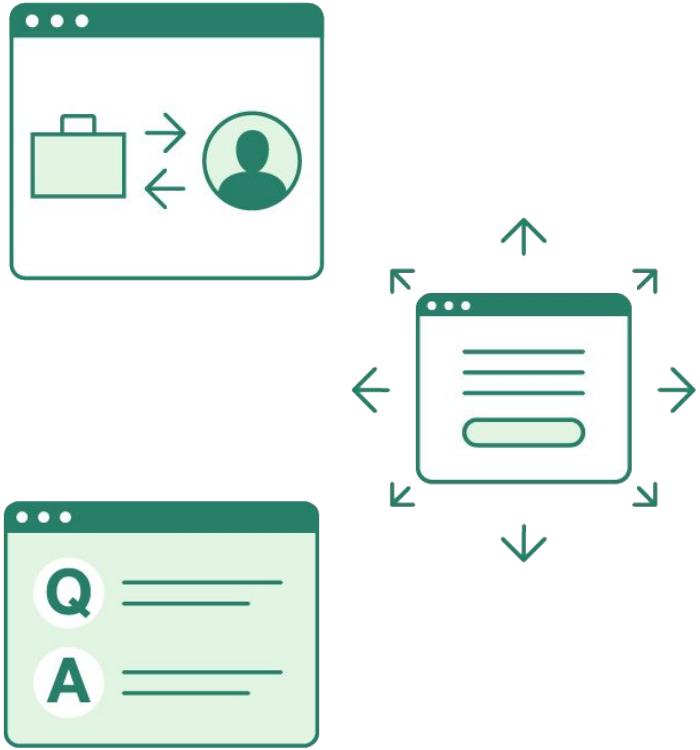

… leads to data pipelines like this

… and problems like this

There are three different versions of this table. Which do I use?

Why do numbers look weird when I join this other table?

I want to do more, but data is a bottleneck.

Where can I find XYZ data?

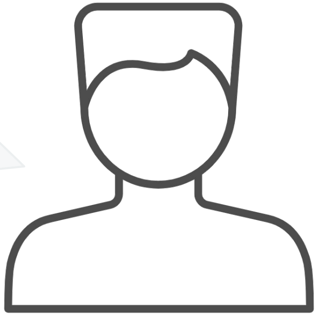

Our data was not built for machine learning

- No single source of truth for core business actions
- Impossible to track changes to data over time
- Questions that should take minutes take hours or days
- Lower visibility into the state of the product overall

The result: hard to validate which models to build, impossible to reproduce once you have them, and really difficult to measure the impact of

What about a data landscape that ✨sparks joy✨?

Recognized that as an ML-driven company, **data is everything.**

Needed a world with

- Easy validation of model ideas
- Reproducible once trained
- No time-traveling
- Easy to measure impact

Enter **data governance.**

We started at the bottom rung of governance maturity

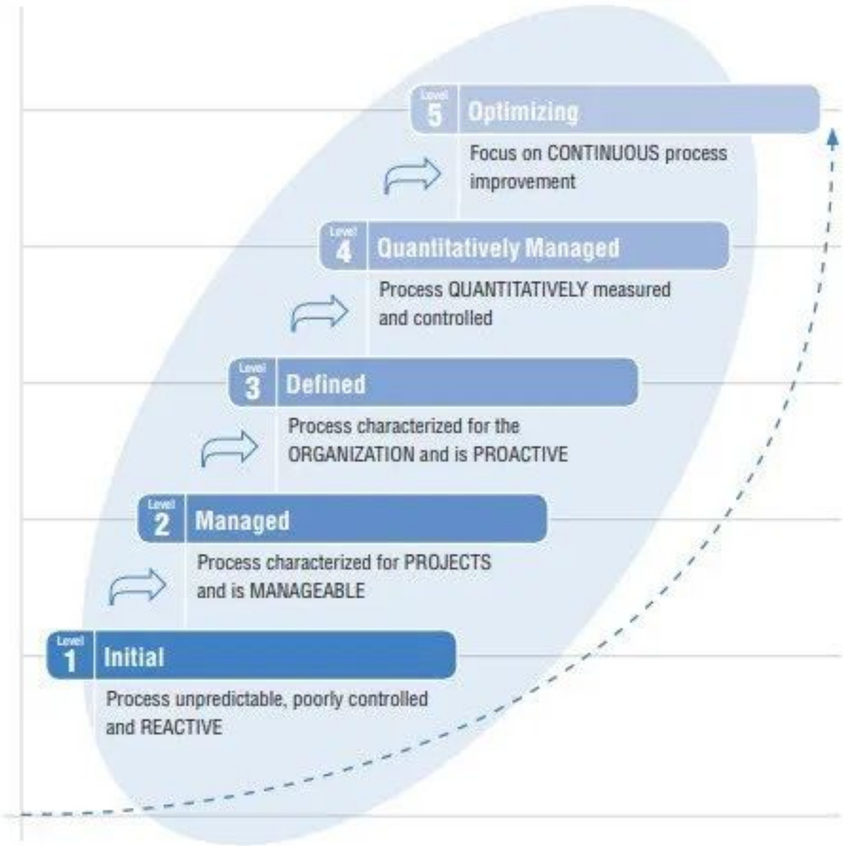

Six rules we've used in untangling the pipelines

1. Make data an unambiguous priority
2. Treat data products as real products
3. Set clear standards
4. Educate at all levels
5. Deliver value early
6. Commit to trust

1. Make data an unambiguous priority

People are policy

Approximately 25% of our product and engineering teams are devoted to data

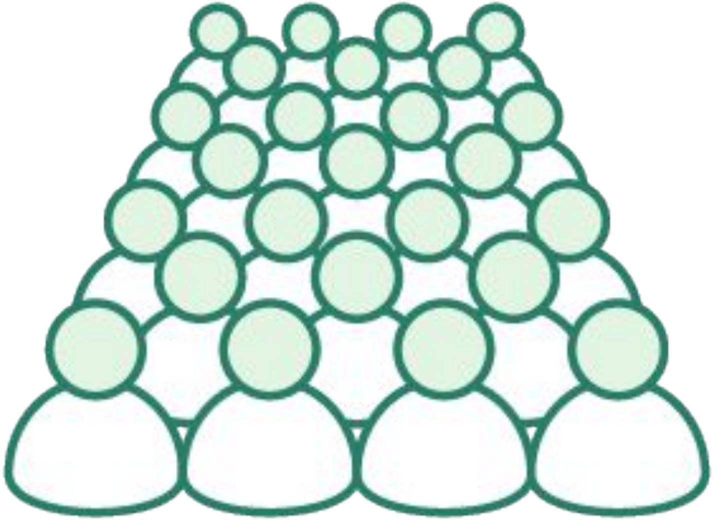

2. Treat data products as real products

We treat data as a specialization within product as well as engineering

- Data sets are products with customers
- Customers inform but do not dictate the roadmap

Define strategy for data collection and use

Specialized guides and templates to get off to a good start

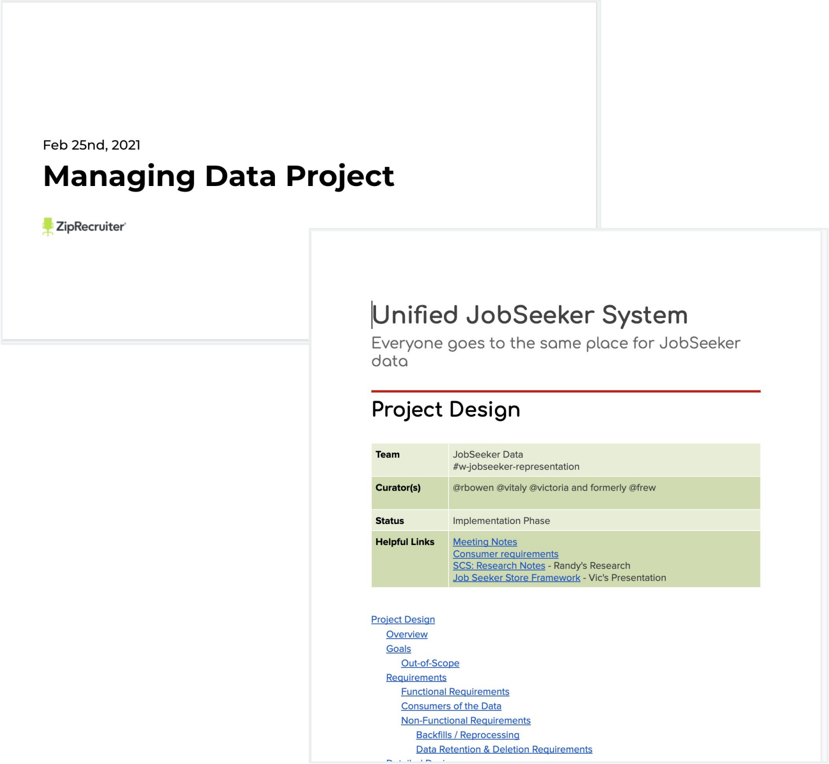
Feb 25nd, 2021

Managing Data Project

ZipRecruiter

Unified JobSeeker System

Everyone goes to the same place for JobSeeker data

Project Design

Team	JobSeeker Data #w-jobseeker-representation
Curator(s)	@rbowen @vitaly @victoria and formerly @frew
Status	Implementation Phase
Helpful Links	Meeting Notes Consumer requirements SCS: Research Notes - Randy's Research Job Seeker Store Framework - Vic's Presentation

3. Set clear standards

Aligned early on what our data needs to do for us

- For model development
- For general analytics use cases

Established shared vocabulary, best practices, and support tiers of different data sets.

Guarantees Ponds

EXAMPLE SLIDE

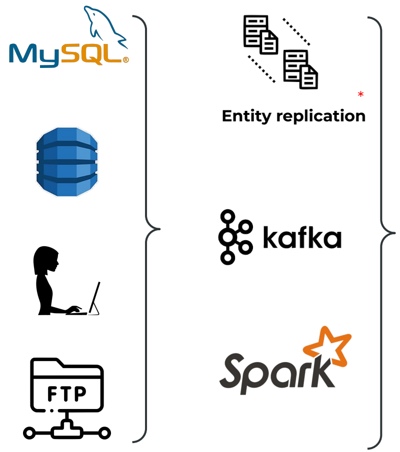

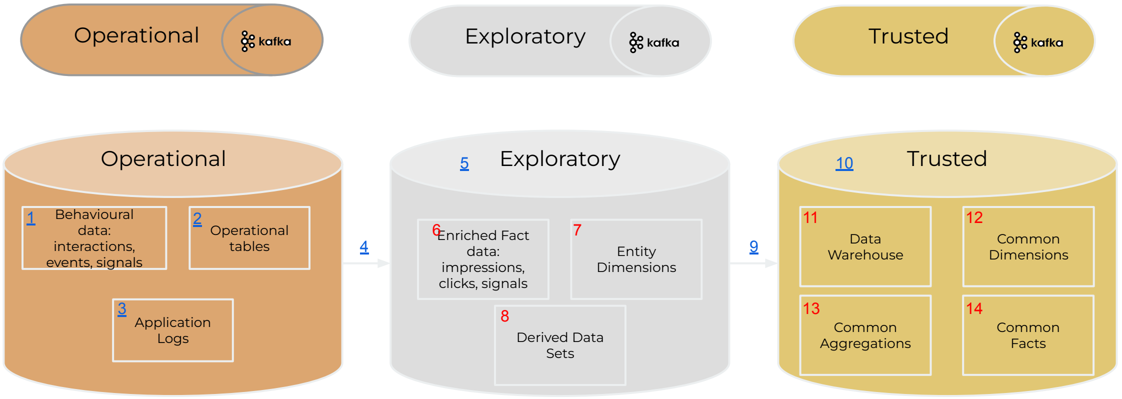

- Unstructured
- Not Verified
- Write optimized
- TTL of X days

- Schema Enforcement
- Filtered
- Cleaned
- Enriched with entity level business logic
- Denormalized across business entities: job category, Companies, campaign,
- PII, GDPR, CCPA
- Read optimized

- Guarantees: SLA, Completeness, Schema
- Self Served :
  - Discoverable
  - Star Schema
  - Read optimized
  - Denormalized
- Enriched with business logic

* should be replaced with CDC+Delta Lake

EXAMPLE SLIDE

What is Project Konmari, you ask?

The plan to organize and scale our data, using a data-lake-first approach.

4. Educate at all levels

From new grad to C-suite

- Tools and abstractions reduce toil, with detailed how-to guides
- Demos and office hours to empower all functions
- Product walkthroughs for executives

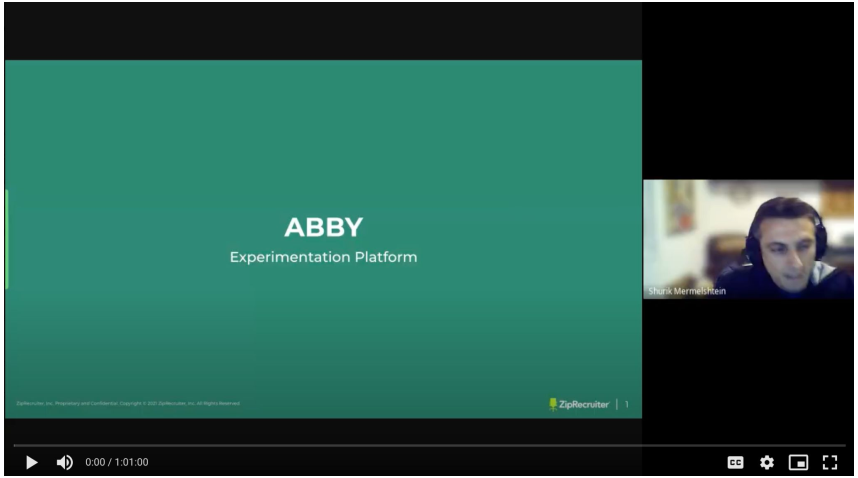

5. Deliver value early

Worst case: execute on long project to build a cool new classifier or centralized data store, then no one uses it.

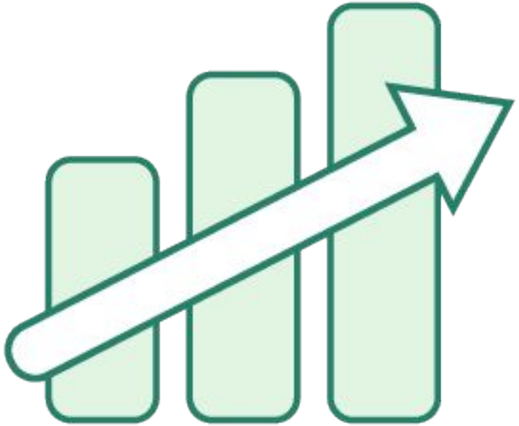

Validate need and impact **early and often** (in hacky ways if needed)

6. Commit to trust

The power of data is that it is **objective**.

Empower all functions to push back if they feel a decision is being made that will compromise trust in the product you're building.

- Encourage diligence in product quality
- Reward persistence in tracking down problems

Progress over time

2022

2021

2020

2019

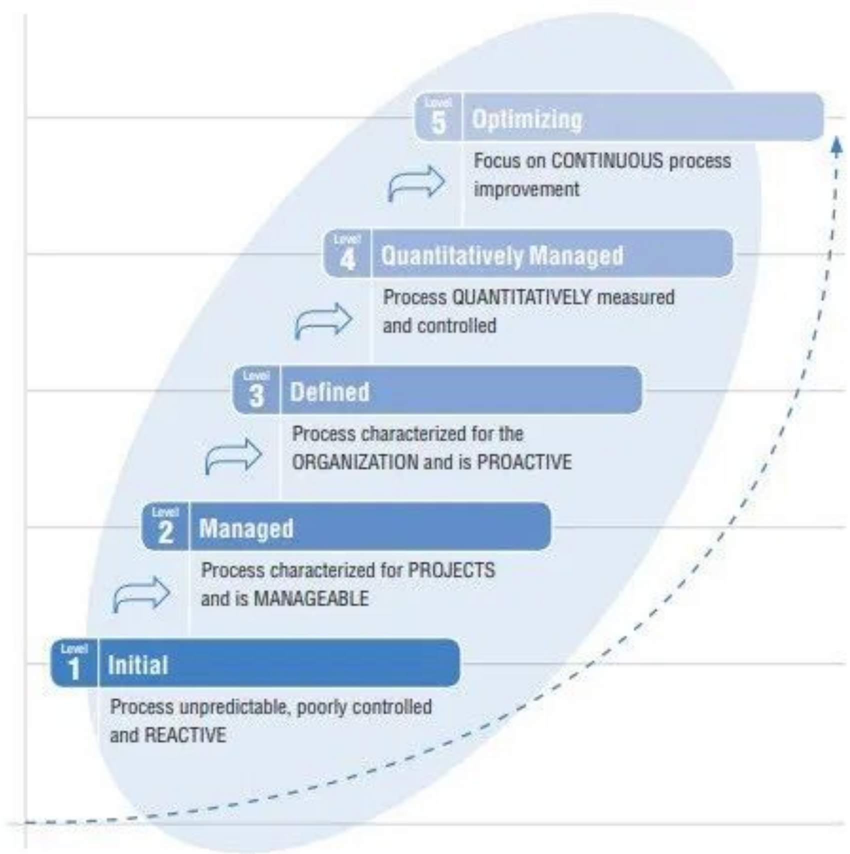

Thank You

Towards Human-AI Teaming:

Intelligence Ecosystems to Tackle High Stakes Use Cases

Clodéric Mars - VP of Engineering @ AI Redefined

AIs actions don’t align with humans intents because they are not aware of context

AIs don't learn from collaborating with humans

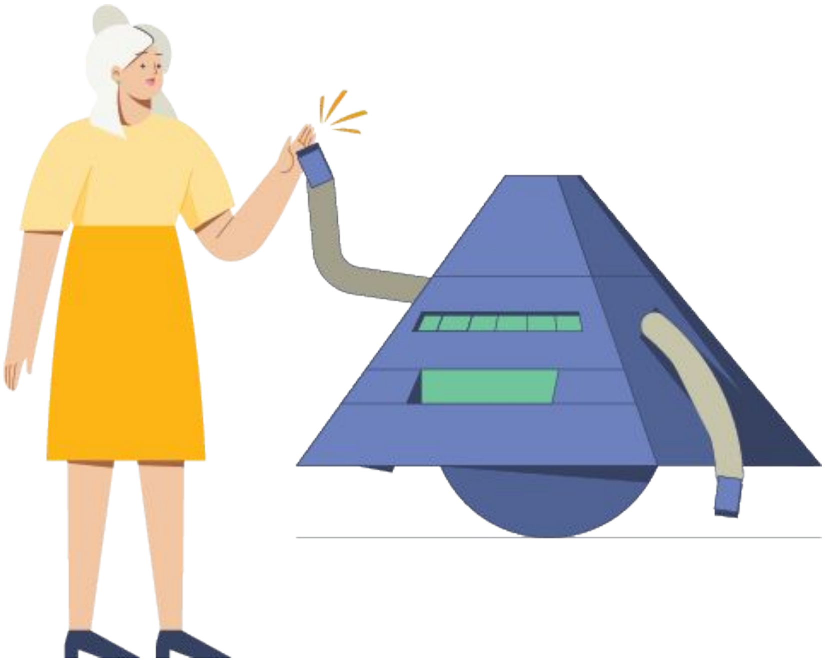

Achieve Human + AI Synergy
Intelligence ecosystems continuous learning through shared experiences

air

Behavior Cloning / Imitation Learning

Humans demonstrating how to achieve a task

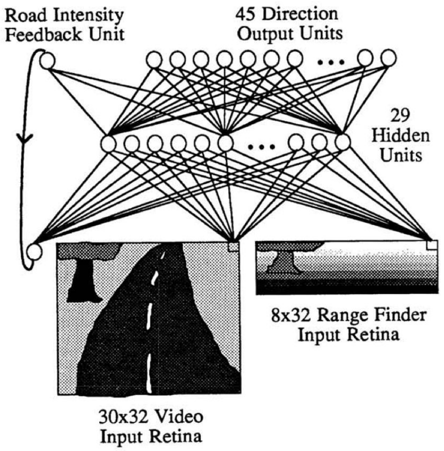

ALVINN: An Autonomous Land Vehicle In A Neural Network
1988 - Pomerlau
https://proceedings.neurips.cc/paper/1988/file/812b4ba287f5ee0bc9d43bbf5bbe87fb-Paper.pdf

- Fast initial training bootstrapping & continuous interactive refinements
- Good alignment
- No additional human skills required
- Bounded performances

Reinforcement Learning with Human Feedback

Humans evaluating how a task is completed

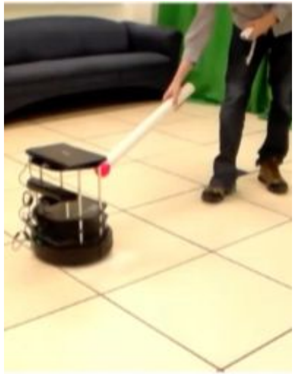

Interactive Learning from Policy-Dependent Human Feedback (COACH)
2017 - MacGlashan et al.
https://arxiv.org/abs/1701.06049

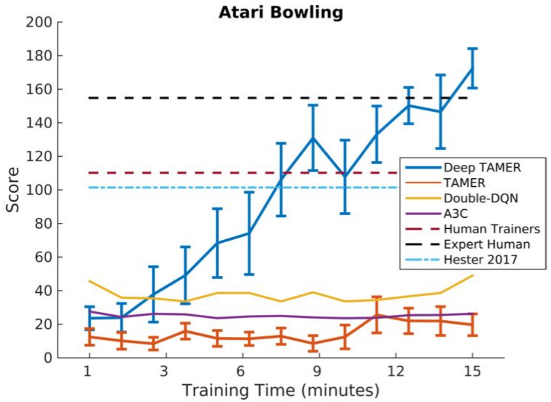

Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces
2018 - Warnell et al.
https://arxiv.org/abs/1709.10163

air

Reinforcement Learning, a very short aside

Discovering instead of reproducing

supervised learning

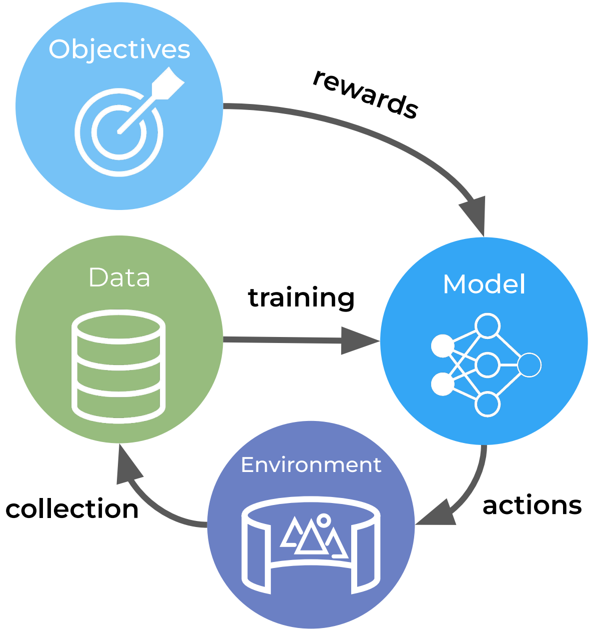

reinforcement learning

air

Reinforcement Learning with Human Feedback

Humans evaluating how a task is completed

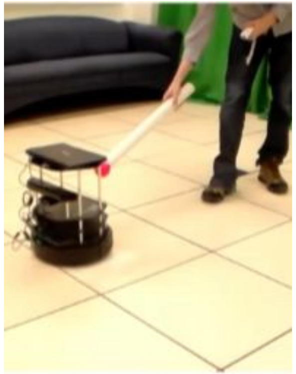

Interactive Learning from Policy-Dependent Human Feedback (COACH)
2017 - MacGlashan et al.
https://arxiv.org/abs/1701.06049

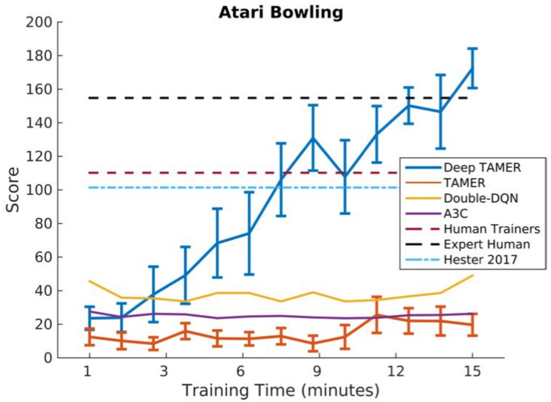

Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces
2018 - Warnell et al.
https://arxiv.org/abs/1709.10163

- Superhuman (even optimal) performances
- Indirect alignment
- Some specialized skills required & labor intensive
- Requires the ability to deal with the lag between an action and its evaluation
- Requires safe environment

AI apprentice: dual control

Interactive human demonstrations to accelerate exploration

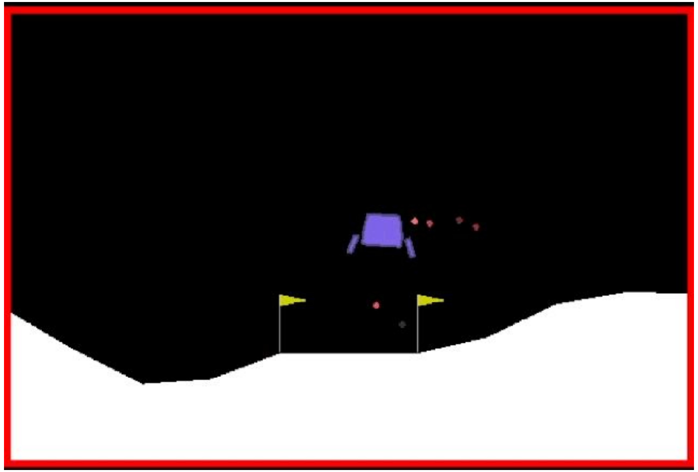

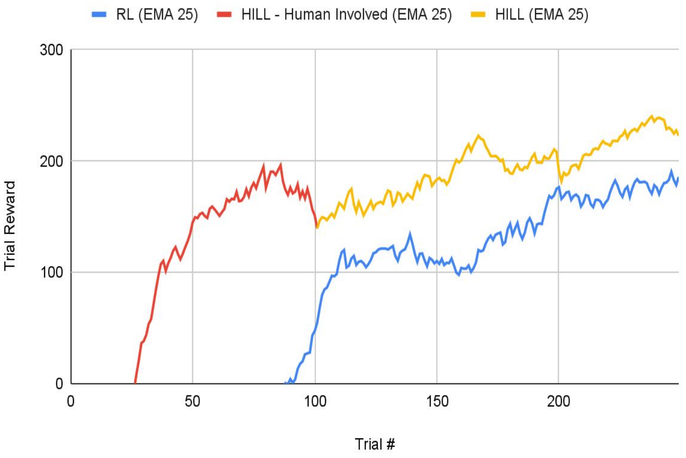

Superhuman (even optimal) performances

Indirect alignment

No additional human skills required

Requires collaborative UX during training and operation

Powered by cogment

air

Human-guided summarization

Combining Human-in-the-Loop Learning with language models to improve Human/AI alignment

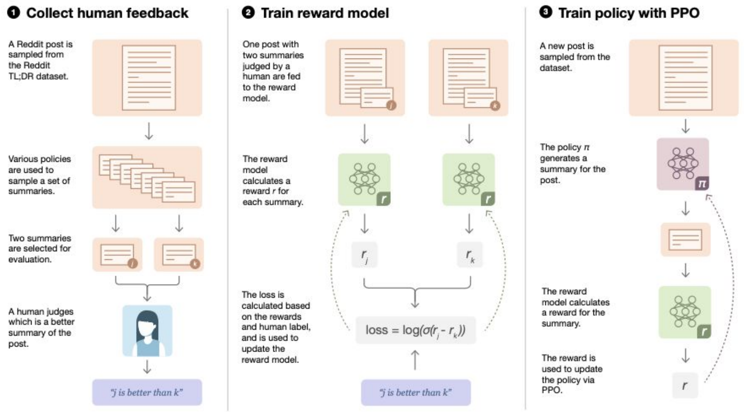

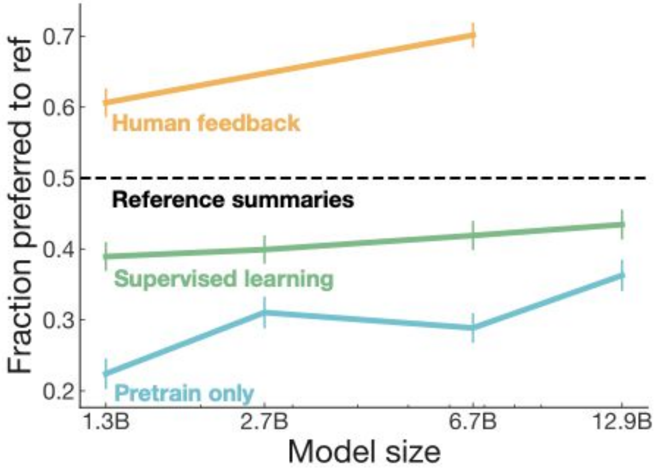

Learning to summarize from human feedback
2020 - Stiennon et al.
https://arxiv.org/abs/2009.01325

- Fine tune offline trained model
- Hybrid technique: demonstration + feedback
- Leverage human expertise on language task
- Improved alignment

Man-Machine Teaming: Multiple AIs and Humans collaborating to complete a task

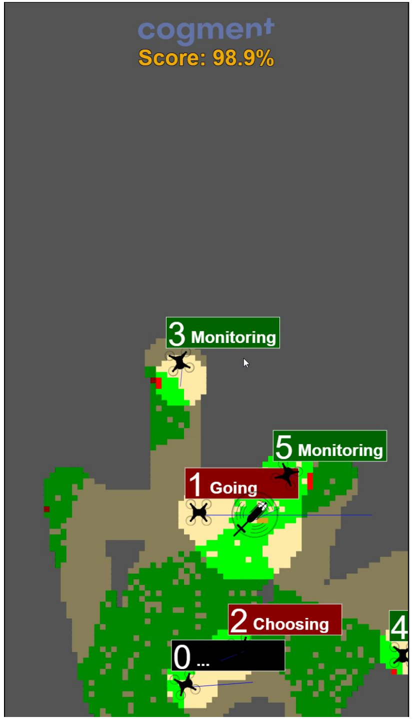

Training AI agents to coordinate together and with humans and vice versa

Powered by

cogment

- Implicitly train AIs from their interaction with humans
- Train Humans alongside AIs
- Requires bootstrapped AIs
- Enables human supervision

air

Intelligence Ecosystem: Heterogeneous actors collaborating

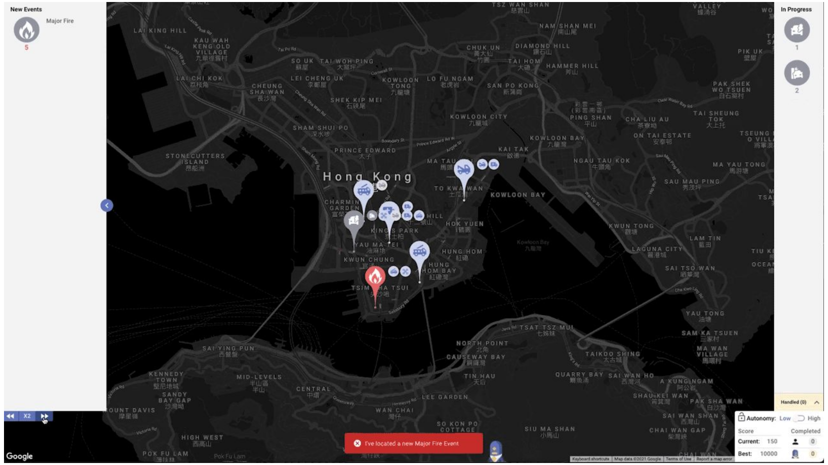

Training and operating complex topologies of roles and tasks within a common environment

- Faster deployment and iterations
- Enable supervision & learning by keeping humans in the loop
- Compliance & accountability where it matters

Applications to other verticals

Utilities

Education

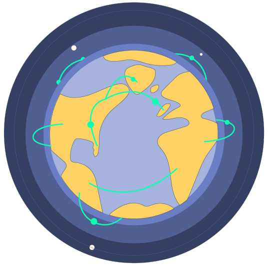

Logistics
Manufacturing

Transportation

Health

Intelligence Ecosystem: Benefits

Leverage human expertise: training, supervision, collaboration

Deploy faster & continuously get better

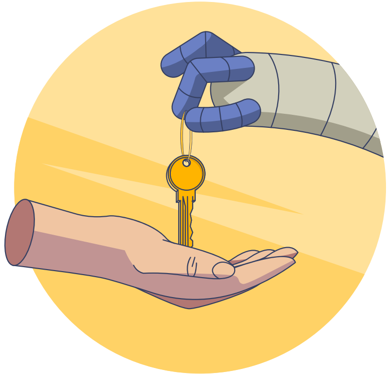

Trust, compliance & accountability

air

Intelligence Ecosystem: Requirements

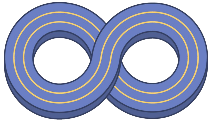

Continuous learning from building to operation

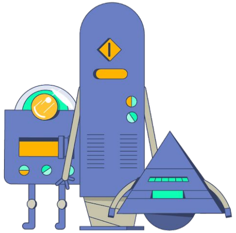

Multiple AI techniques

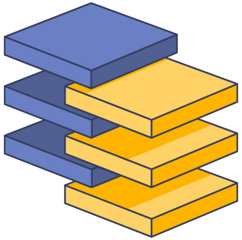

Tech agnostic

air

Cogment: Build, train, and operate AI agents in simulated or real environments shared with humans

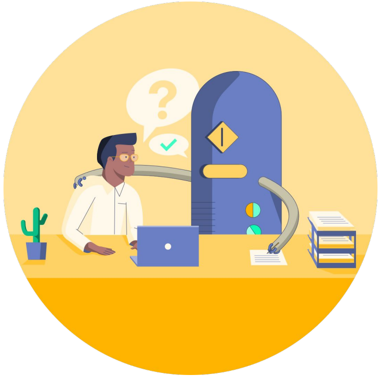

Continuously train AIs & Humans together

Less data required
Real time adaptation
Faster training
Fostering trust

Operate intelligence ecosystems

Best of human & AI capabilities
Human supervision when it matters
Hybrid AI: compliance and high performance
Modular approach: reduce compute usage & facilitate validation

Iterate smoothly from sim to real

Safe and simple design and training in simulation
Progressive deployment to real environment
Real environments, digital twins, numerical simulations, etc.

Available open-source & with further information at https://cogment.ai

air

Lessons learned designing AI-enabled products

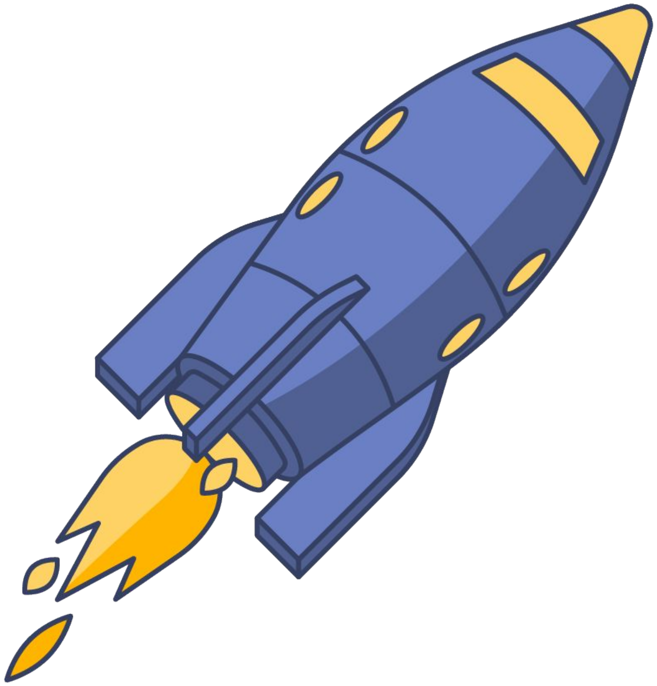

- Take into account the Human/AI relationship
- Consider AI apprenticeship approach
- Think in terms of intelligence ecosystem

AI Redefined: Humans and AI elevating each other

P.S. We are hiring!

Clodéric Mars

cloderic@ai-r.com

@cloderic

ai-r.com

cogment.ai

github.com/cogment

Exploring Data through Natural Language Conversations

Anand Ranganathan
Co-Founder & Chief AI Officer
anand@unscrambl.com

UNSCRAMBL

How do people consume data & analytics today?

– through charts & dashboards

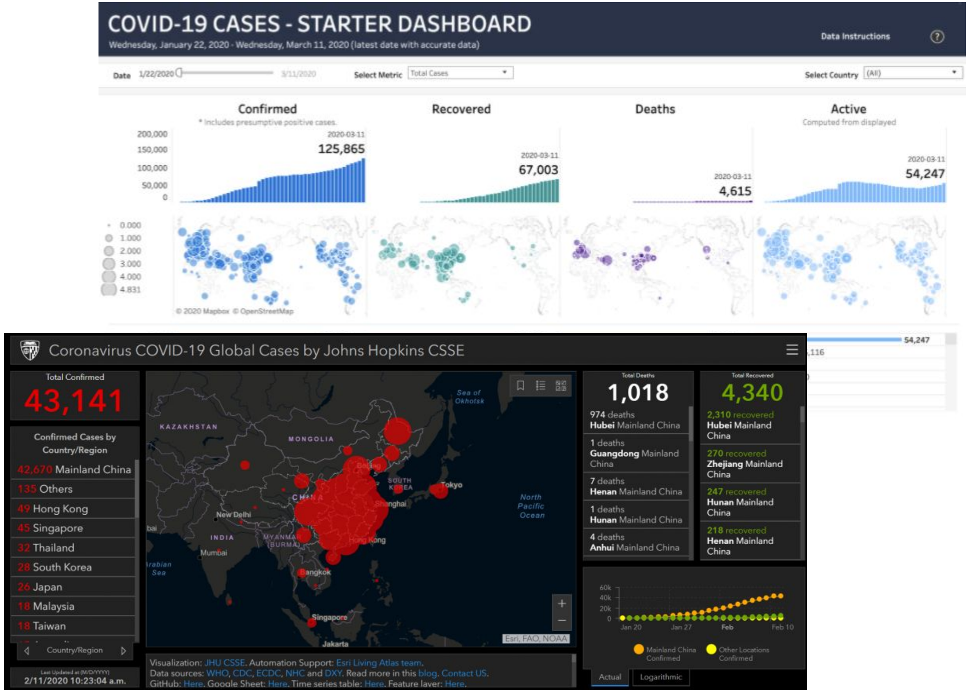

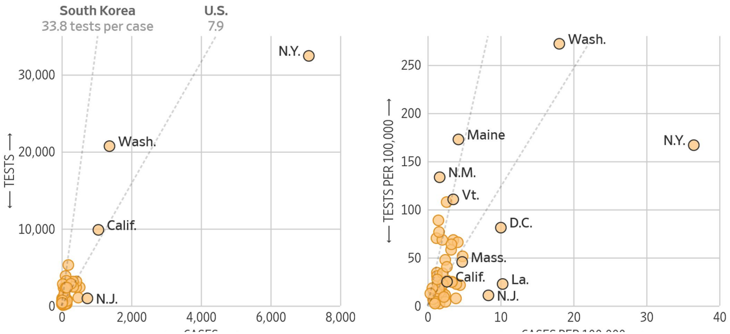

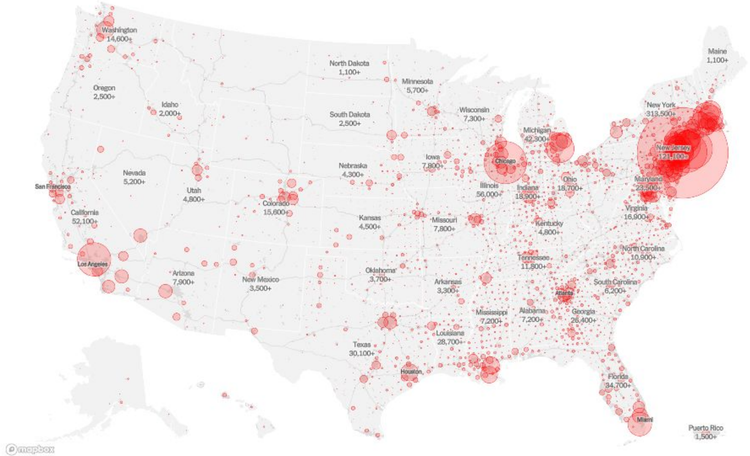

What's wrong with dashboards

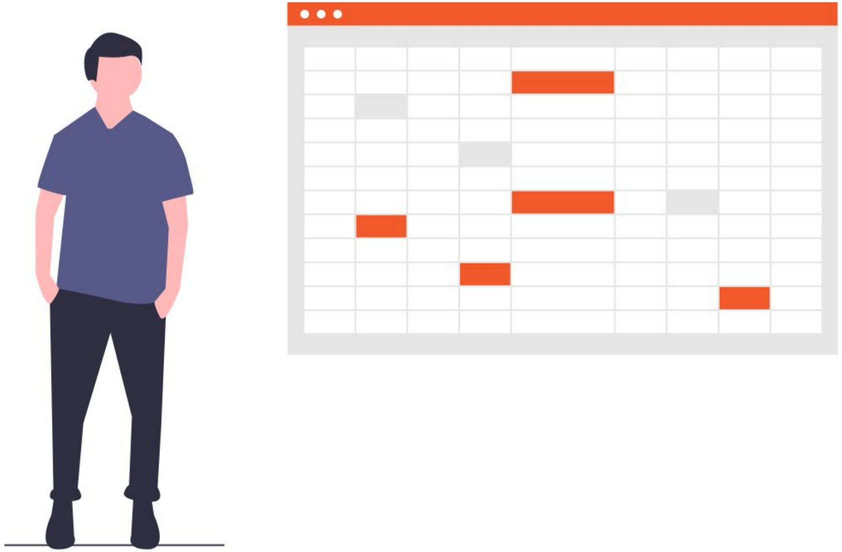

- Limited or no drill-downs
- Don't know how the data for the dashboards is produced
- Can't ask a slightly different question
- Representative of an opinion; easy to cherry-pick stats
- Can be misleading

UNSCRAMBL

UNSCRAMBL

What is Conversational Analytics?

Allow any user to ask text or voice questions of their data

and

receive back a natural language + visual analysis
of statistically relevant and actionable insights for that user.

**Note that in this talk, we focus on structured data stored in relational format (e.g. SQL databases, Excel sheets, etc)*

UNSCRAMBL

Qbo: Natural language conversations with data within collaboration platforms

qbo

AI-Powered,
Data Analyst

Connect
Converse
Collaborate

Why have sales dropped in the North region this month?

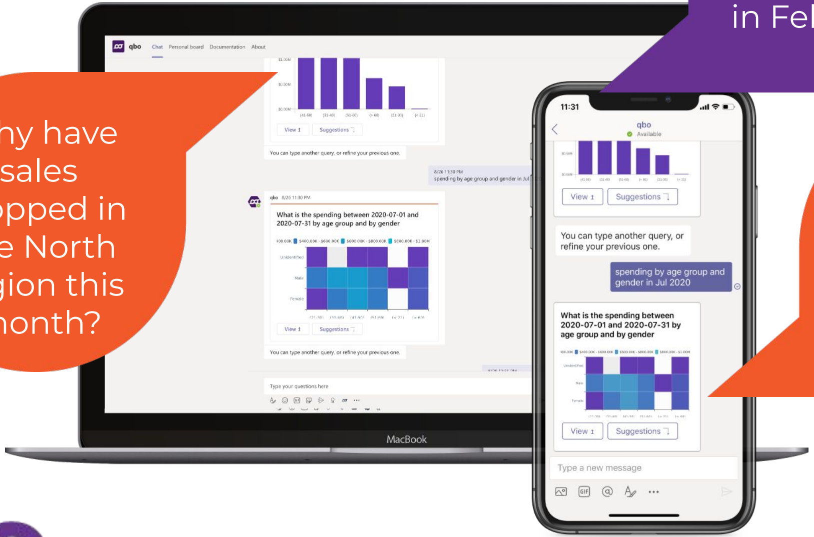

Hey **QBO**, why were new acquisitions in Feb lower?

Hey **QBO**, how many policies are expiring in Oct 2022?

UNSCRAMBL

Qbo sits between users and disparate, siloed datasets

Support 20+ data connectors, and access via a web interface or Microsoft Teams

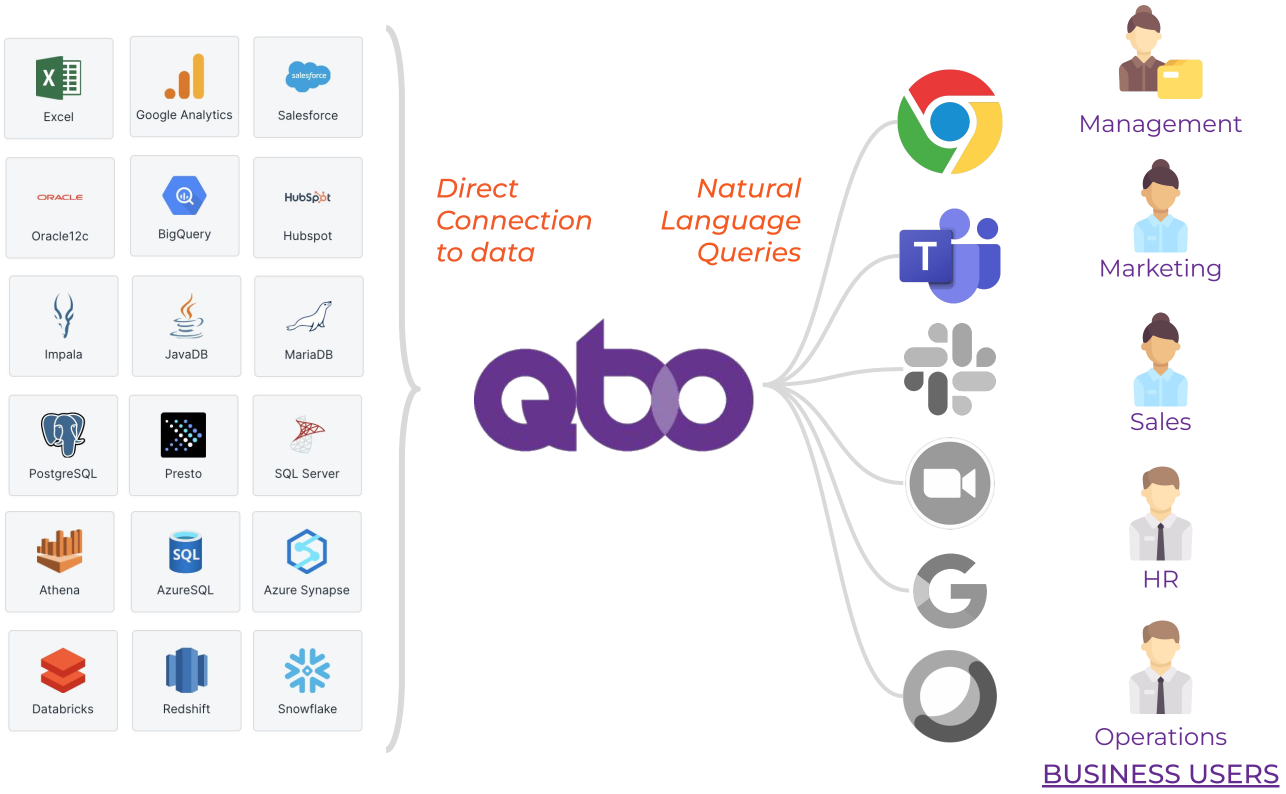

UNSCRAMBL

Unscrambl Qbo Demo : on Austin 311 Data

Anand Ranganathan
Co-Founder & Chief AI Officer
anand@unscrambl.com

UNSCRAMBL

A (very simplified) overview of NLU pipeline

Anand 7:29 PM

number of trips in winter 2017 by age and gender

Entity Recognition & Construction

number of trips in winter 2017 by age and gender

Identification of Query type and mapping to known concepts in DB

Type: Aggregation Query on Trips table with a group-by and a filter; **age ->** derived from birth year attribute; **gender ->** gender attribute; **in winter 2017 ->** 2017-12-23 and 2018-03-19 (filter)

Generate DB-specific SQL query

```
SELECT anon_1."age group", anon_1.gender, count(*) AS "Count"
FROM (SELECT "TRIP_ANALYSIS".end_station_id AS "end station id", "TRIP_ANALYSIS".program_id AS "program id", "TRIP_ANALYSIS".start_station_id AS "start station id",
"TRIP_ANALYSIS".bikeid AS bikeid, CASE WHEN (:birth_year_1 - "TRIP_ANALYSIS".birth_year < :param_1) THEN :param_2 ELSE CASE WHEN (:birth_year_2 -
"TRIP_ANALYSIS".birth_year < :param_3) THEN :param_4 ELSE …, "TRIP_ANALYSIS".gender AS gender
FROM "TRIP_ANALYSIS"
WHERE ….
```

Get results, decide on visualization and narratives, and present back to user

qbo 7:29 PM

What is the total number of the trips between 2017-12-23 and 2018-03-19 by age group and by gender

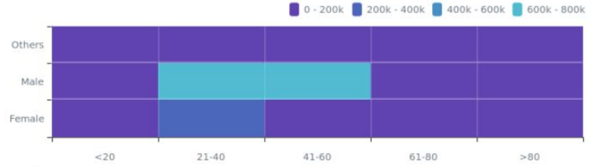

View

UNSCRAMBL

Key Challenge : Bridging the gap between users and data

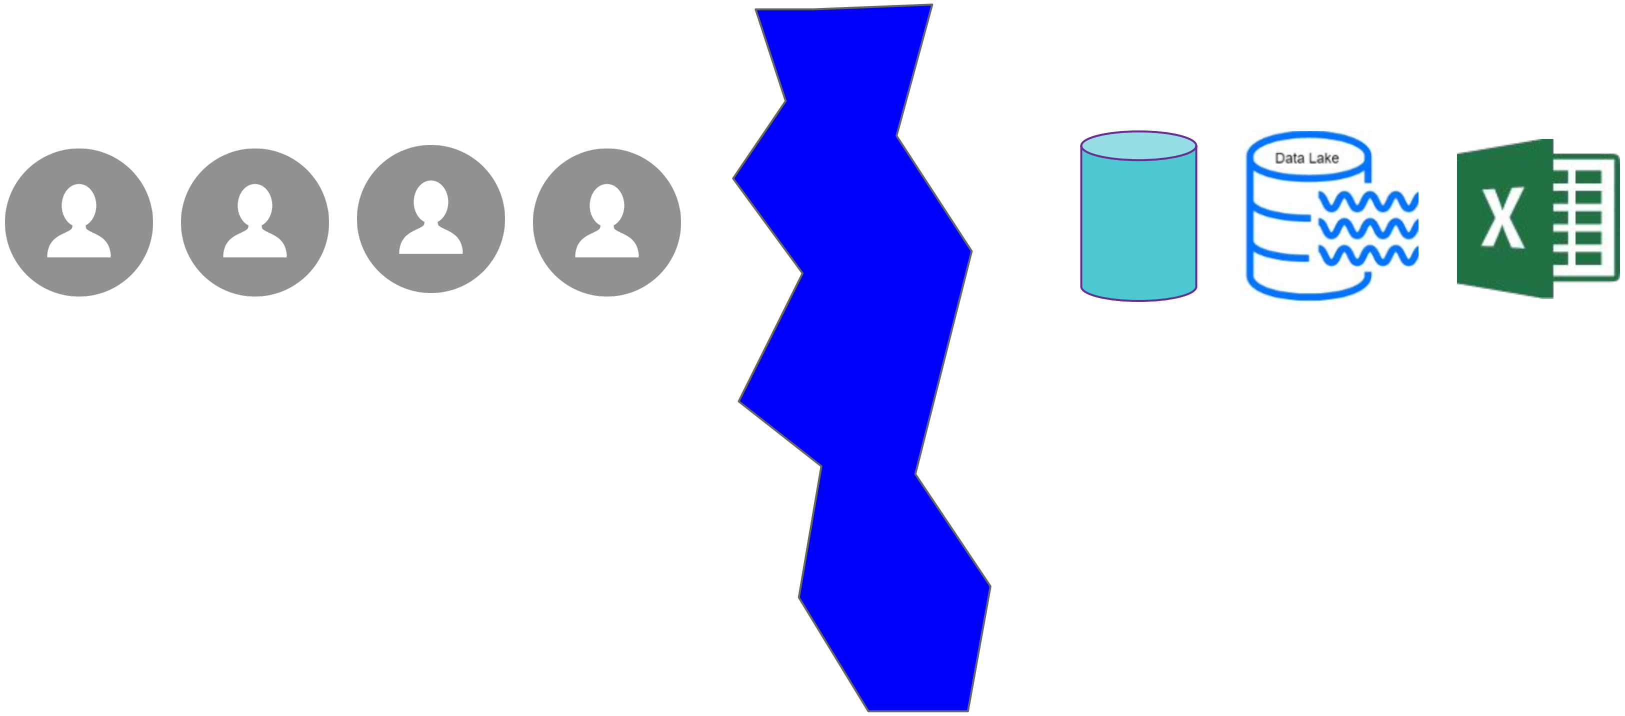

UNSCRAMBL

Key Challenge : Bridging the gap between users and data

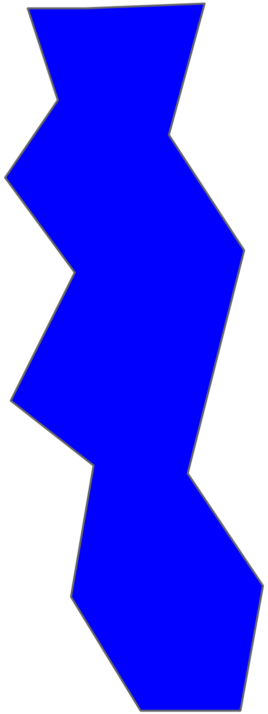

- Users don't know what to ask
- Users don't know how to ask
- Users may pose questions in an ambiguous manner
- Users may use terms not in the dataset

Key Challenge : Bridging the gap between users and data

- Users don’t know what to ask
- Users don’t know how to ask
- Users may pose questions in an ambiguous manner
- Users may use terms not in the dataset

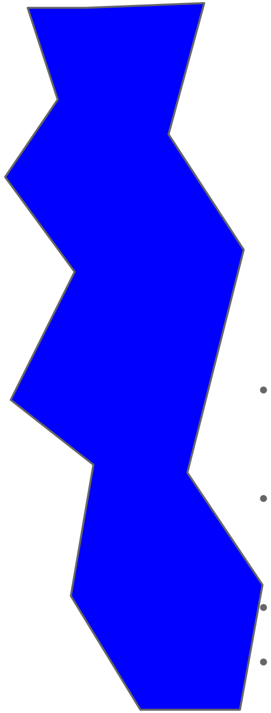

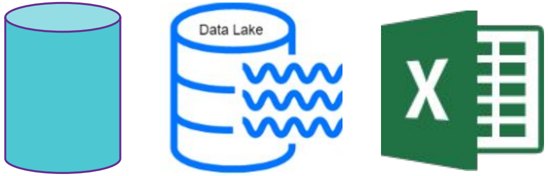

- Data may be modeled in a variety of ways
- Hidden semantics and assumptions behind different tables and columns
- Data may be incomplete, unclean
- Data may be spread across silos

UNSCRAMBL

Users ask questions the way they want -- & not be constrained by the physical data model

What is the total number of rides by age group

What is the ADPU in the last 6 months?

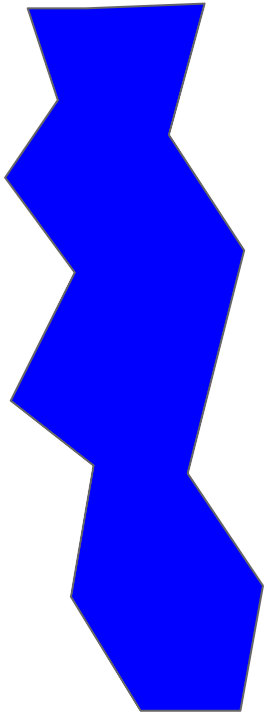

SEMANTIC GAP --

MUST BE BRIDGED BY THE PERSONAL DATA ANALYST

PHYSICAL DATA MODEL

Trip
starttime
stoptime
tripduration
birth_year
bikeid
usertype

UNSCRAMBL

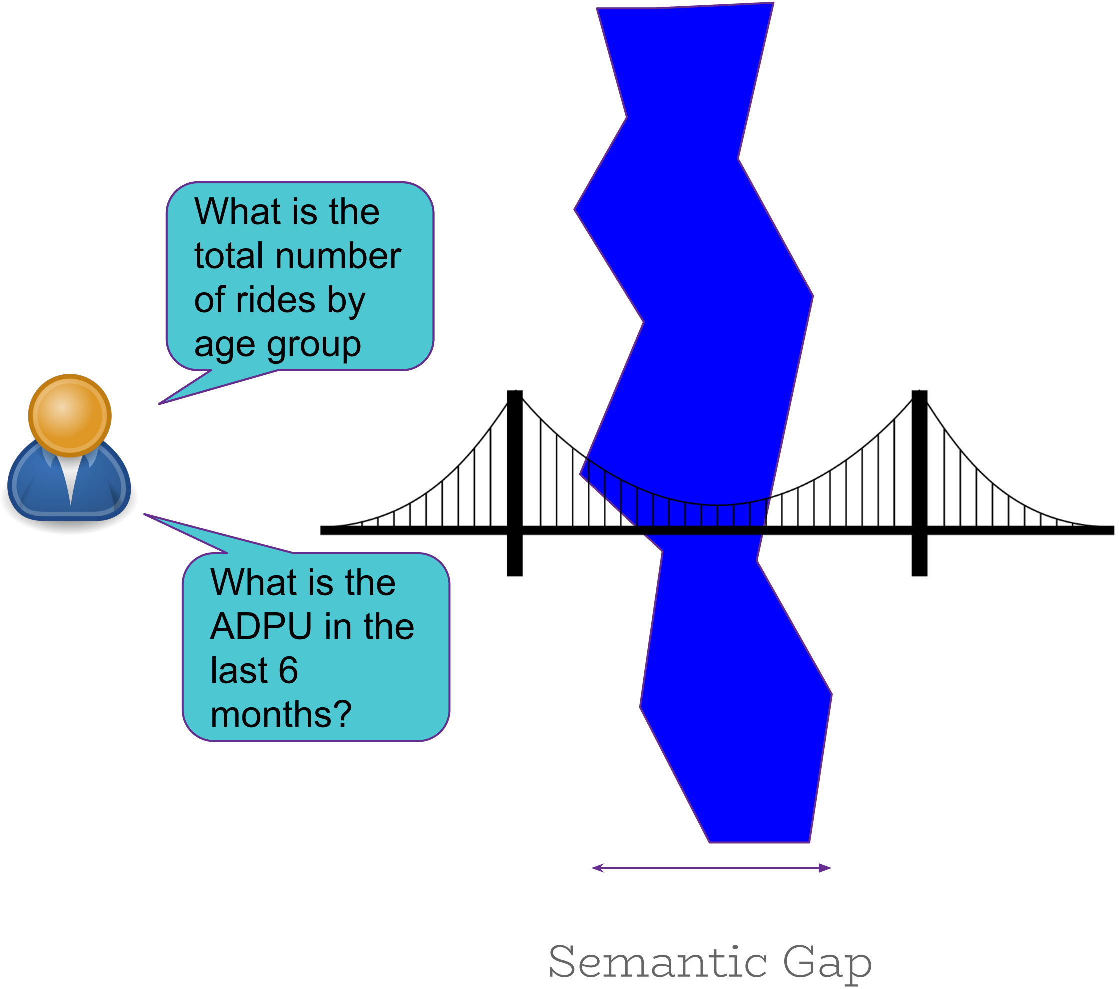

Physical Data Model

Trip
starttime
stoptime
tripduration
birth_year
bikeid
usertype

UNSCRAMBL

Logical Data Model

(the bridge between the user and the physical data model)

Trip (= Ride)

Attributes :

- bike id
- start time (default date attribute)
- stop time
- day of week (synonym: weekday)
- season
- age (unit : year)
- age range
- user type
- trip duration (unit : second, default sort attribute)

Named Aggregates (or Metrics) :

- ADPU (synonym: Average duration per user)

Named Filters (or Segments) :

- very old customers

Physical Data Model

Trip
starttime
stoptime
tripduration
birth_year
bikeid
usertype

UNSCRAMBL

Our approach to solving some of these challenges

UNSCRAMBL

Step 1. Initial Data Discovery & Configuration

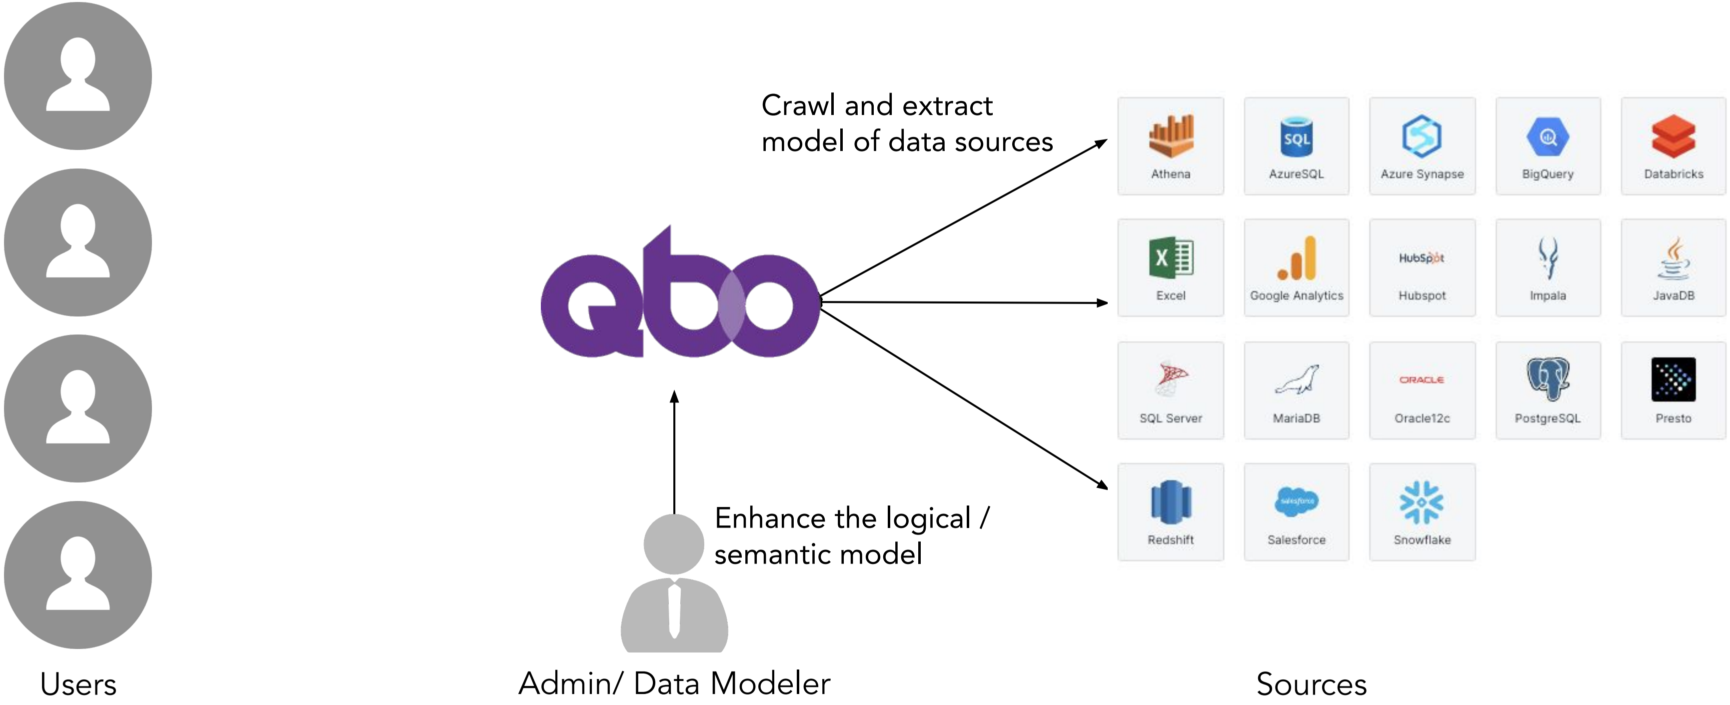

UNSCRAMBL

Step 2. Users can converse with qbo about their data

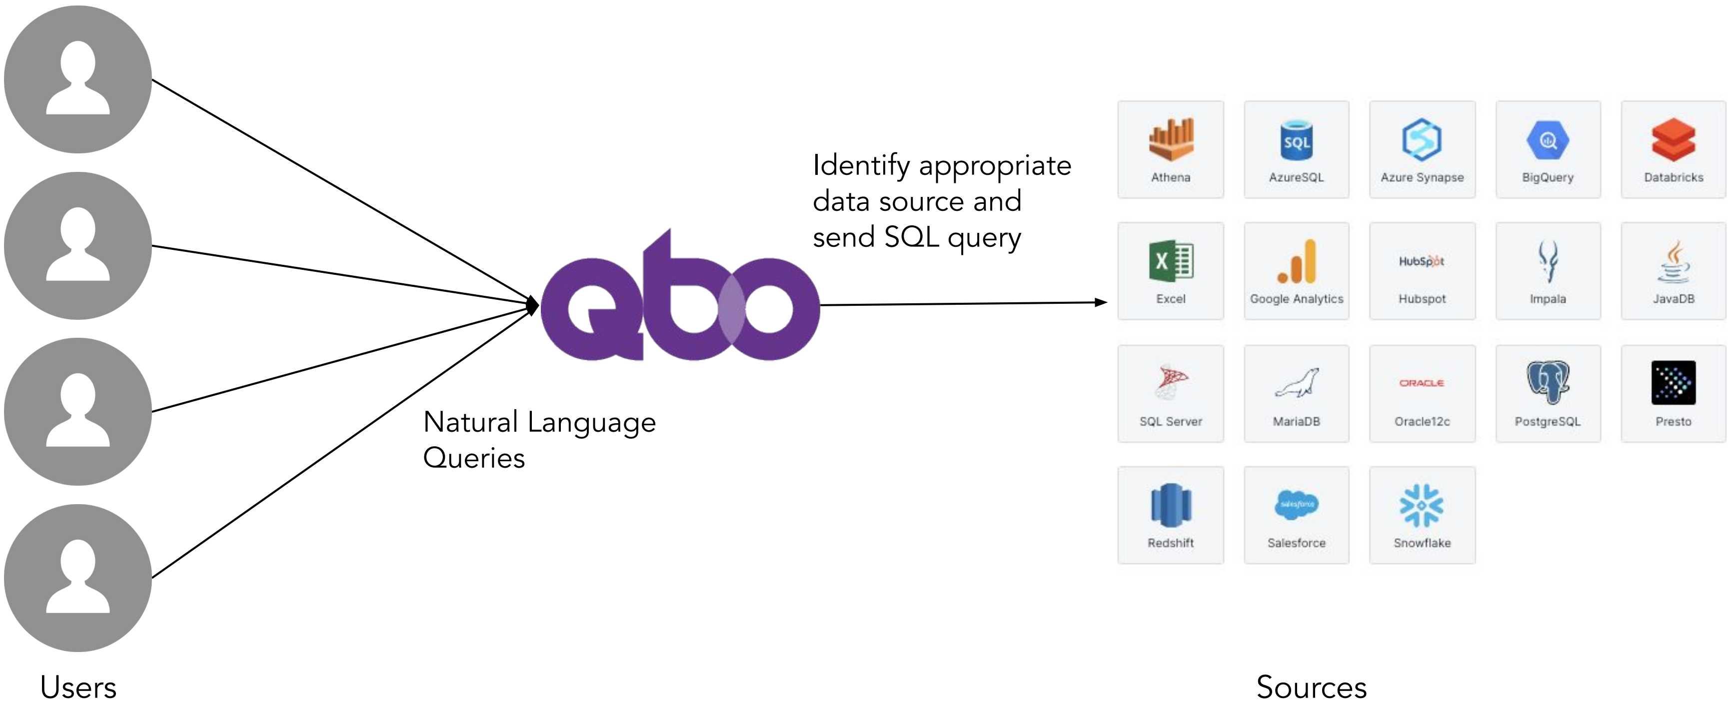

UNSCRAMBL

Step 2. Users can converse with qbo about their data

Imagine...

#futureofwork

#futureofdata

Try for free at qbo.ai

Get it from Microsoft AppSource

UNSCRAMBL

Building “Responsible AI”:

Best Practices Across the Product Development Lifecycle

Susannah Shattuck

Head of Product, Credo AI

credo | ai

Who Am I?

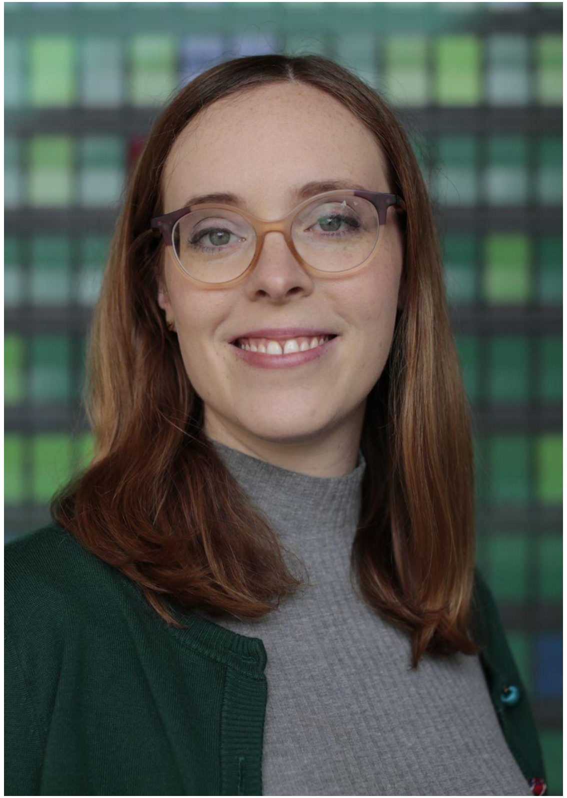

Nice to meet you, I'm Susannah.

👩‍💻 Head of Product, Credo AI

🏢 Formerly IBM, Google X, Arthur AI

🧠 Algorithmic bias detection and mitigation, setting up governance structures within AI development organizations, AI regulation & public policy

😍 Speculative design, science fiction, cooking

@shshattuck


B I N G O
12 27 44 57 69
5 29 38 47 72
11 30 FREE 4702 SPACE 53 67
3 26 32 59 73
2 24 42 56 63

The six key tenets of Responsible AI.

FAIRNESS	TRANSPARENCY	SAFETY & SECURITY
PRIVACY	SOCIAL & ENVIRONMENTAL WELL-BEING	ACCOUNTABILITY

Responsible AI considerations need to be integrated into the ML development lifecycle.

How does Responsible AI assessment fit into the ML development lifecycle?

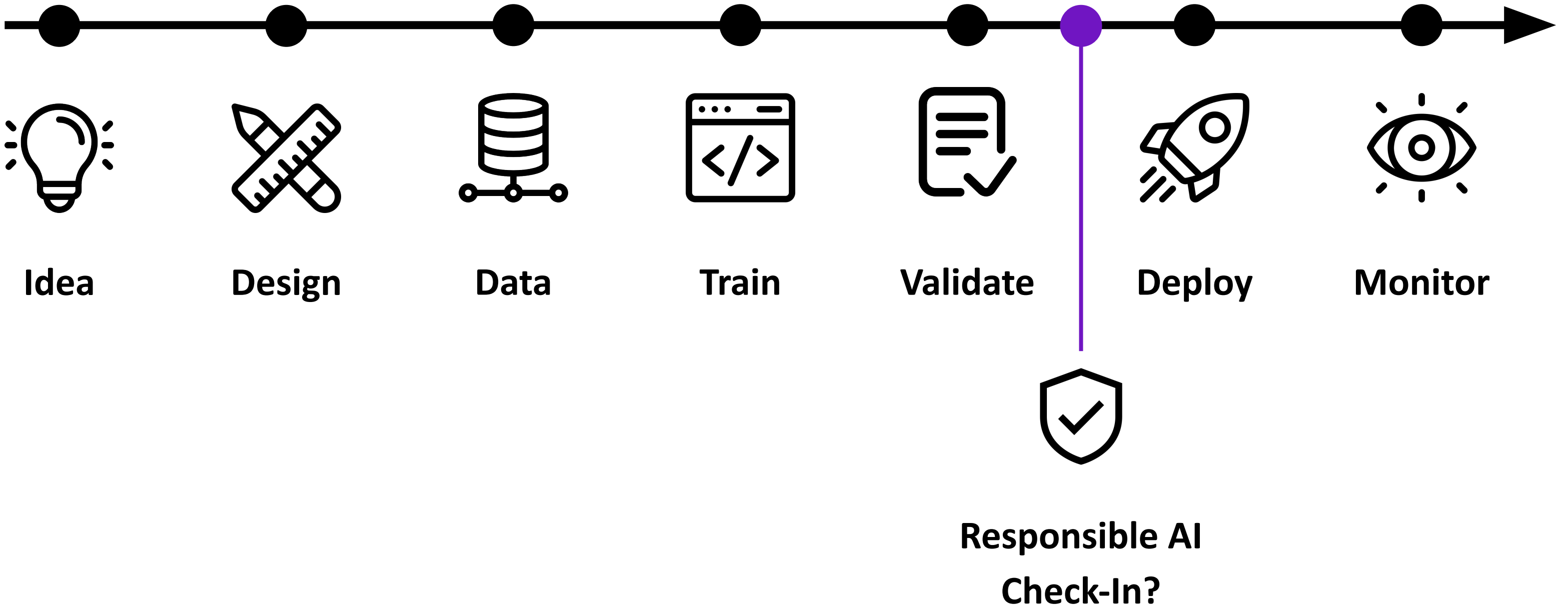

How does Responsible AI assessment fit into the ML development lifecycle?

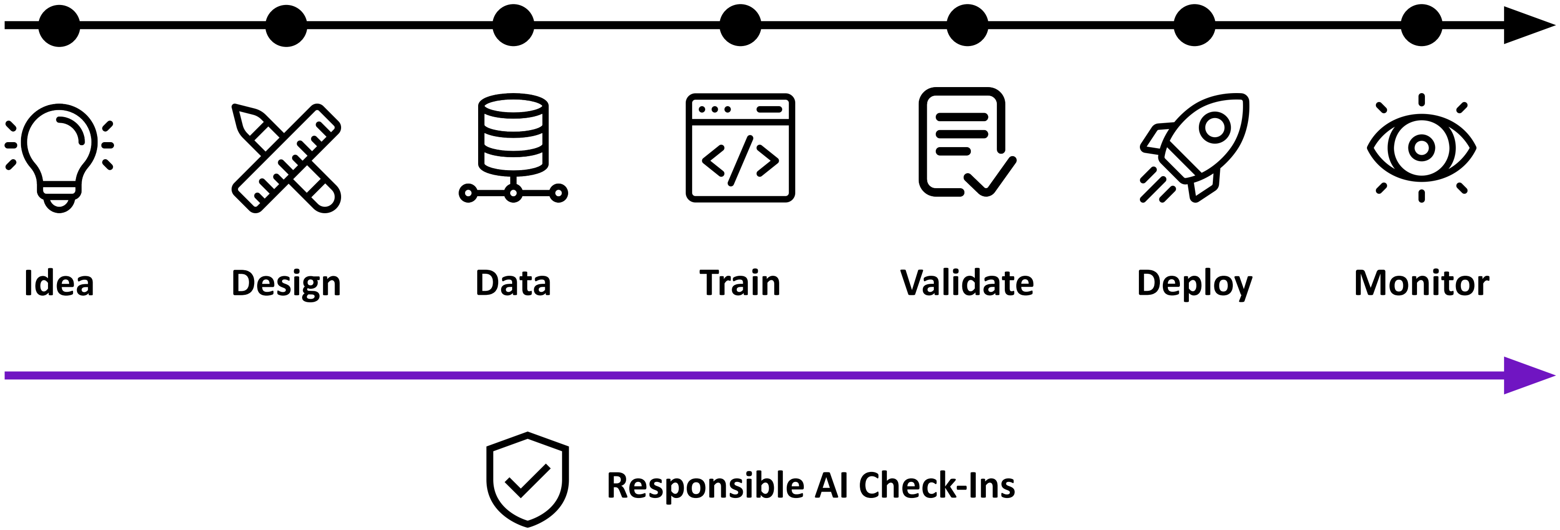

How does Responsible AI assessment fit into the ML development lifecycle?

Tl;DR—you need to evaluate the “responsibility” of your AI system at every step of the development lifecycle.

🎨 **During Design:** identify potential risks of the use case and how to measure them

🛠️ **During Development:** prioritize Responsible AI metrics during training and testing

🚀 **During Deployment:** monitor Responsible AI metrics, conduct regular audits

🙅‍♀️ **This is not something that you can or should do alone!** Getting input from different perspectives is key—Responsible AI is a multi-disciplinary problem.

DESIGNING RESPONSIBLE AI SYSTEMS

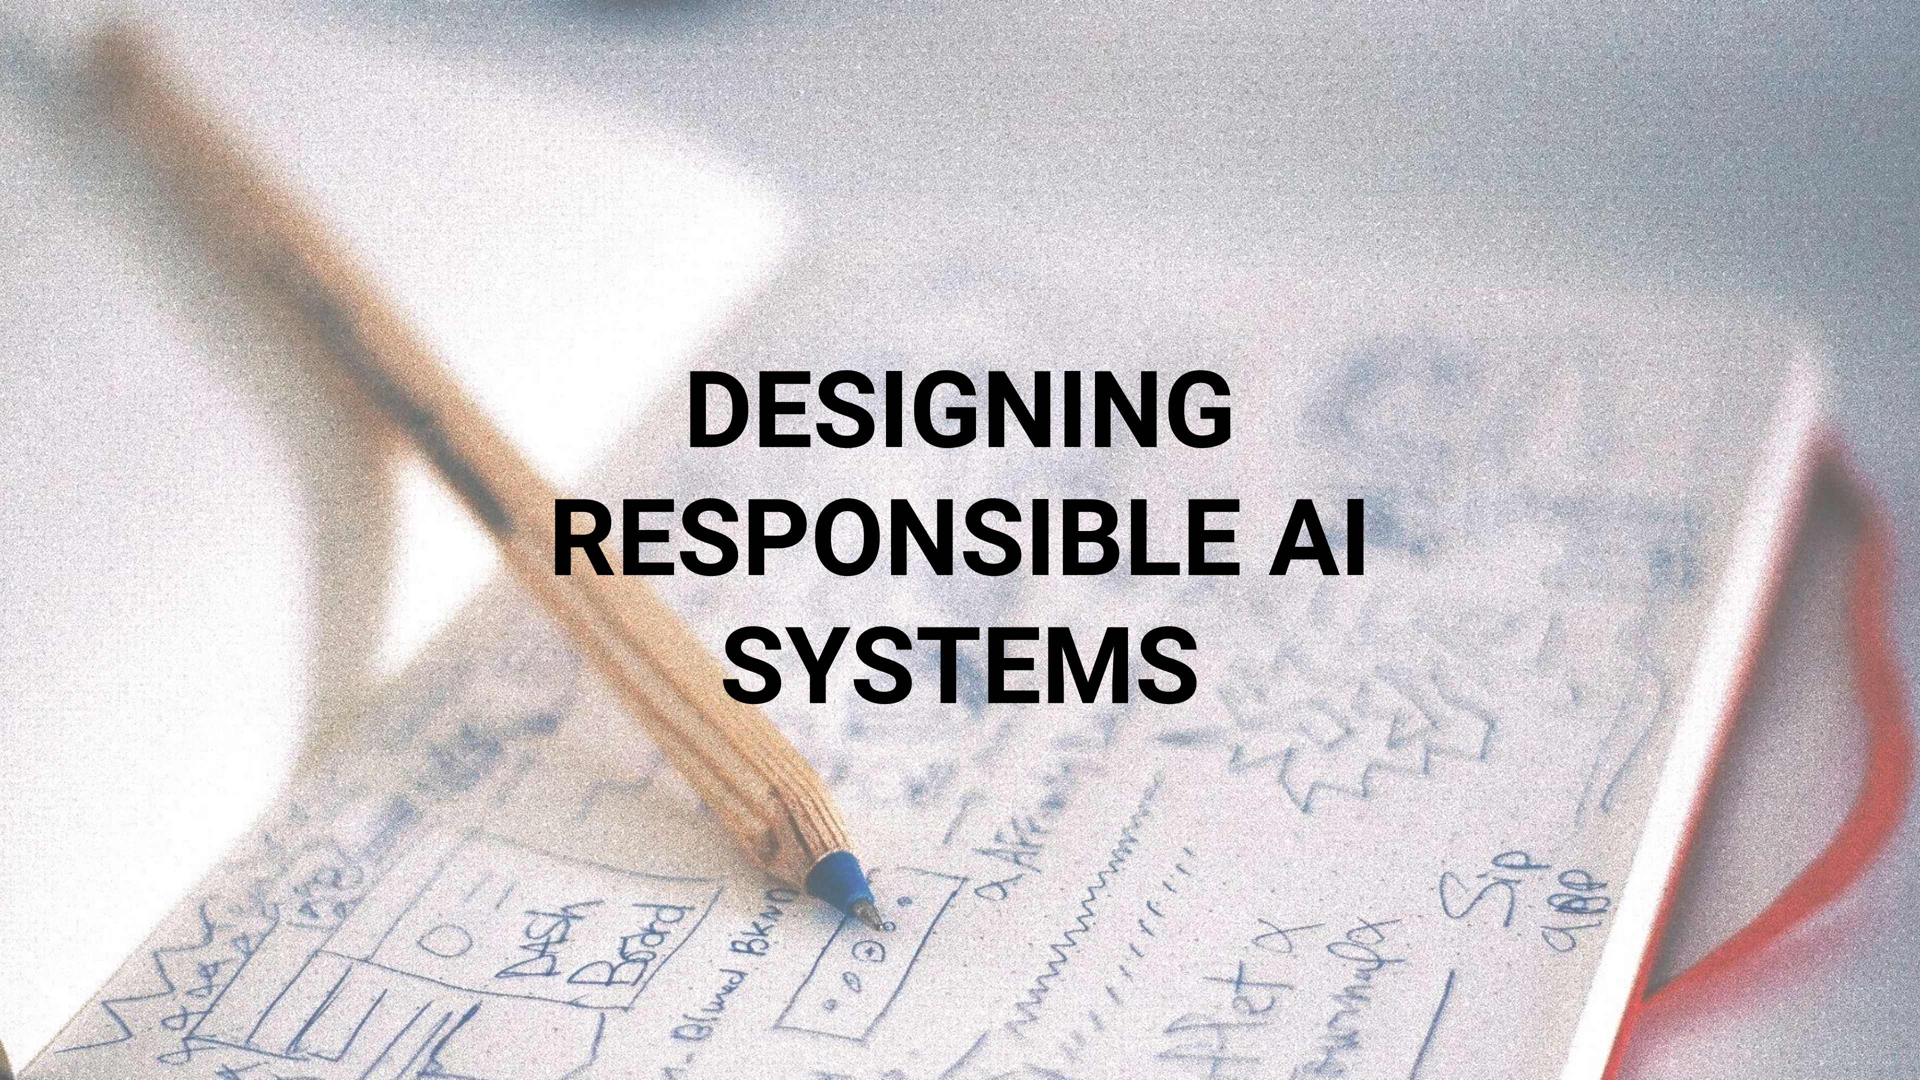

Evaluating your use case: a multi-stakeholder project.

👤 **Who is going to be impacted?** Think about both direct and indirect users; identify all of the groups that will be affected by use of your AI system.

🙁 **What are the potential negative impacts on these people/groups?** Talk to people. Do real user research. Invite impacted groups to participate in the design process.

⚖️ **What is the regulatory context?** Are there any rules, regulations, or standards that need to be followed based on your use case?

📏 **How might we measure and mitigate negative impacts?** Develop a Responsible AI Assessment Plan that will address negative impacts and regulatory requirements.

Tools that help with Responsible AI Alignment:

- AEQUITAS Framework
- Industry standards and benchmarks (NIST, IEEE, etc.)
- Credo AI

DEVELOPING RESPONSIBLE AI SYSTEMS

Measuring Responsible AI during, not after, development.

🎯 **Include Responsible AI metrics in your objective function.** Optimize for the most performant model that meets your RAI requirements.

📊 **Don’t just evaluate your models; evaluate your data.** Fairness and privacy assessments should happen at the dataset level *before* the model level.

🛑 **Rule out model methodologies that don’t meet requirements from the start.** Is explainability a regulatory requirement? Don’t waste time building a black box model.

📄 **Document your development decisions.** Transparency and accountability are made possible by good documentation; create consistent artifacts during development.

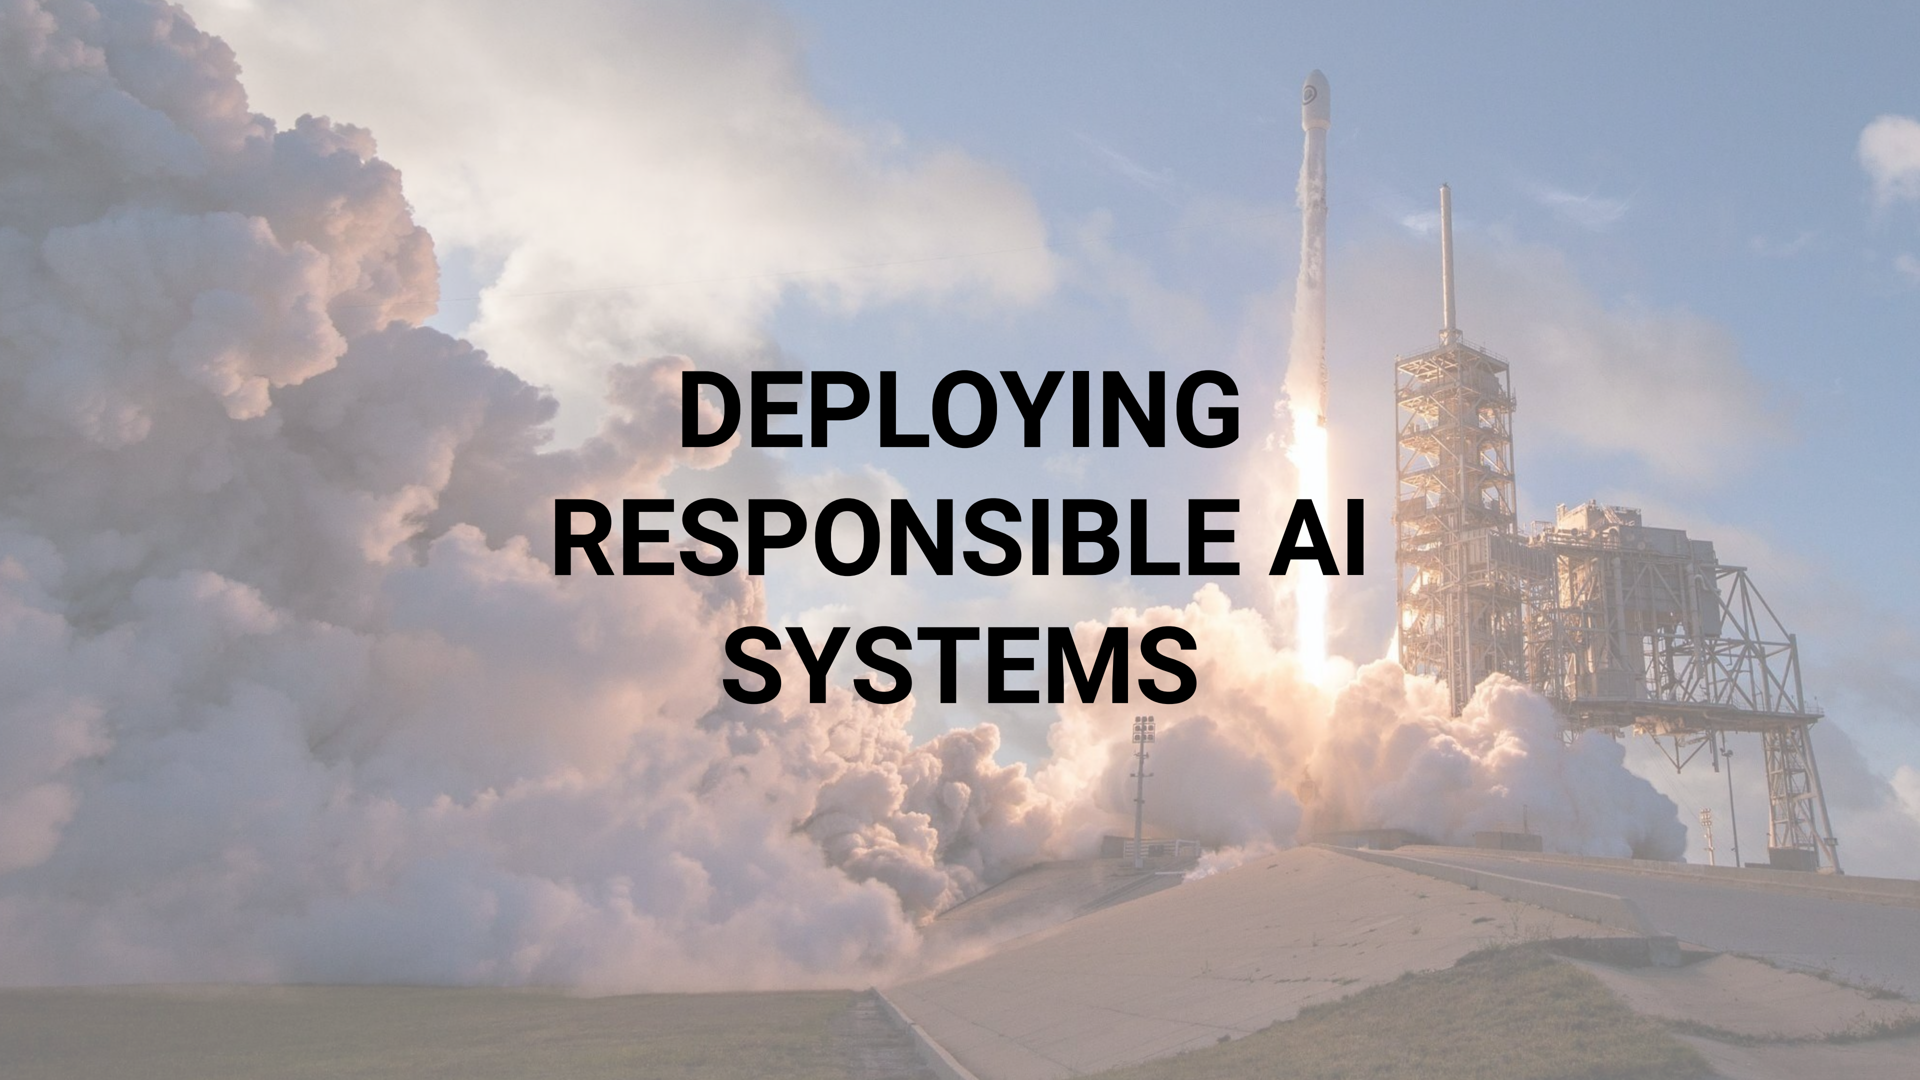
DEPLOYING
RESPONSIBLE AI
SYSTEMS

Continue monitoring and managing Responsible AI in production.

🔎 **Include Responsible AI metrics in your monitoring plan.** Don’t just monitor performance and drift; make sure you’re tracking fairness metrics, too.

📅 **Conduct regular stress tests and audits.** Regulations are increasingly requiring regular audits or reports on AI systems’ behavior over time.

📢 **Build Responsible AI feedback mechanisms.** Get feedback from your users and the communities impacted by your AI system—and act on that feedback regularly.

🕵️ **Have a plan in place if something goes wrong.** Who is responsible for fixing a problem, when it arises? What is your mitigation plan for Responsible AI issues?

CREDO AI LENS

What you can’t observe, you can’t control.

Bringing Responsible AI Assessment tools together.

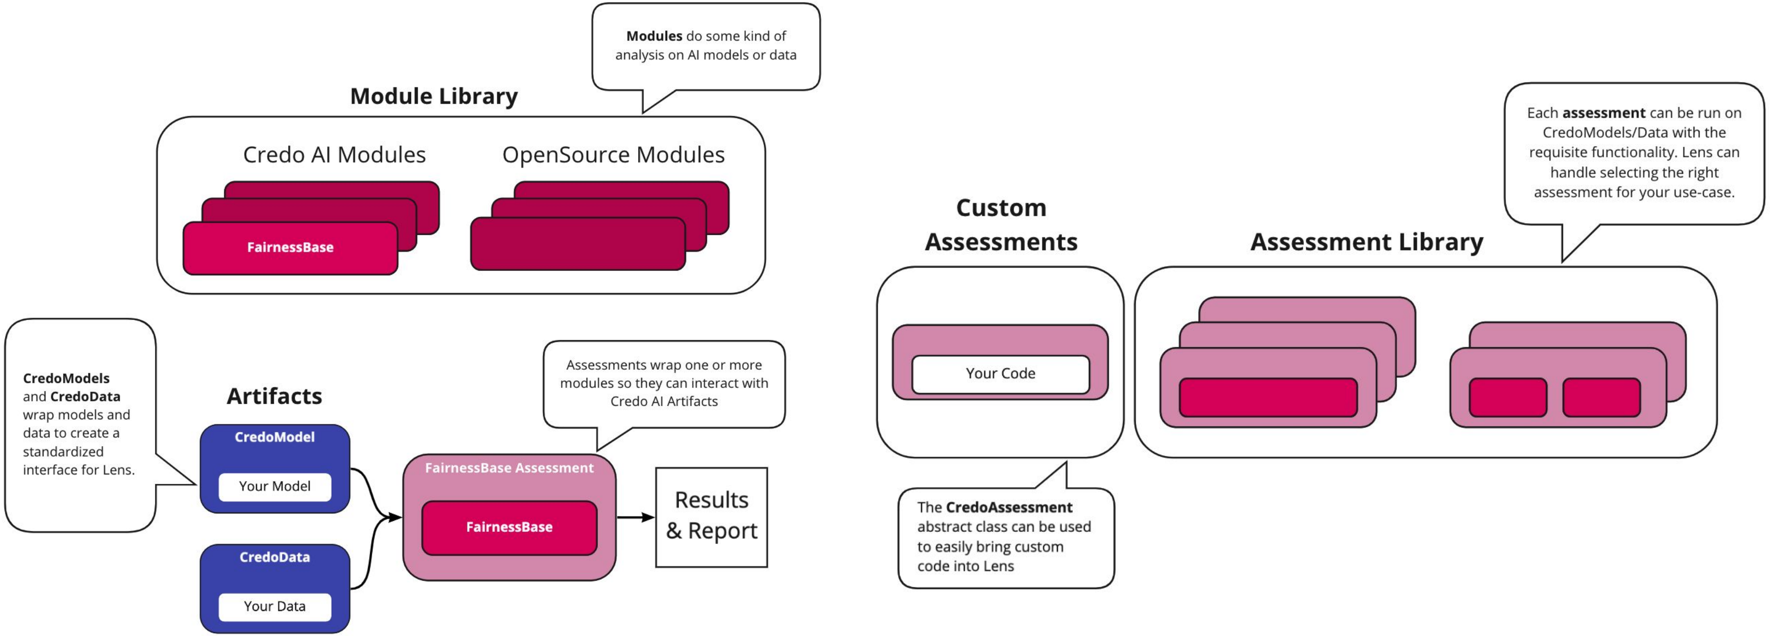

Current Credo AI Lens Assessment Capabilities:

⚖️ **Fairness assessments.** Easily assess parity metrics like disparate impact, equal opportunity difference, etc. for binary classification models.

📈 **Dataset assessments.** Detect proxy variables for protected attributes and get demographic parity analysis of your datasets.

📝 **Custom NLP assessments: toxicity, profanity, verbosity.** For large language models, run a variety of NLP-specific assessments to identify negative model behavior.

👨‍👩‍👧 **Disaggregated performance assessment.** Easily compare disaggregated performance of your model across groups of interest.

https://credoai-lens.readthedocs.io

Thank you!

@shshattuck

Watchful

The **Machine Teaching** Platform for the Enterprise

Contact John
COO & Co-Founder

Problems with **Hand-Labeling**, and the Efficacy of **Automation Techniques**

Major Issues with Hand-Labeling

Introduction of Bias

- Hand-labels are not interpretable or reproducible, and are inherently bias-prone

Prohibitive Costs

- Models that require lots of data or subject matter interpretation of the data are frequently cost prohibitive to build

Ground Truth Is a Lie

- The real world is full of shades of gray
- Hand-labeled data often does not capture the nuance of inter-annotator disagreement

Algorithmic Bias

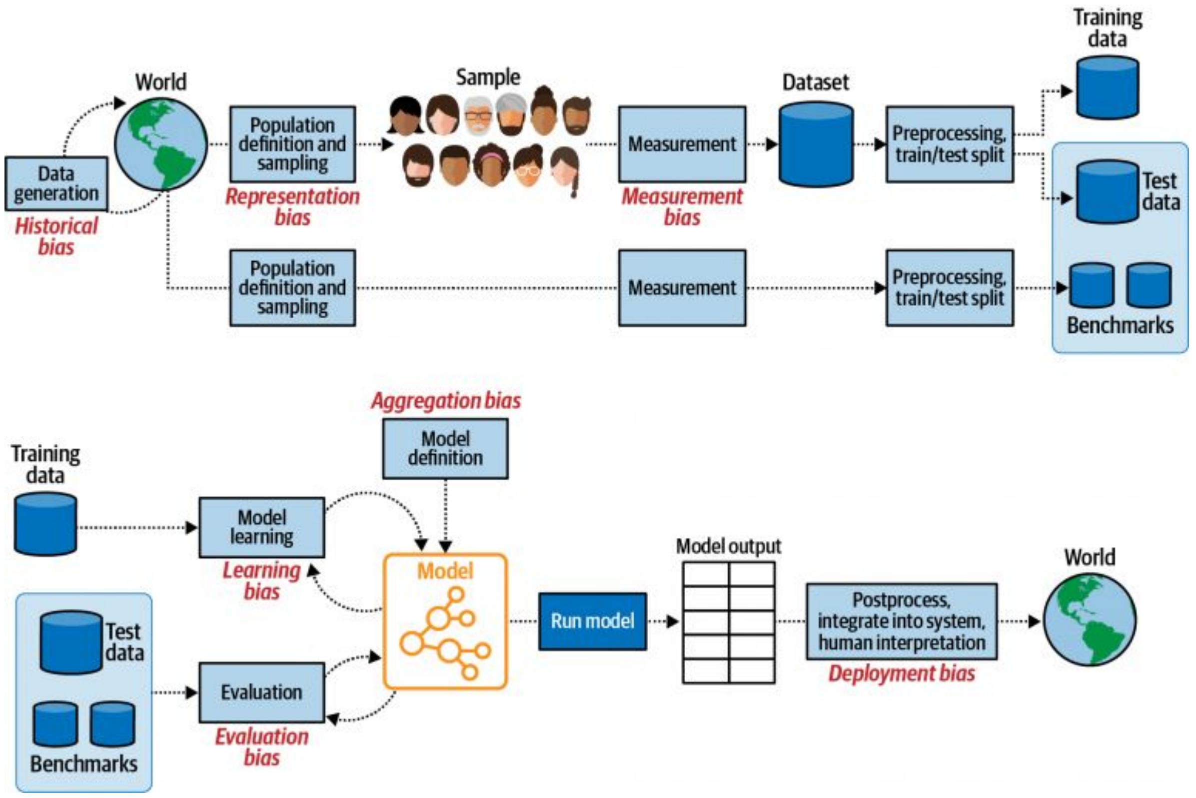

Source: [Hand Labeling Considered Harmful]

Algorithmic Bias

- Often creeps in through data
- Can’t explain hand-labels
- Can’t easily remedy bias in hand-labels

Hand-Labels: Lack of Interpretability & Explainability

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

What Do We Do About the Biases in AI?

by James Manyika, Jake Silberg, and Brittany Presten

October 25, 2019

The World of Alternatives

Low — **Amount Supervision Required** — High

Machine Teaching
Collection of techniques to extract knowledge from humans for model training

Weak Supervision
Noisy heuristics are used to weakly label large amounts of data for machine learning

Semi-Supervised Learning
Combine a small amount of labeled data with a large amount of unlabeled data

Active Learning
Algorithm queries users interactively to label specific segments of the data

Synthetic Data Generation
Building models to generate data points that have the same statistical validity as "real" data

Transfer Learning
Leveraging general pre-trained models to quickly bootstrap specific models

Machine Teaching: Introducing Explainability

- Techniques like weak supervision offer a framework for interpretability in labels
- Often must trade interpretability for quality
- Can combine approaches to achieve the right levels of interpretability, performance, and quality

The Algorithmic Auditing Trap

‘Bias audits’ for discriminatory tools are a promising idea, but current approaches leave much to be desired

Mona Sloane Mar 17 · 7 min read ★

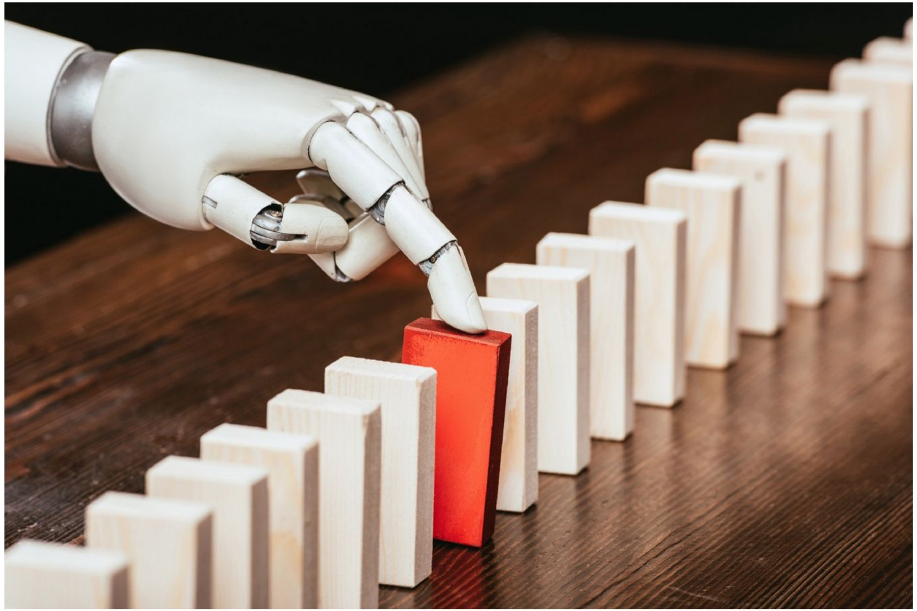

Image: LightFieldStudios/Getty Images

Societal, **Time**, and **Financial** Costs of **Hand-Labeling**

Societal Cost

MOTHERBOARD
TECH BY VICE

Underpaid Workers Are Being Forced to Train Biased AI on Mechanical Turk

Workers who label images on platforms like Mechanical Turk say they're being incentivized to fall in line with their responses—or risk losing work.

AN By Aliide Naylor

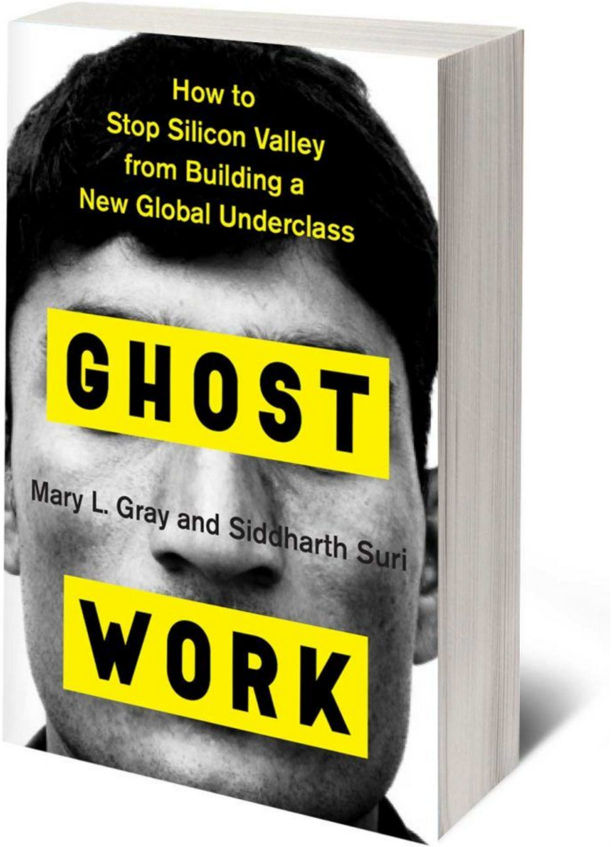

Time of SMEs

- The time of experts is the scarcest resource
- You’re never done labeling
- Time spent by experts must be measured over the lifetime of the model

Time of SMEs: Not Vertical Specific

Healthcare

- Clinical Trial Matching
- Clinical Decision Support

Finance

- Fraud detection
- Contract Intelligence

Insurance

- Risk Classification
- Claims Fraud Detection

Cost to Accuracy, Reality, and Ground Truth

Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks

Curtis G. Northcutt* ChipBrain, MIT
Anish Athalye MIT
Jonas Mueller Amazon

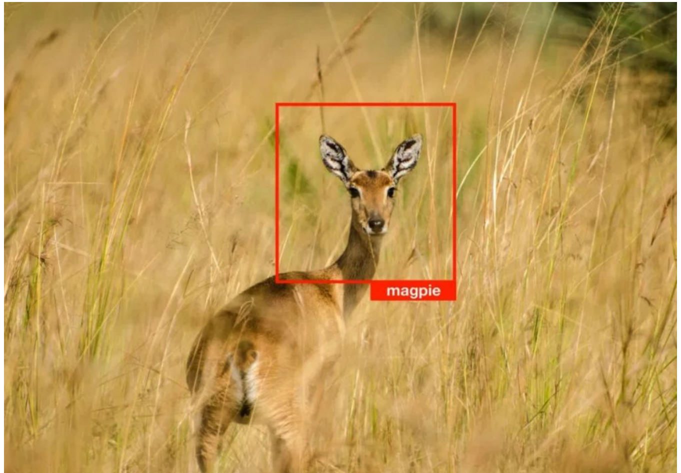

Measuring Model Biases in the Absence of Ground Truth

Osman Aka*
Google

Ken Burke*
Google

Alex Bäuerle†
Ulm University

Christina Greer
Google

Margaret Mitchell‡

Financial Costs of Hand Labeling

- You’re never done labeling
- Class definitions often change as labeling progresses
- Cost of SME time compounds cost of overall pipeline

The Efficacy of Automation Techniques

Scaling to Very Very Large Corpora for Natural Language Disambiguation

Michele Banko and Eric Brill
Microsoft Research
1 Microsoft Way
Redmond, WA 98052 USA
{mbanko,brill}@microsoft.com

Deep Learning Scaling is Predictable, Empirically

Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, Yanqi Zhou

{joel,sharan,ardalaninewsha,gregdiamos,junheewoo,hassankianinejad,
patwarymostofa,yangyang62,zhouyanqi}@baidu.com

Baidu Research

The Efficacy of Automation Techniques

*“We empirically validate that **DL model accuracy improves as a power-law as we grow training sets for state-of-the-art (SOTA) model architectures** in four machine learning domains: machine translation, language modeling, image processing, and speech recognition. These power-law learning curves exist across all tested domains, model architectures, optimizers, and loss functions.”*

-- Hestness et al. 2017

Driving ROI from Data Products

Data and AI in Construction

Alvaro Soto, Director of Product, Data at Procore

A large industry with data everywhere

Construction is a 10 trillion industry (13% of global GDP) that has just started to be digitized in the past 10 years. Procore's platform is at the forefront of this shift helping general contractors, owners, and subcontractors collaborate and manage their projects in real time. 110 TB of data run through the Procore project management and financials platform every month.

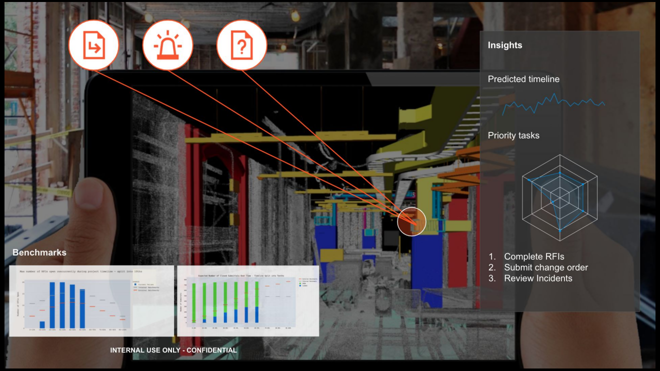

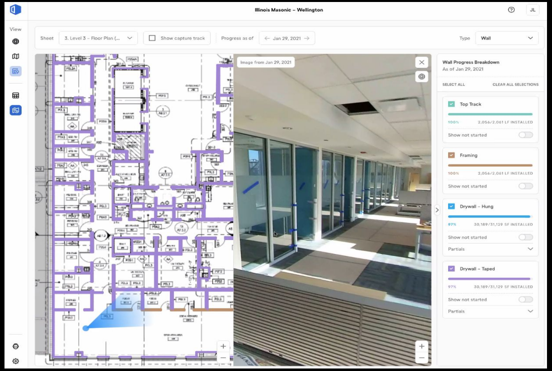

An evolution in data products

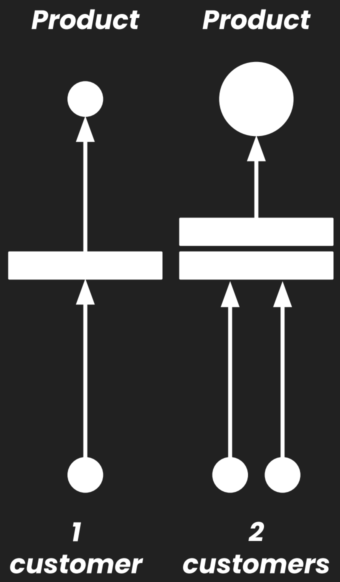

The addition of data (customer or ad-hoc) makes the product more useful

Examples: Dashboards, reports, benchmarks

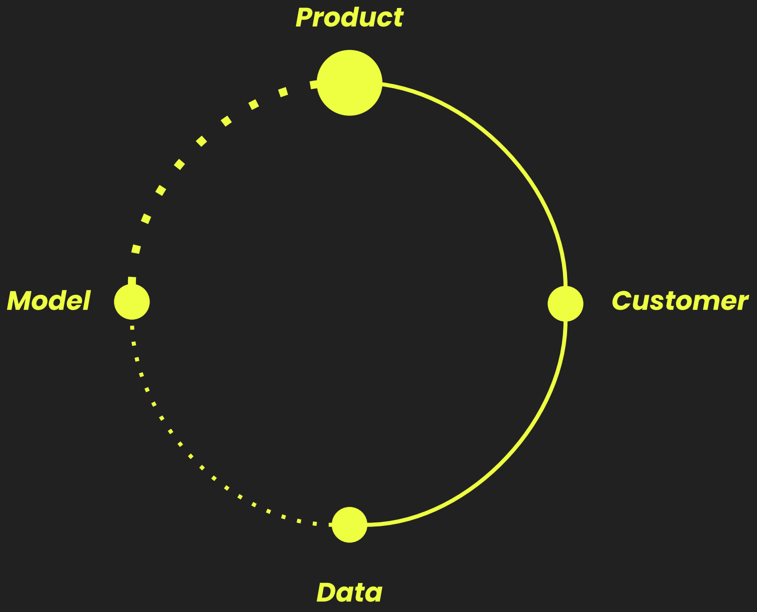

The addition of data (customer or ad-hoc) plus an AI model makes the product more useful.

Examples: recommendations, predictions, automation

Establishing data learning loops

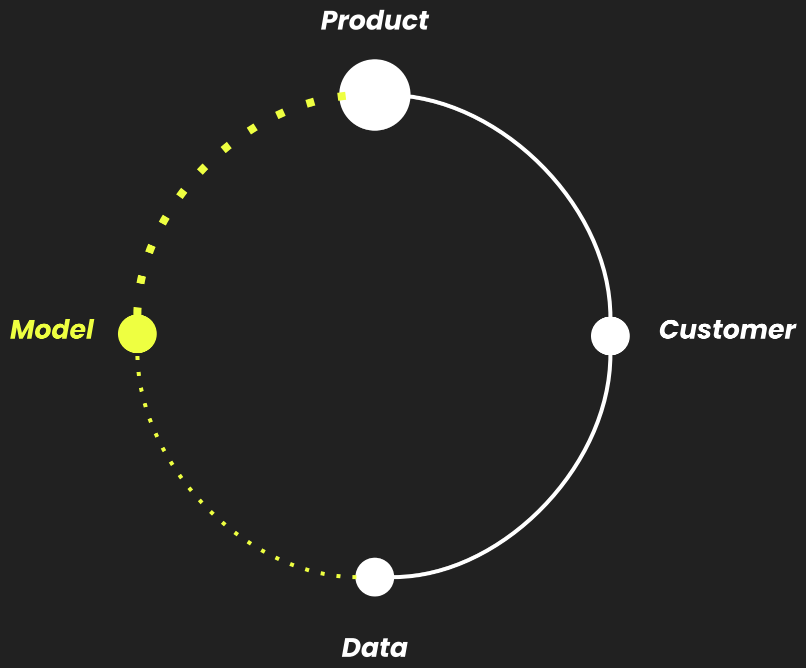

Every customer brings new data, and new data improves the model. Ultimately improving the product in a self-sustaining learning loop.

Driving ROI from data products is a long journey with deep valleys

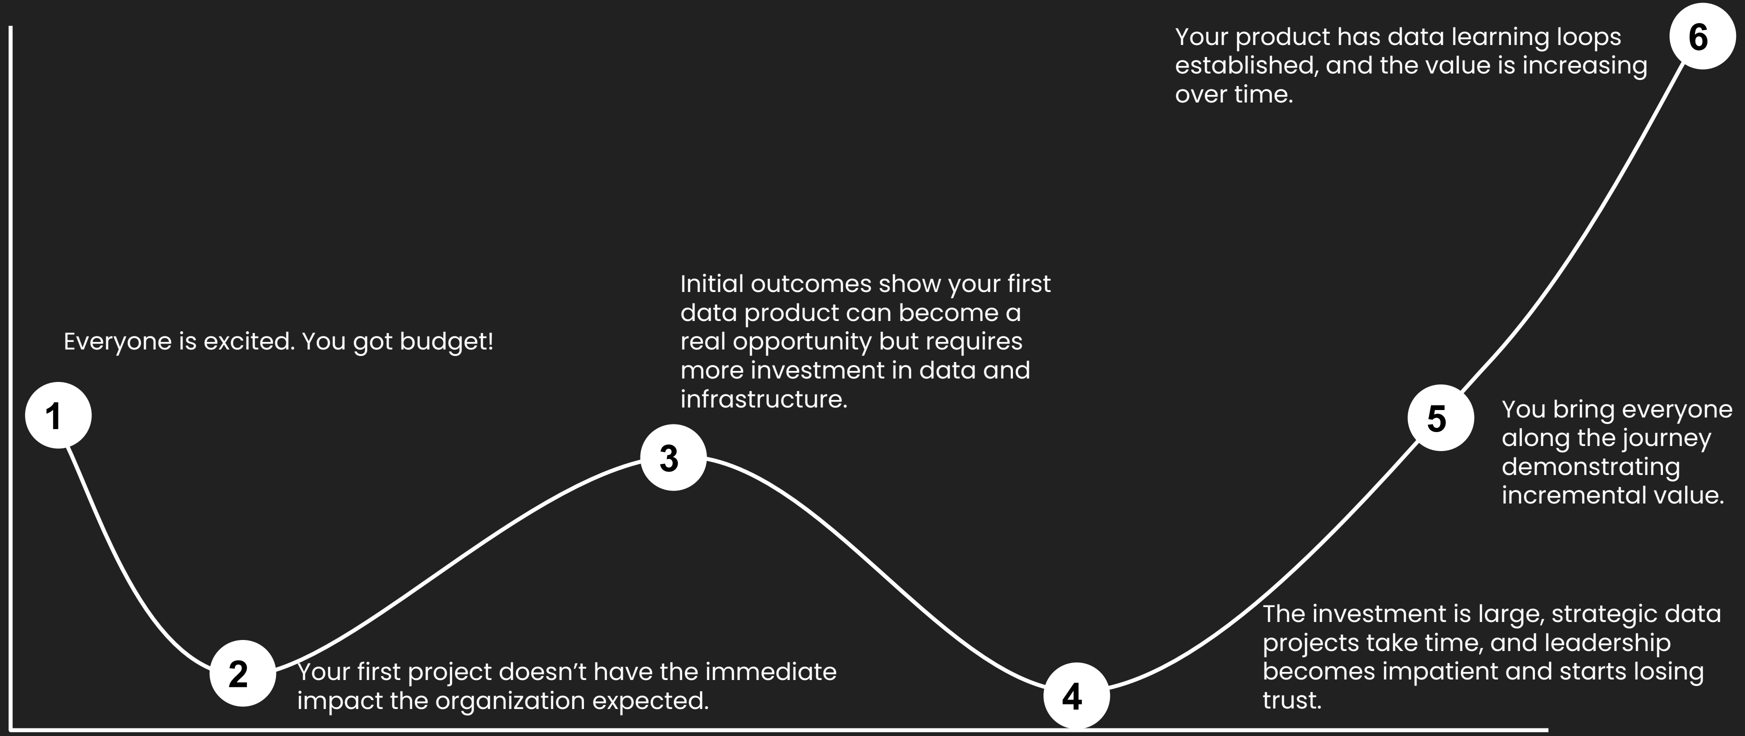

What you need to get right

1. A vision, and product strategy with a clear path to enabling data learning loops.
2. A northstar with a goldilocks starting point.
3. Stakeholder management through customer outcomes.

Building a strategy

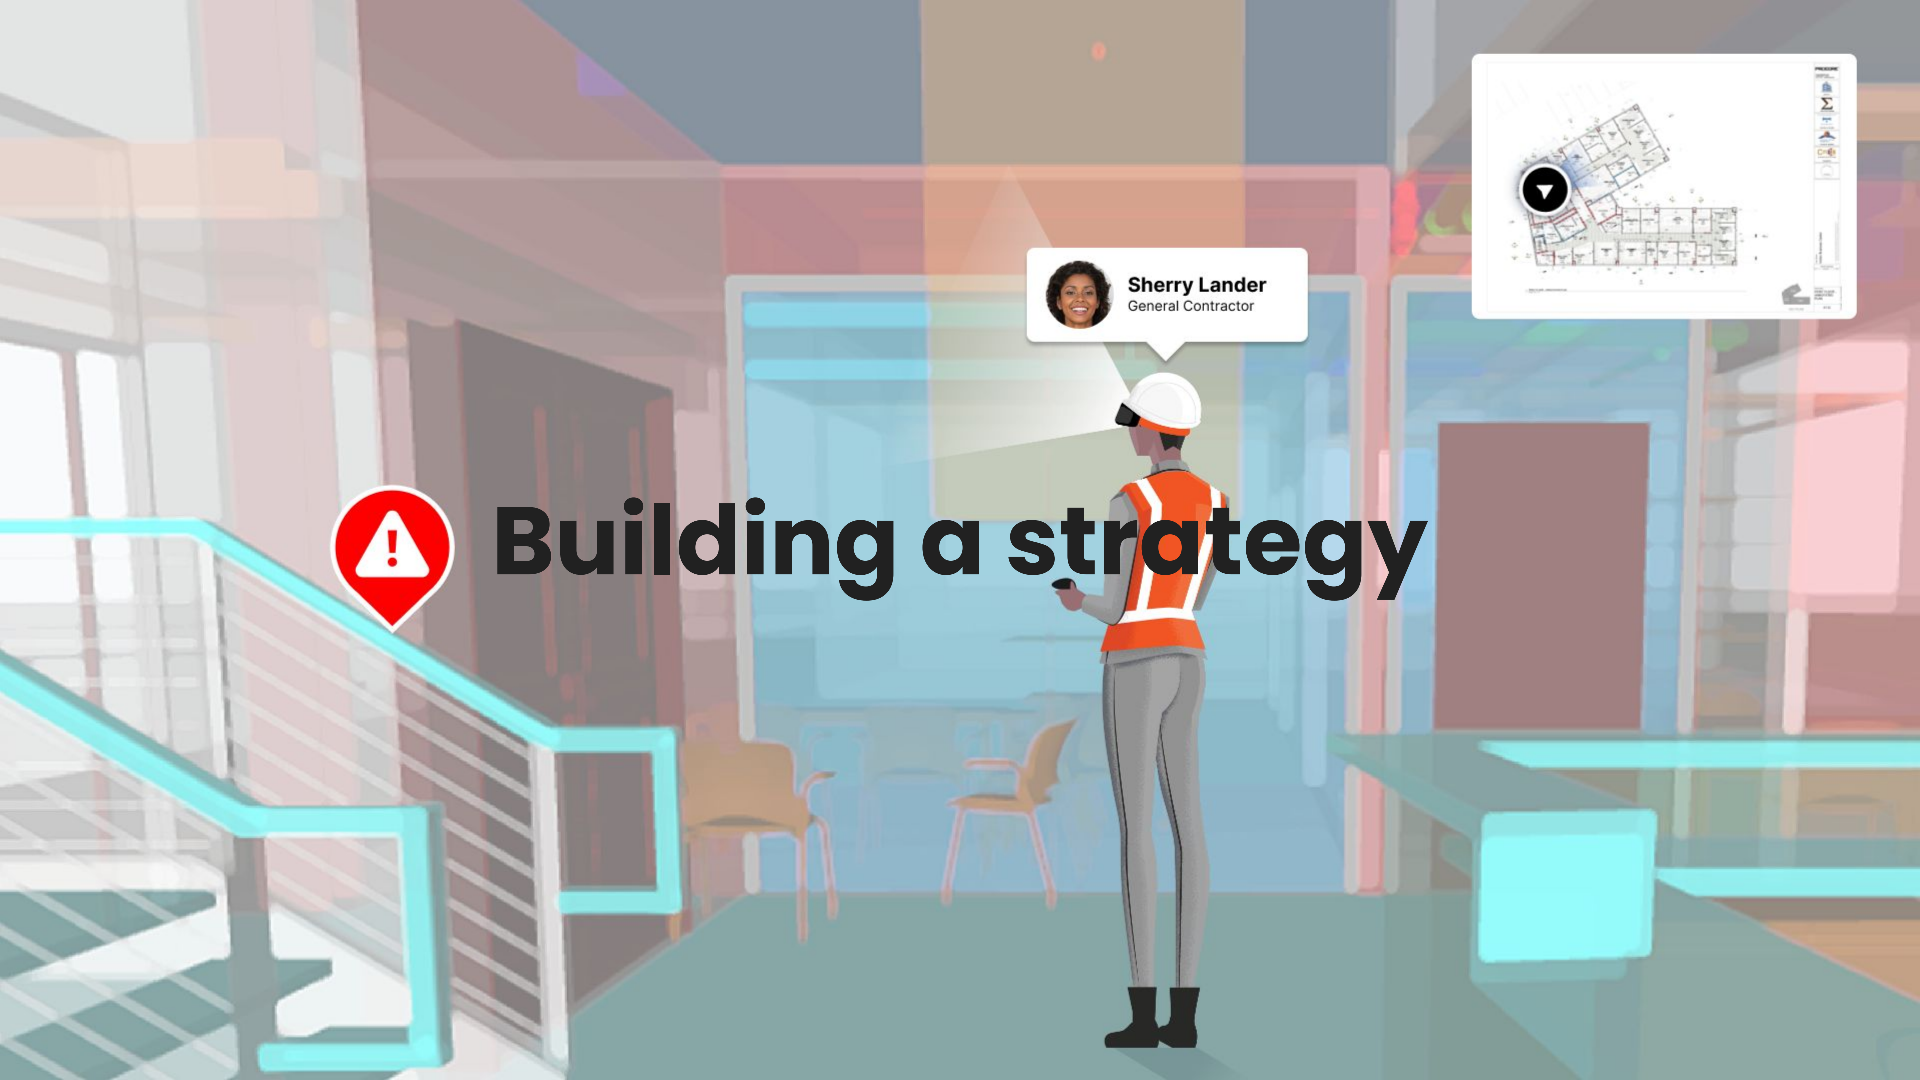

Vision + Tactics

What are you going to do to win?

1. What is your product **vision?**
2. What are the top **tactics** to pursue?

To connect everyone in construction on a single Platform.

In the next 3 years we will...

1. Connect All Stakeholders.
2. Accelerate growth with financial products.
3. Win Preconstruction.

Identify your north star questions

1. **Connect All Stakeholders.**

   What are the best companies I can hire for this project?

2. **Accelerate growth with financial products.**

   What do we need to do to ensure highest profit on this project?

3. **Win Preconstruction.**

   Do I have a competitive bid to win this project?

North stars questions are a goal post

They frame the problem, and challenge your team to take three critical steps:

1. Understand how customers attempt to answer these questions today.
2. Identify the data needed to answer these questions.
3. Create systems to acquire, process, and prepare the data.

Choosing your Goldilocks

Don’t burn your shots

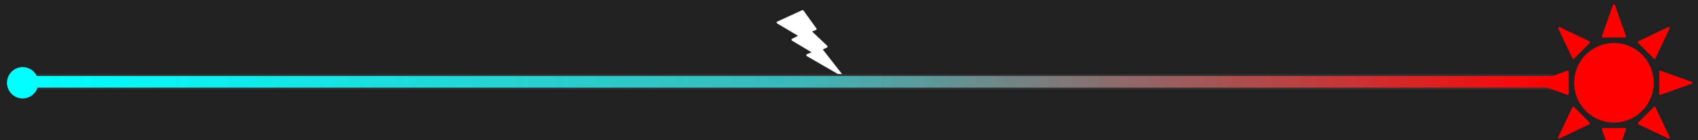

Your first project must not be too far from your goal posts, nor too close. You must be able to identify incremental progress and have room for failure and experimentation.

Your first data project

1. A real customer problem tied to one of your north star questions.
2. Data readily available from one data source.
3. Measurable from the onset.
4. Should test your end to end development cycle.
5. Removed from a critical customer workflow.

A construction Goldilocks project

North star: *what do we need to do to ensure highest profit on this project?*

1. *Is my contingency budget right for this project?*
2. *Are we at risk of general conditions funds overspend?*
3. *What type of projects are more profitable for us?*

Build, define success, and measure

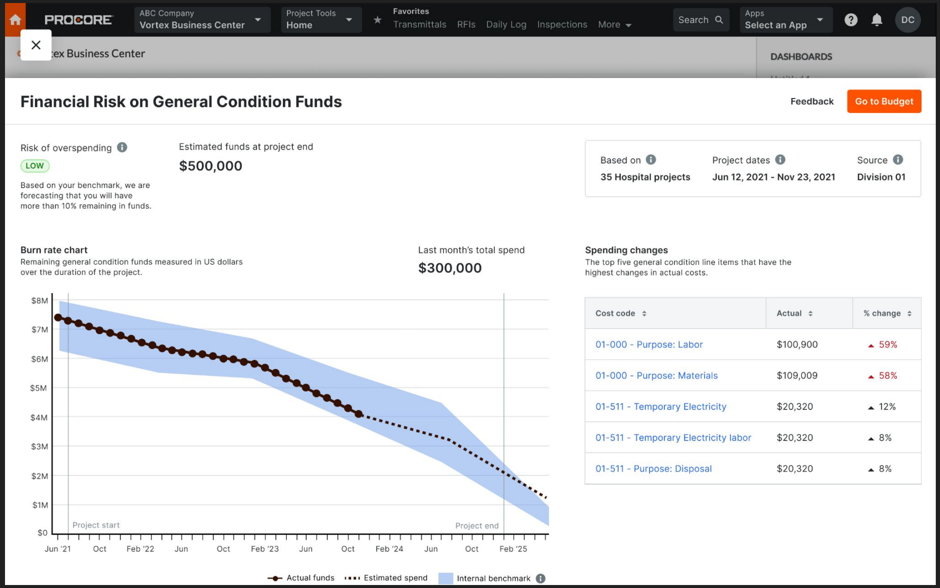

% of users who take direct action in the budget tool.

Build learning loops with your users

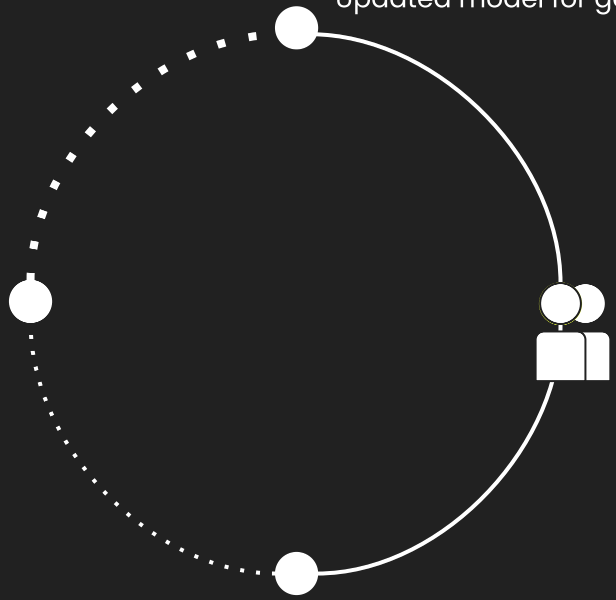

Updated model for general conditions

Retrain model: new data + feature tuning

New customers or existing customers start new projects with new budgets.

1. Additional general conditions data is added to the system.

2. Users provide feedback on model predictions


Manage your
stakeholders

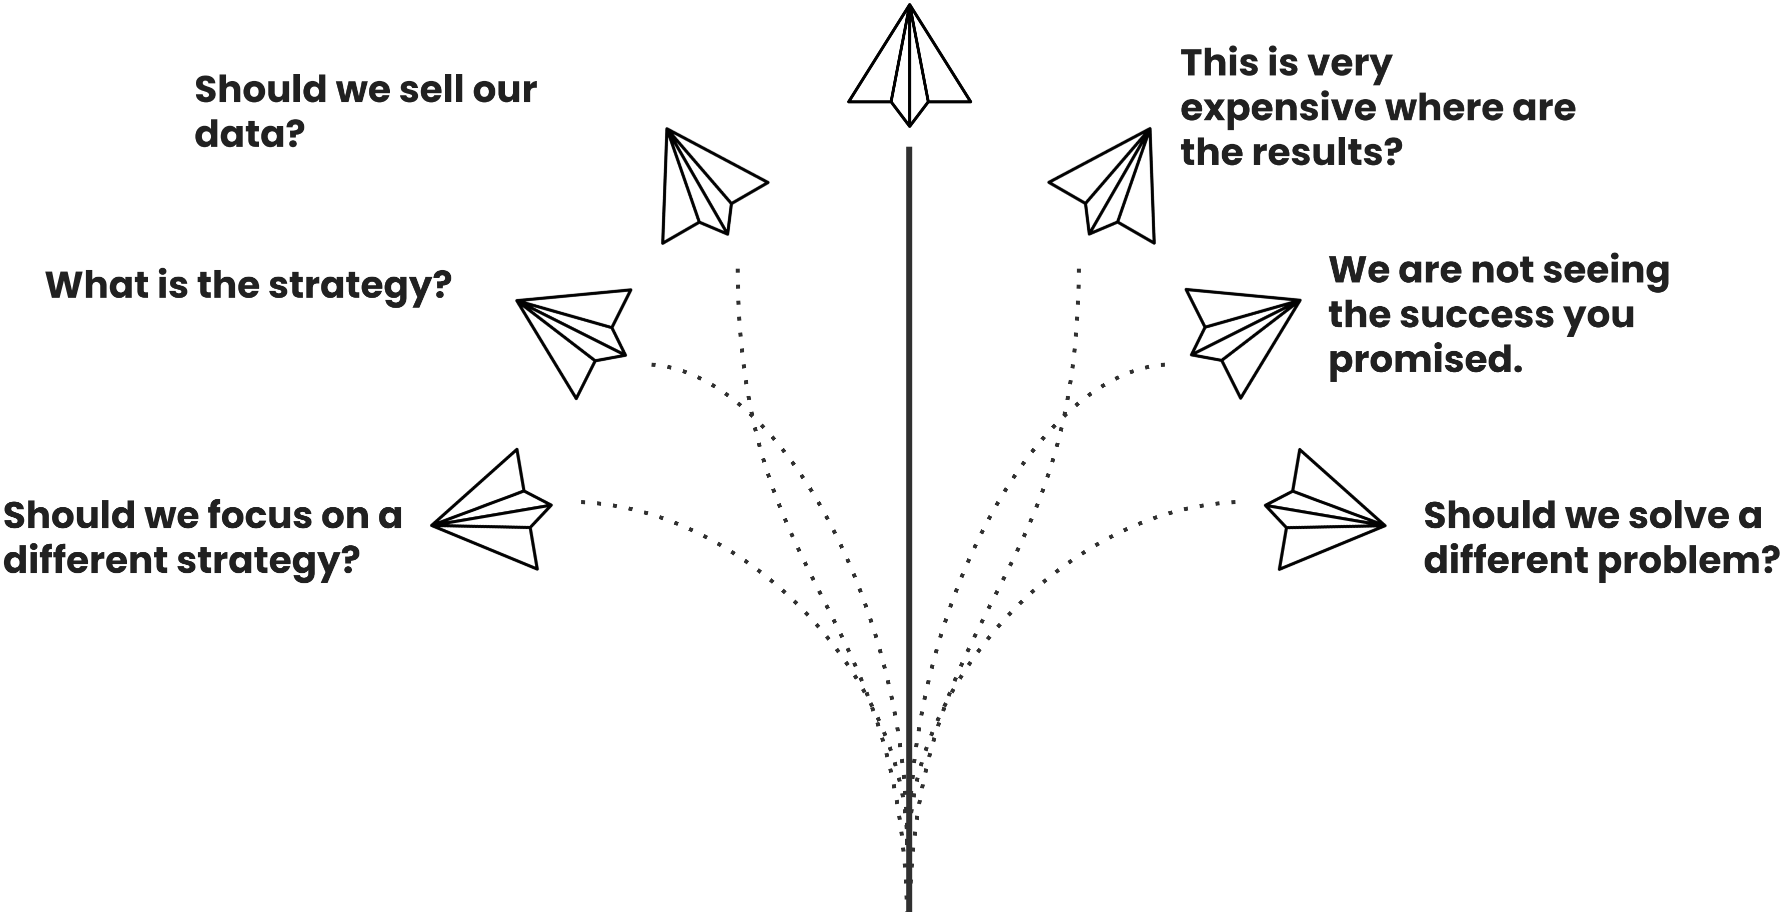
Should we sell our data?
What is the strategy?
Should we focus on a different strategy?
This is very expensive where are the results?
We are not seeing the success you promised.
Should we solve a different problem?

Whittle out distractions ...

1. New opportunities will emerge. Keep your organization and your team focused on the north stars of the strategy.
2. Communicate often. Use your goldilocks project as a means for engaging your stakeholders throughout the process.
3. Keep your team motivated with the "**why**". Bring customer stories to life with your goldilocks project.