DRISHTI

How MLOps Helped AI Rise

A from-the-trenches view of MLOps

MANUFACTURING LEADERSHIP COUNCIL
2021 Awards for Ford, DENSO + Hella

THE TOYOTA WAY 2ND EDITION
"A revolution in TPS"

Forbes
2020 Forbes AI 50

NVIDIA
2020 Top AI Startup

WORLD ECONOMIC FORUM
2019 Tech Pioneer

Dr. Prasad Akella
Founder & Chairman

Sujay Narumanchi
Founding Engineer

A picture is worth a thousand words.

A video is worth a million.

DRISHTI

Software 2.0: It's all about data

ML systems are fundamentally different

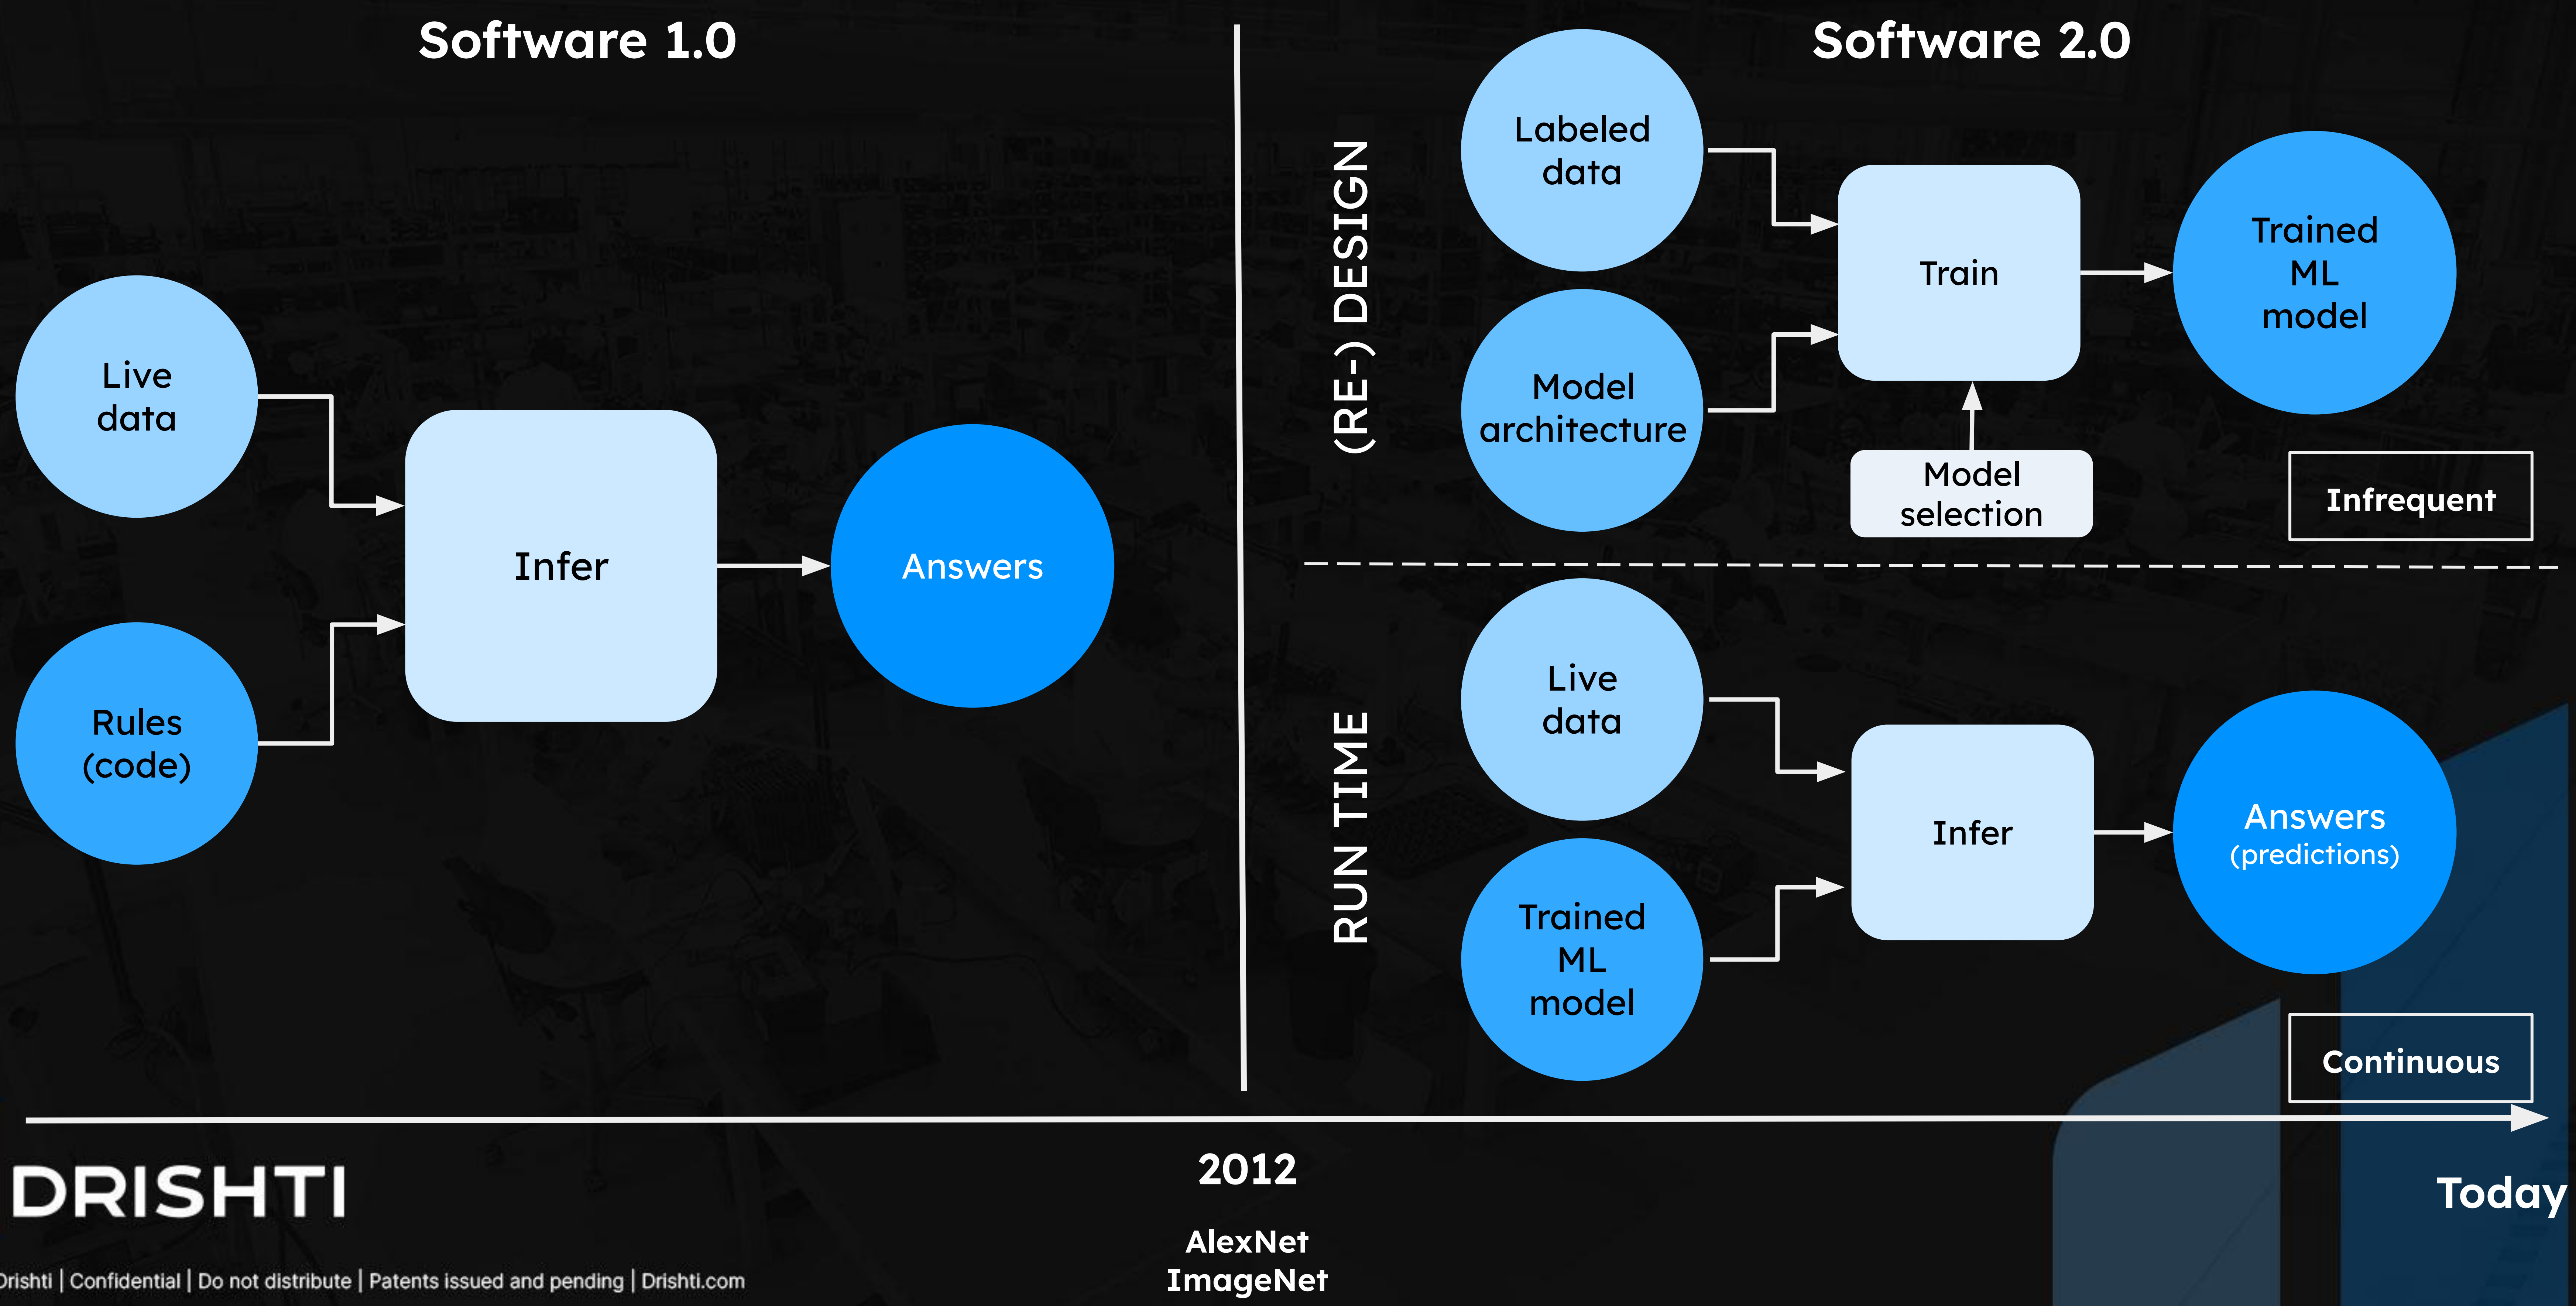

DRISHTI

Software 2.0: The building blocks

The ML lifecycle from data to insights

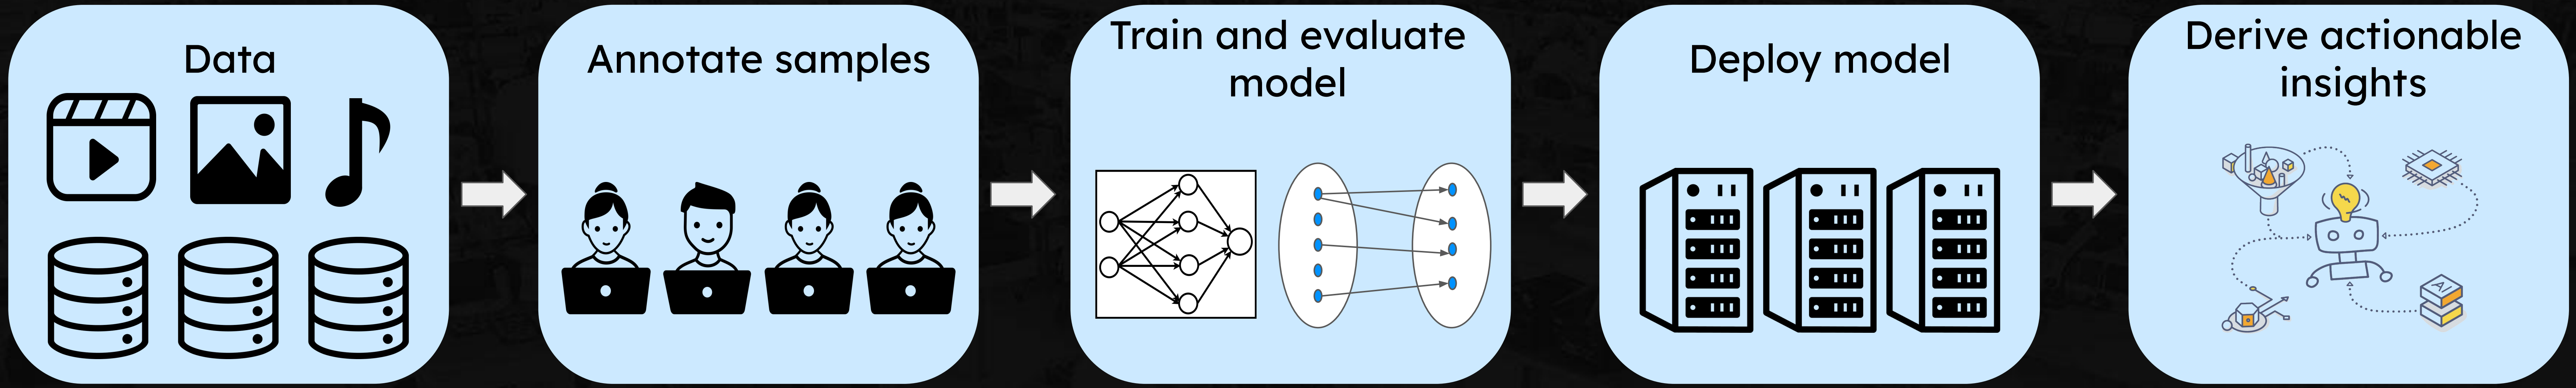

Where does it break?

DRISHTI

Many challenges, one source: Data quality

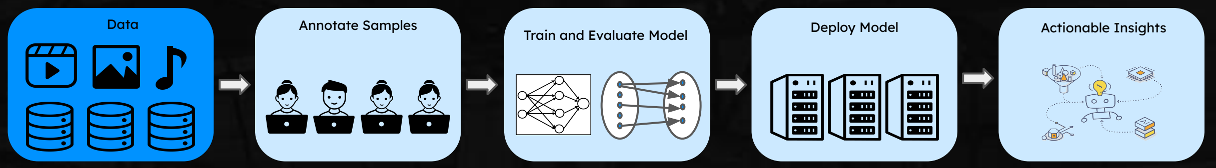

1 Data holes: network transfer, disk I/O corruption

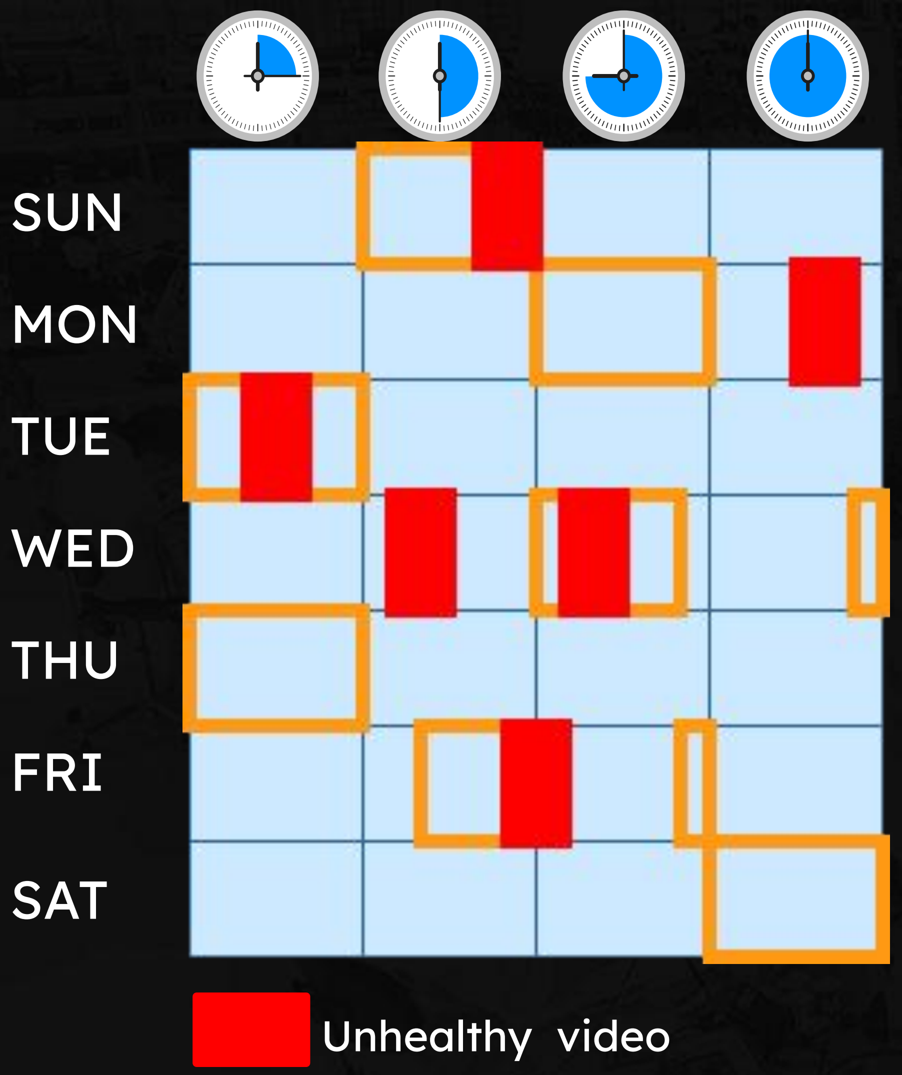

Robust data pipeline, network and edge monitoring

2 The world looks different: camera viewpoint changes

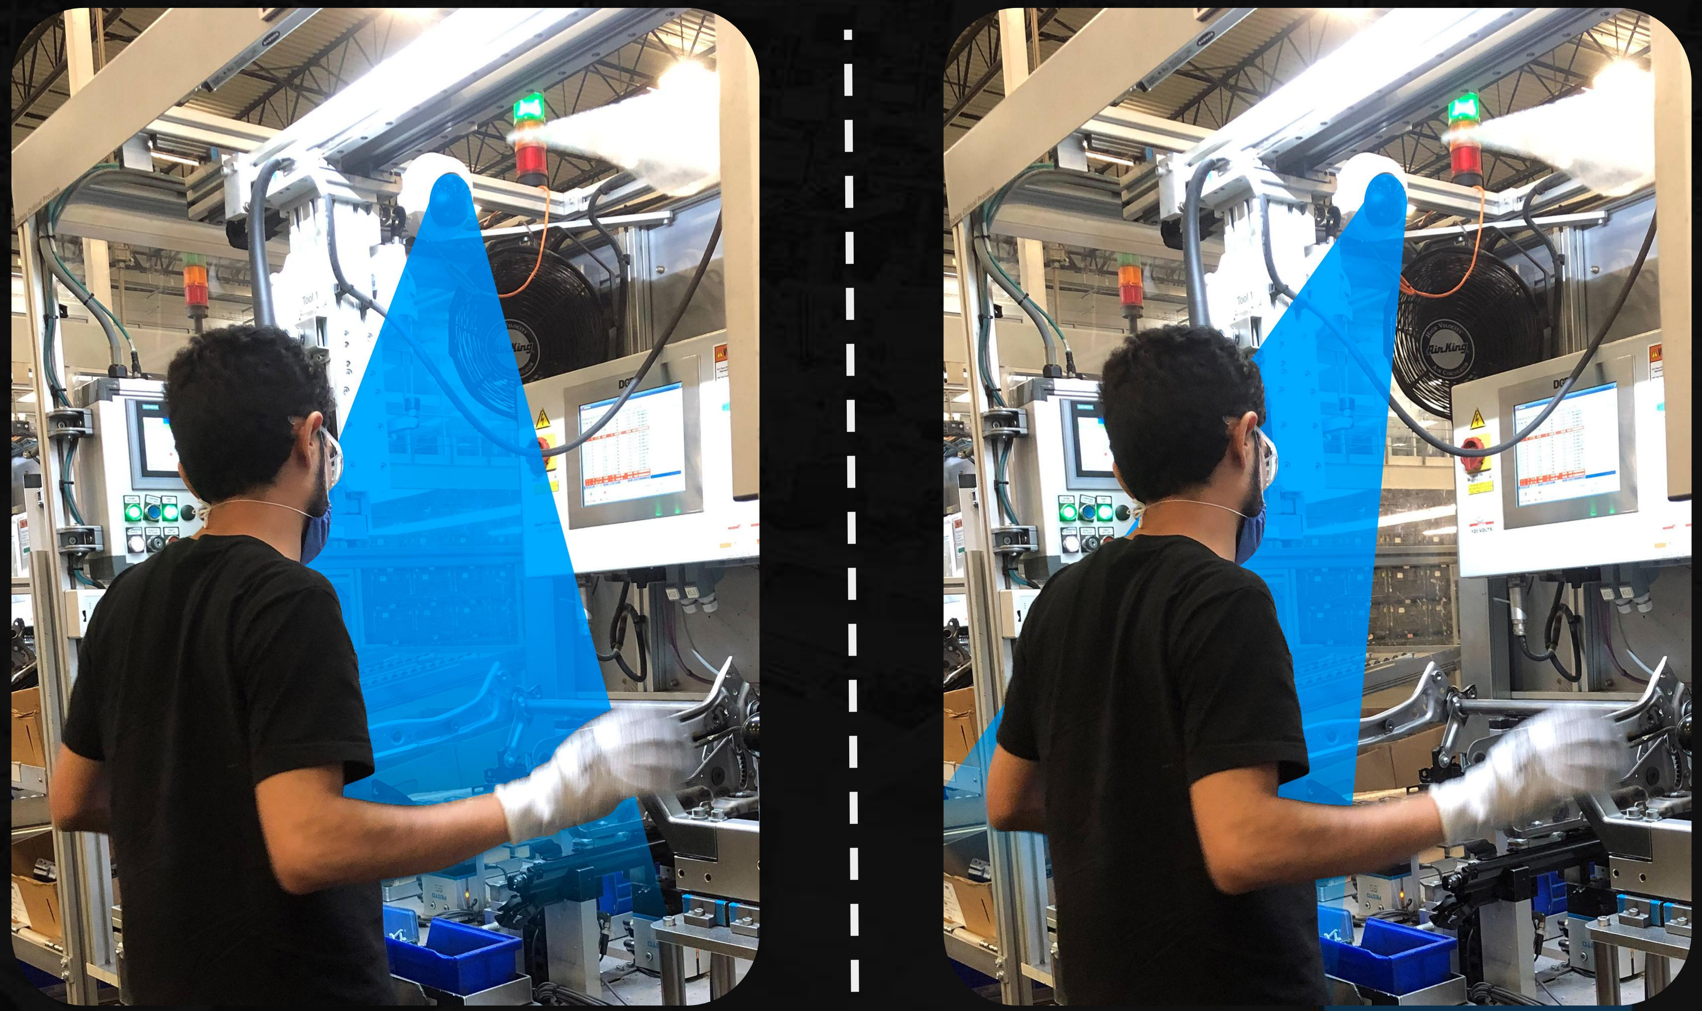

Camera auto-calibration and monitoring

DRISHTI

Ambiguity is omnipresent: Labeling

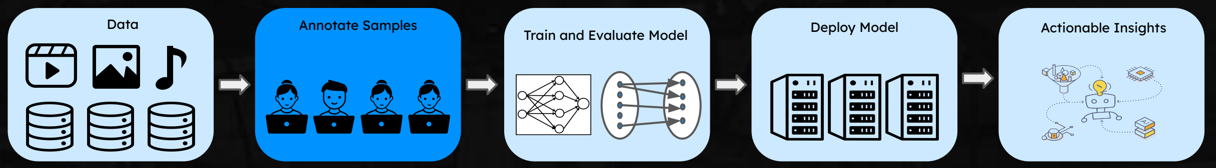

1 Garbage in, garbage out: Labeling accuracy

Annotator 1

Unit

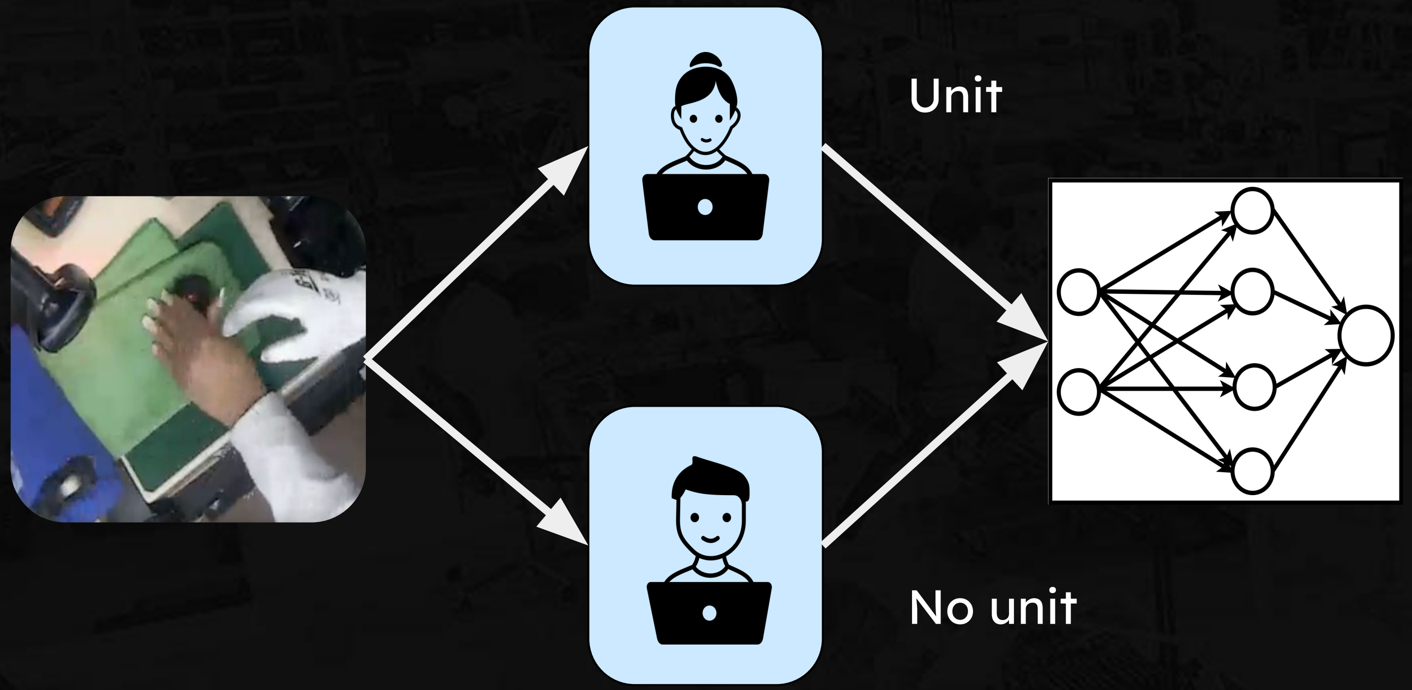

Confused neural net

???

No unit

Annotator 2

Audit workflows for QC, redundancy + 3-way voting

2 Ambiguity in task definition

Unit

Unit ?

Occluded unit

Multiple units

Partial unit in view

Tighter collaboration between ML and Annotation teams

Document standard processes, edge cases

DRISHTI

Much variance, much data: Training

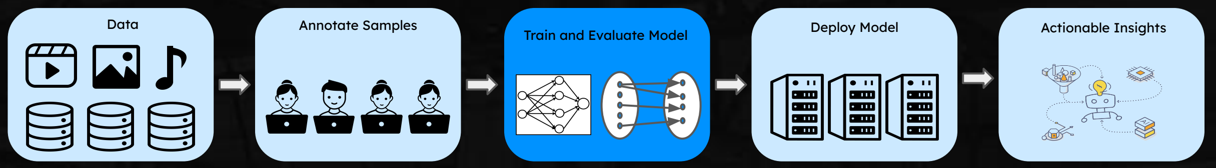

1 Extreme variance in the real world

Variation in unit size | Indeterminate location | Multiple units in field of view

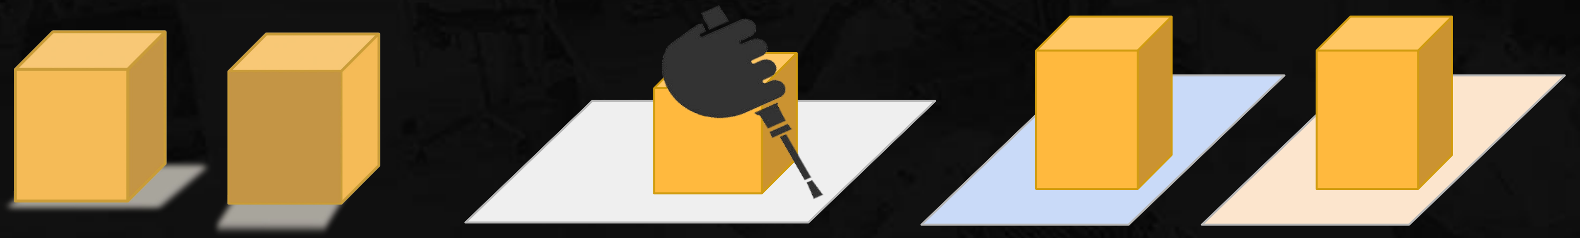

Lighting changes | Occlusion | Background variation

Smart data sampling

Generalizable neural networks

2 Data hungry deep learning models

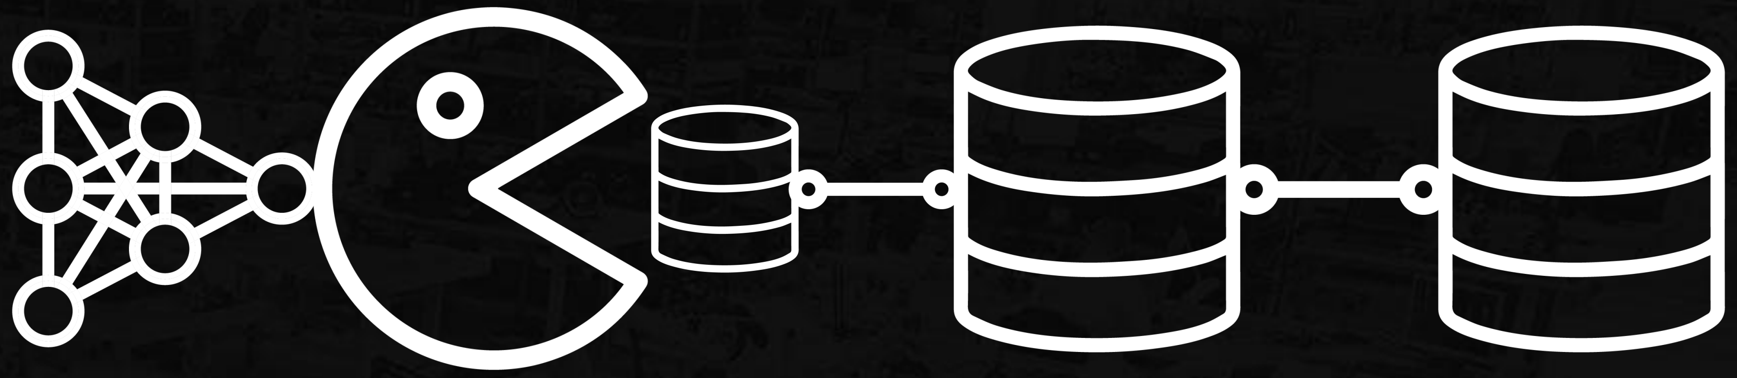

Neural network

Data

Data costs ↑

Annotation resources ↑

Time to deploy ↑

Semi-supervised learning

Unsupervised learning

DRISHTI

Change is the order of the day: Data drift

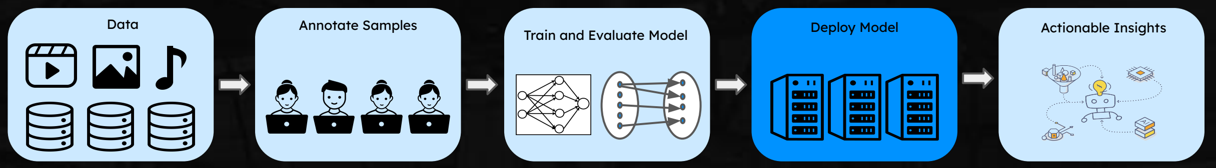

1 Can you spot the difference?

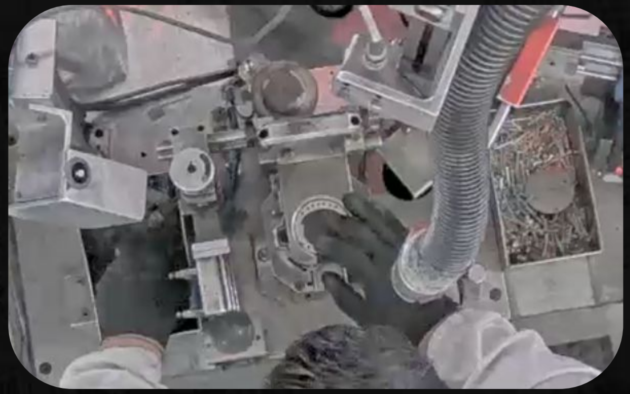

Black gloves

Black gloves

Black gloves

Gray gloves

2 System accuracy reduces

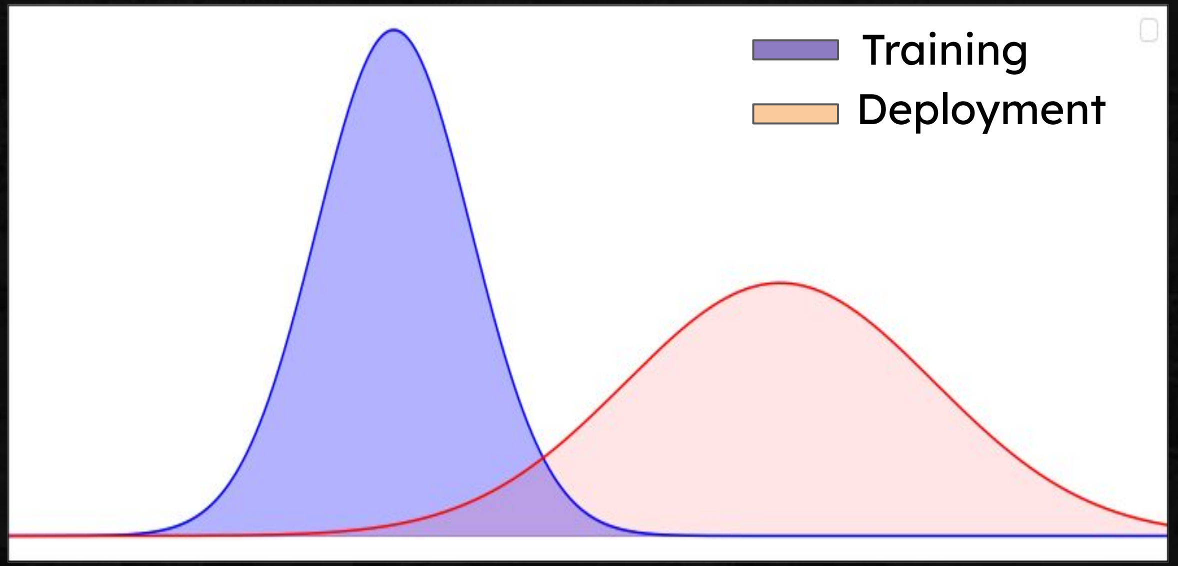

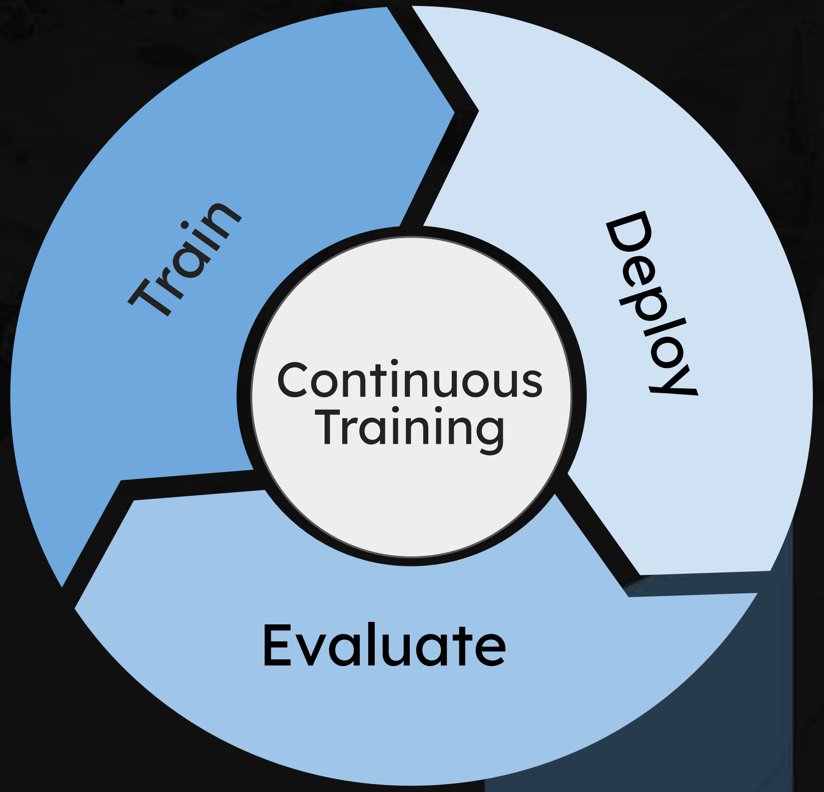

DRISHTI

From 99% to 99.9% accuracy: Reliable insights

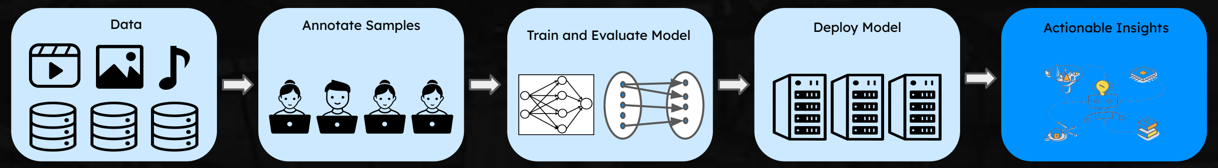

1 Ensembling models

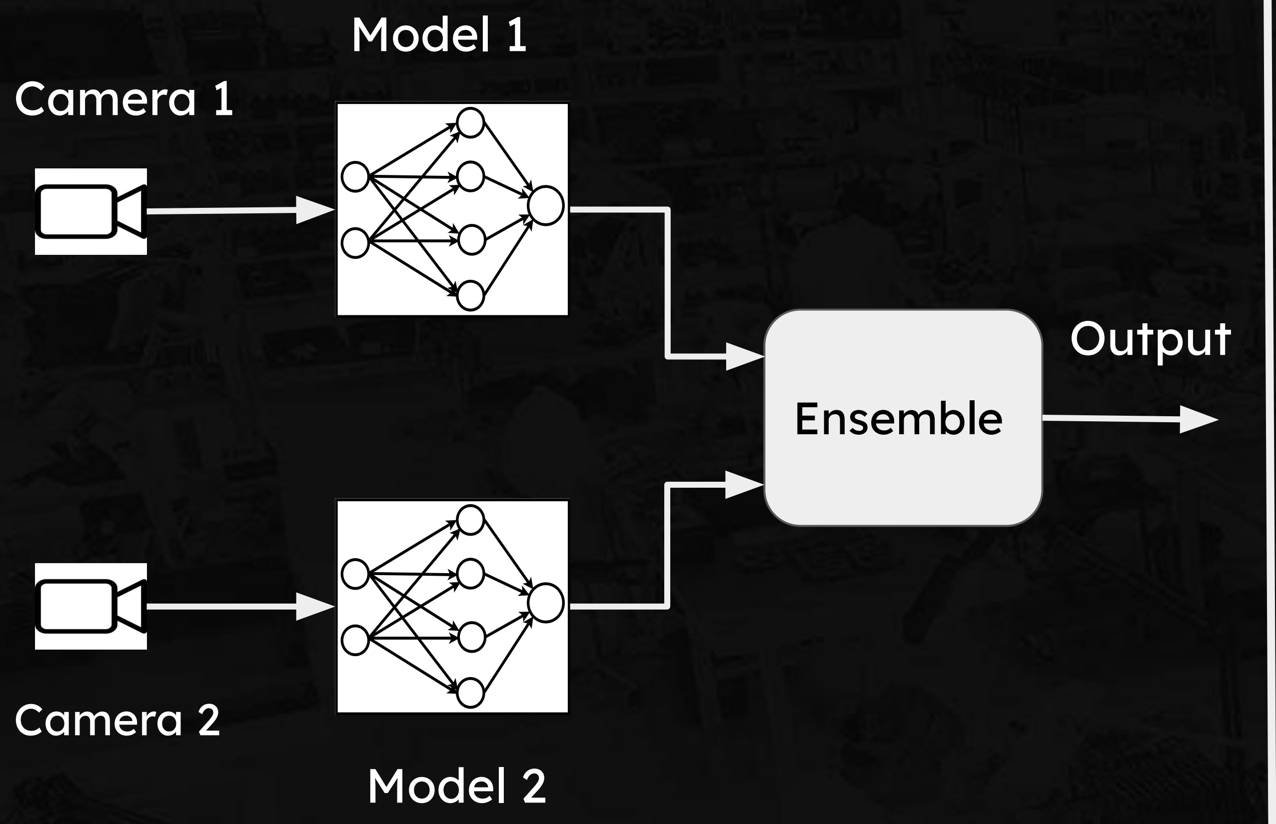

Suite of models ⇒ Model management

2 Heuristics: Finite state machines

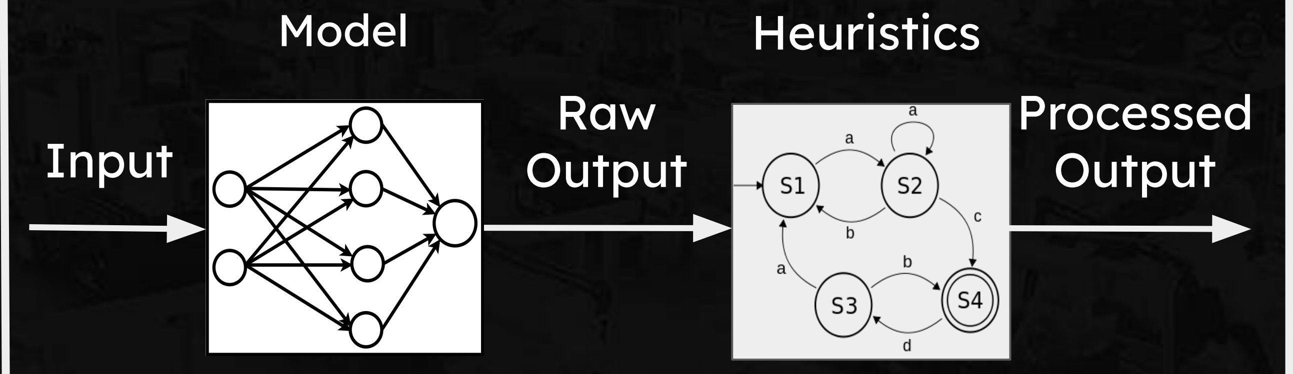

Heuristics configuration management

3 Human in the loop

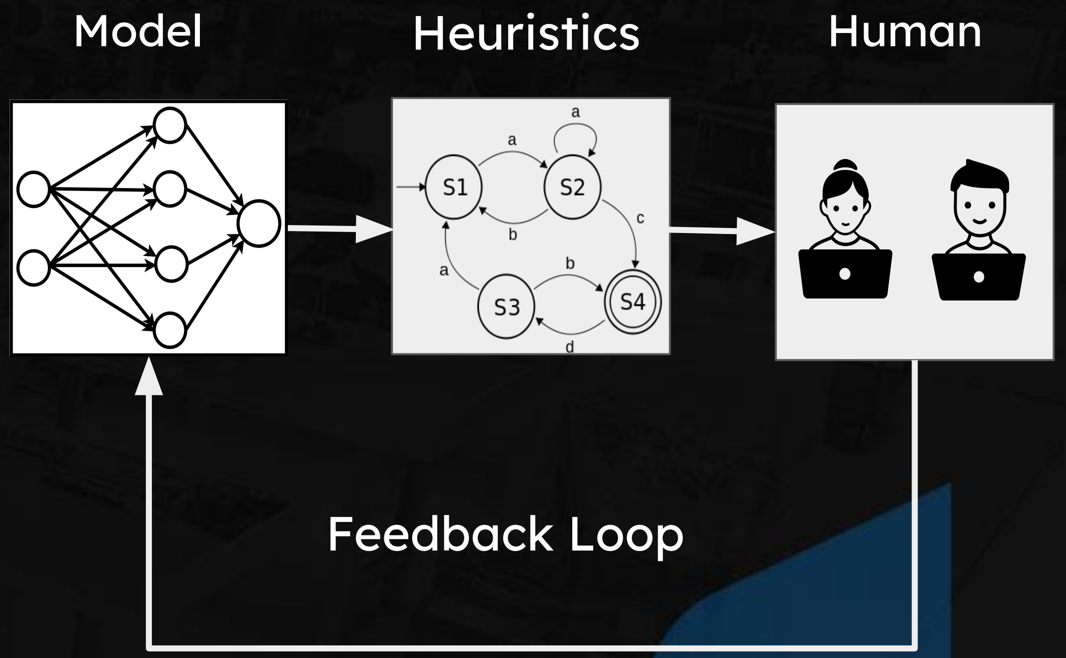

Correct mistakes, feed back to model

Software 2.0: Towards self-service AI

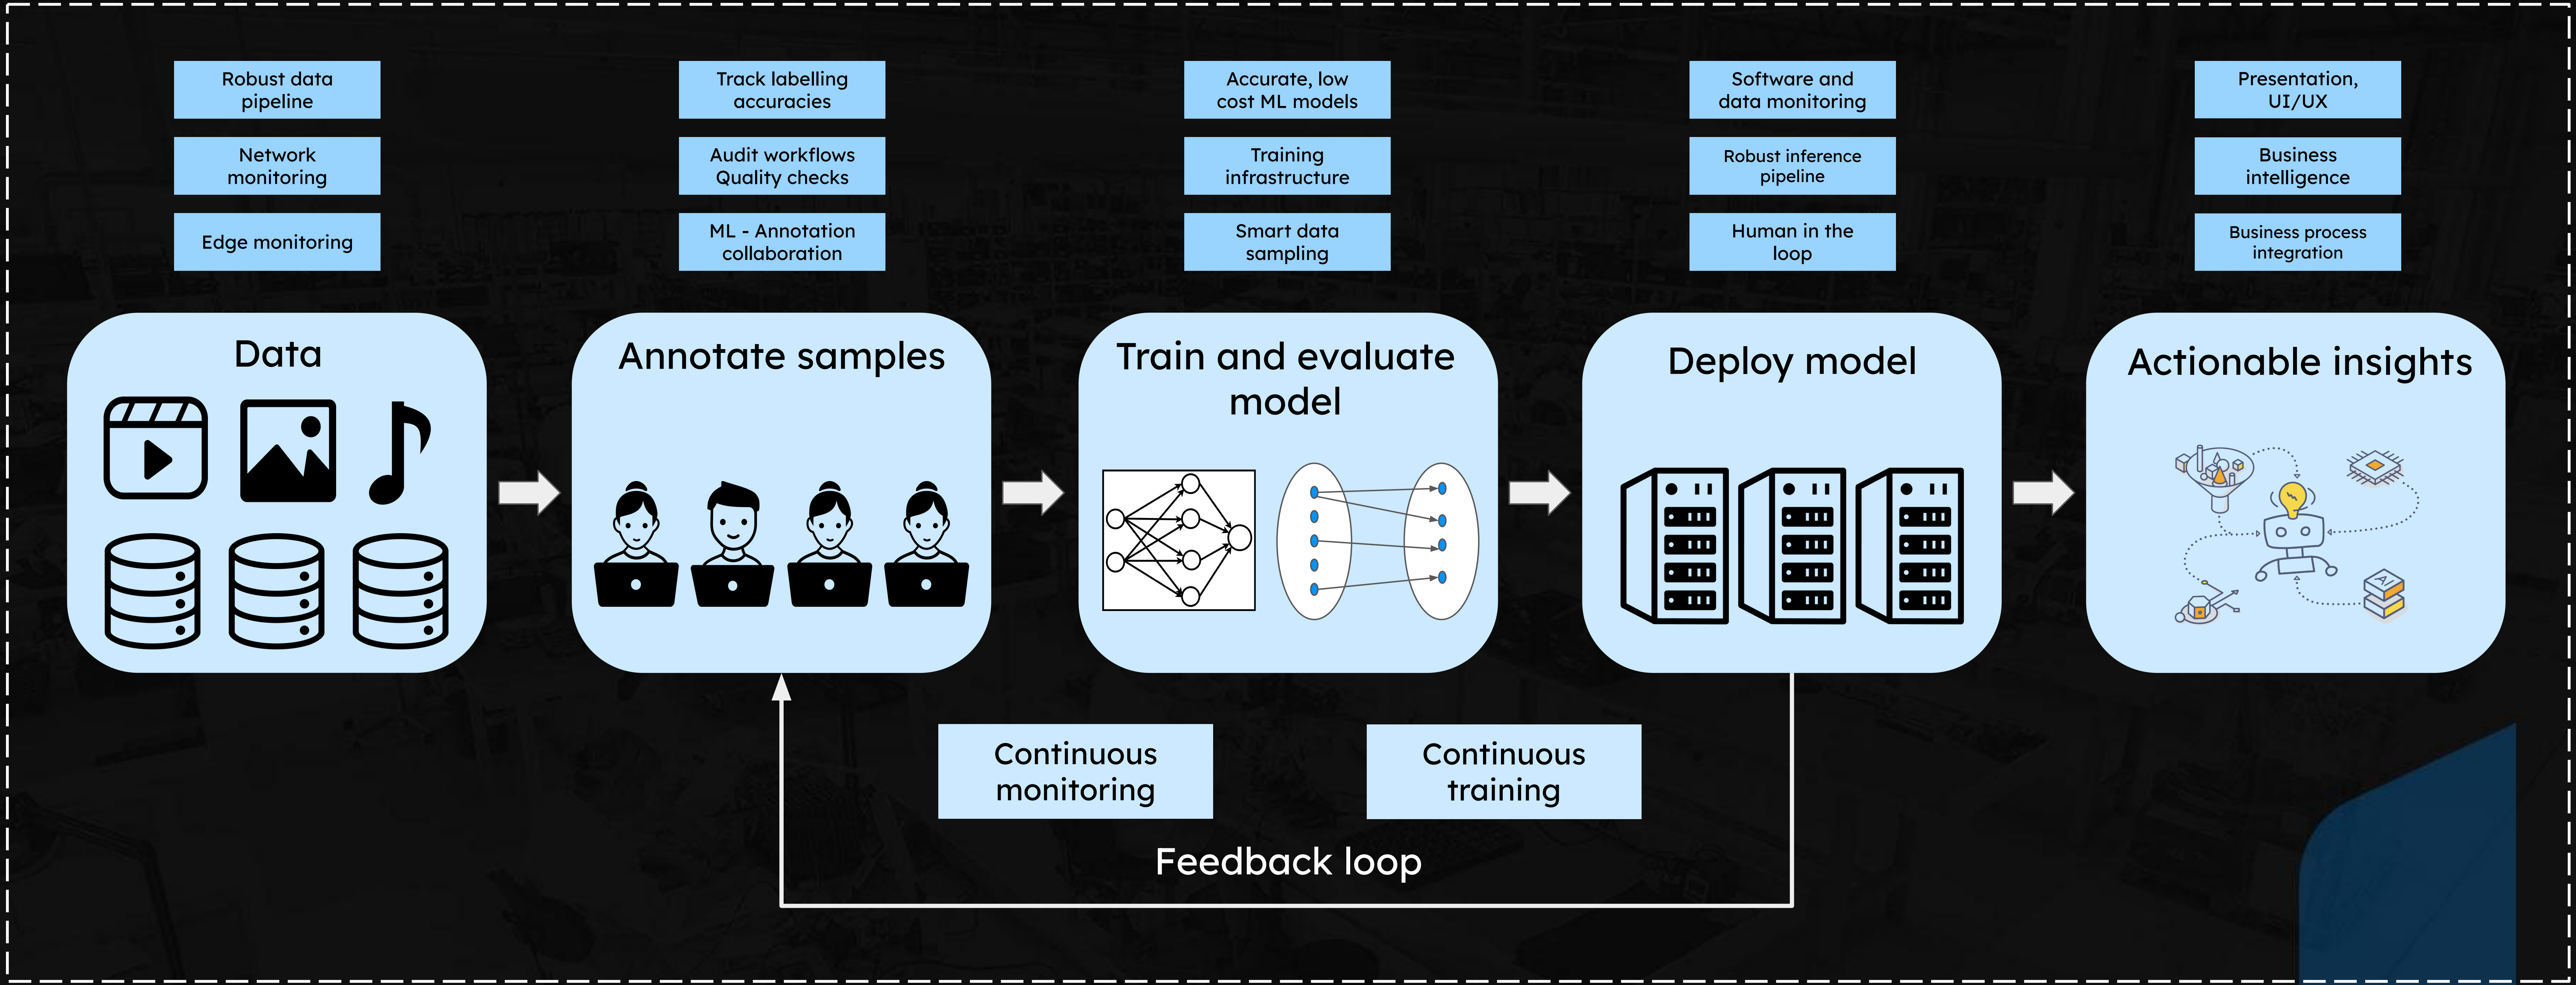

What if the customer wants control? → Self-service AI

DRISHTI

Software 2.0: Your team has to be organized differently

The workflows determine the organization including the need for skills and proximity

Software 1.0

ENGINEERING

Development

Software

- Code
- Unit testing
- Integration testing

Quality

Software QA

- Functional testing
- Standards development

Ops

DevOps

- Release
- Deployment
- System monitoring

Software 2.0

ENGINEERING

Data

Labelling

- Data annotation
- HITL execution

Development

Software | ML

Software 1.0 + Software 2.0:

- ML algorithms
- ML tuning
- ML tooling

Quality

Software QA | Data QA

Software 1.0 + Software 2.0:

- Data testing
- Data standards

Ops

DevOps | MLOps

Software 1.0 + Software 2.0:

- Continuous monitoring
- Continuous training

DRISHTI

Introducing "The MLOps House"*

Real world changes:
Product
Process
People
Environment

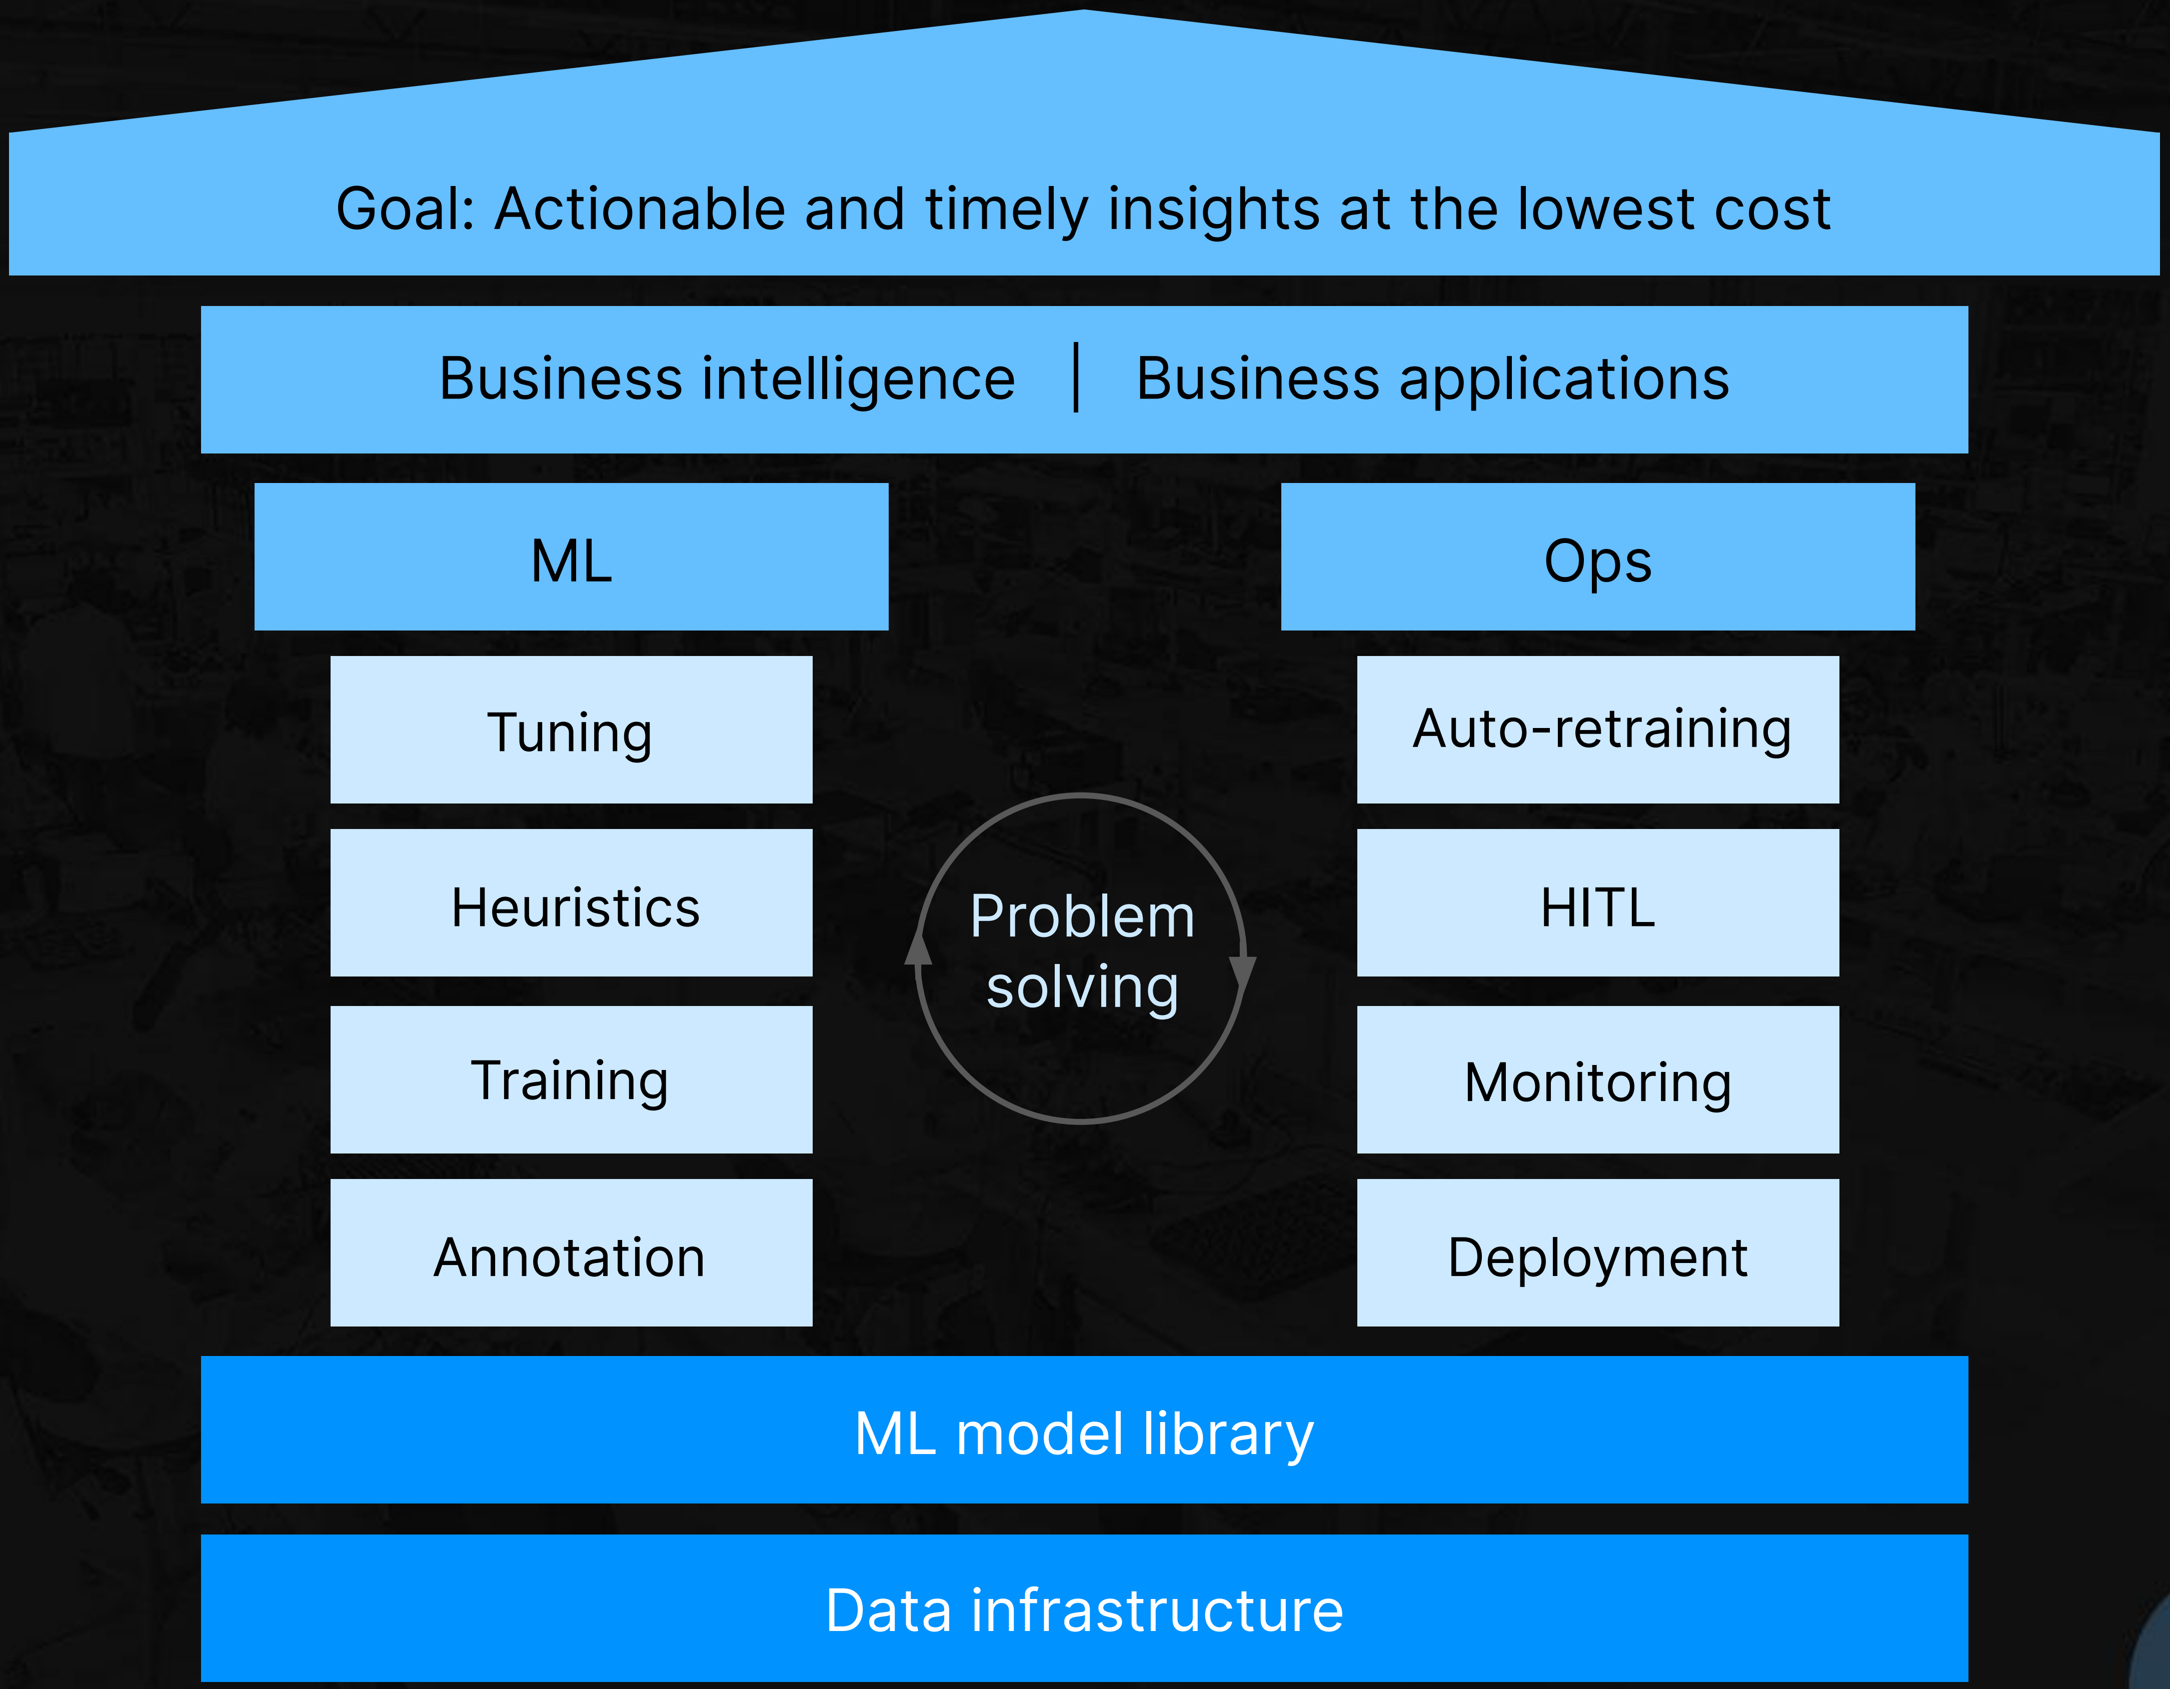

Produce impactful insights

Be robust to changes in a dynamic world

Stably create quality data

DRISHTI

*Credit: The Toyota Production System House

The implementer's roadmap

- **Design your entire ML system for YOUR use cases**
  - [Example] Inference accuracy drives the annotation quality and the neural network accuracy
- **Ensure that domain expertise permeates the teams**
  - Organize teams for seamless communication, especially if you are running a real time inferencing system
- **Details matter: Attention is all you need**
  - AI is non-intuitive and debugging is hard; so, paying attention to details will pay off in spades

DRISHTI

Reaching us

Dr. Prasad Akella
Founder & Chairman
p.akella@drishti.com

Sujay Narumanchi
Founding Engineer
s.narumanchi@drishti.com

DRISHTI

APPENDIX

MLOps: Bringing it all together

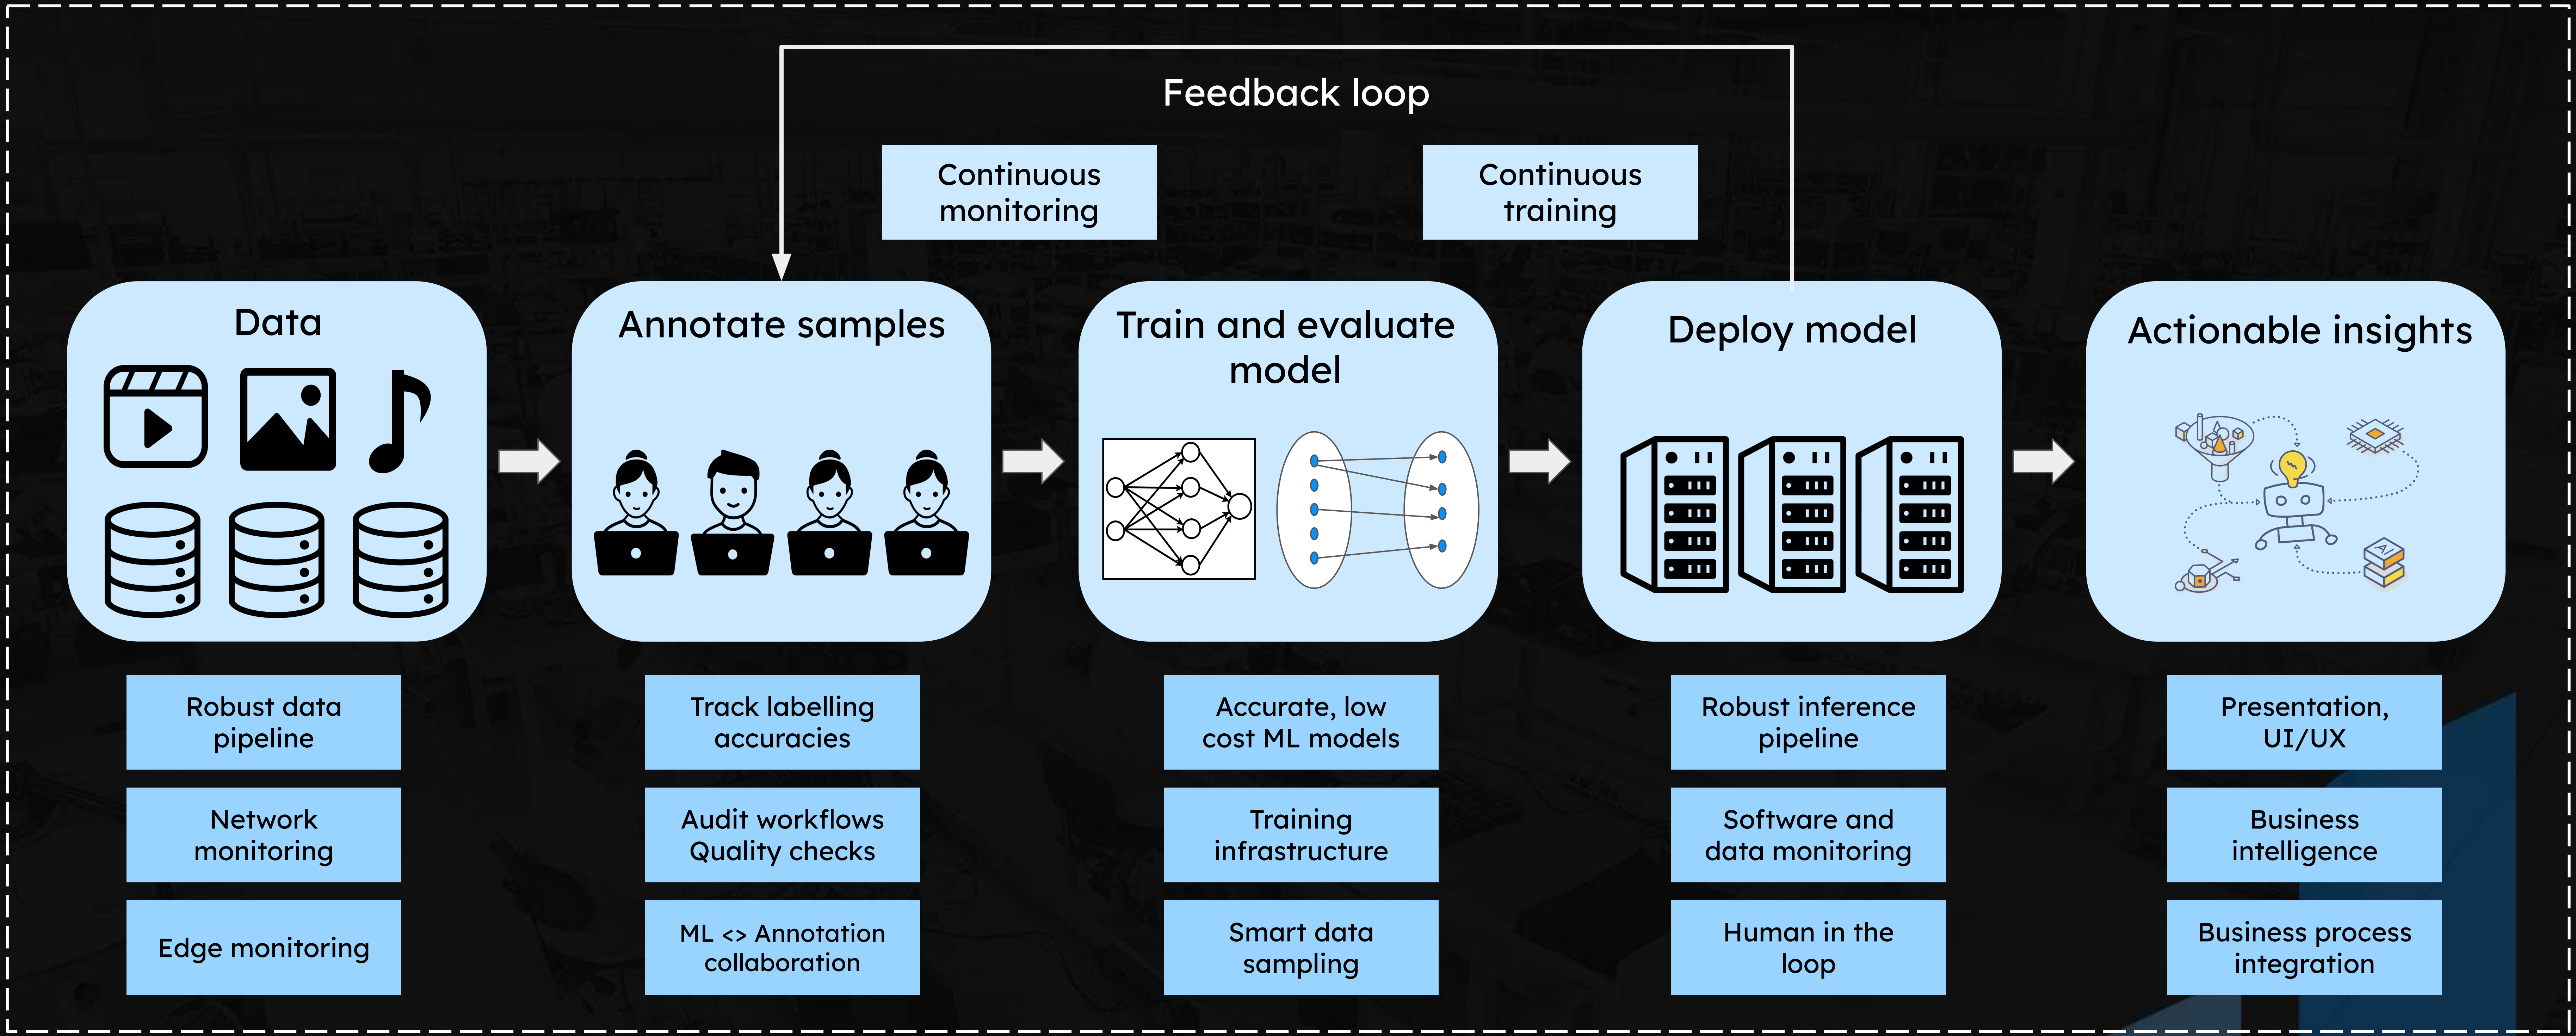

DRISHTI

What if the customer wants control? → Self-service AI

MLOps house — The vendor landscape

DRISHTI

MLOps House - Business view

Goal: Actionable, timely insights @ lowest cost

Deliver actionable insights integrated into business processes

ML to Model	Man & Machine	Ops to Track
Apply heuristics		Update heuristics
Infer	Monitor / Experiment	Use diff model
Train		Retrain
Annotate	Scientific experimentation	New data

Build accurate ML models

Construct a wholesome data infra

Product
Process
People
Environment

To produce impactful insights

Robust in the face of a dynamic world

Creating quality & timely data (data operations stability)

DRISHTI

MLOps House - Business view

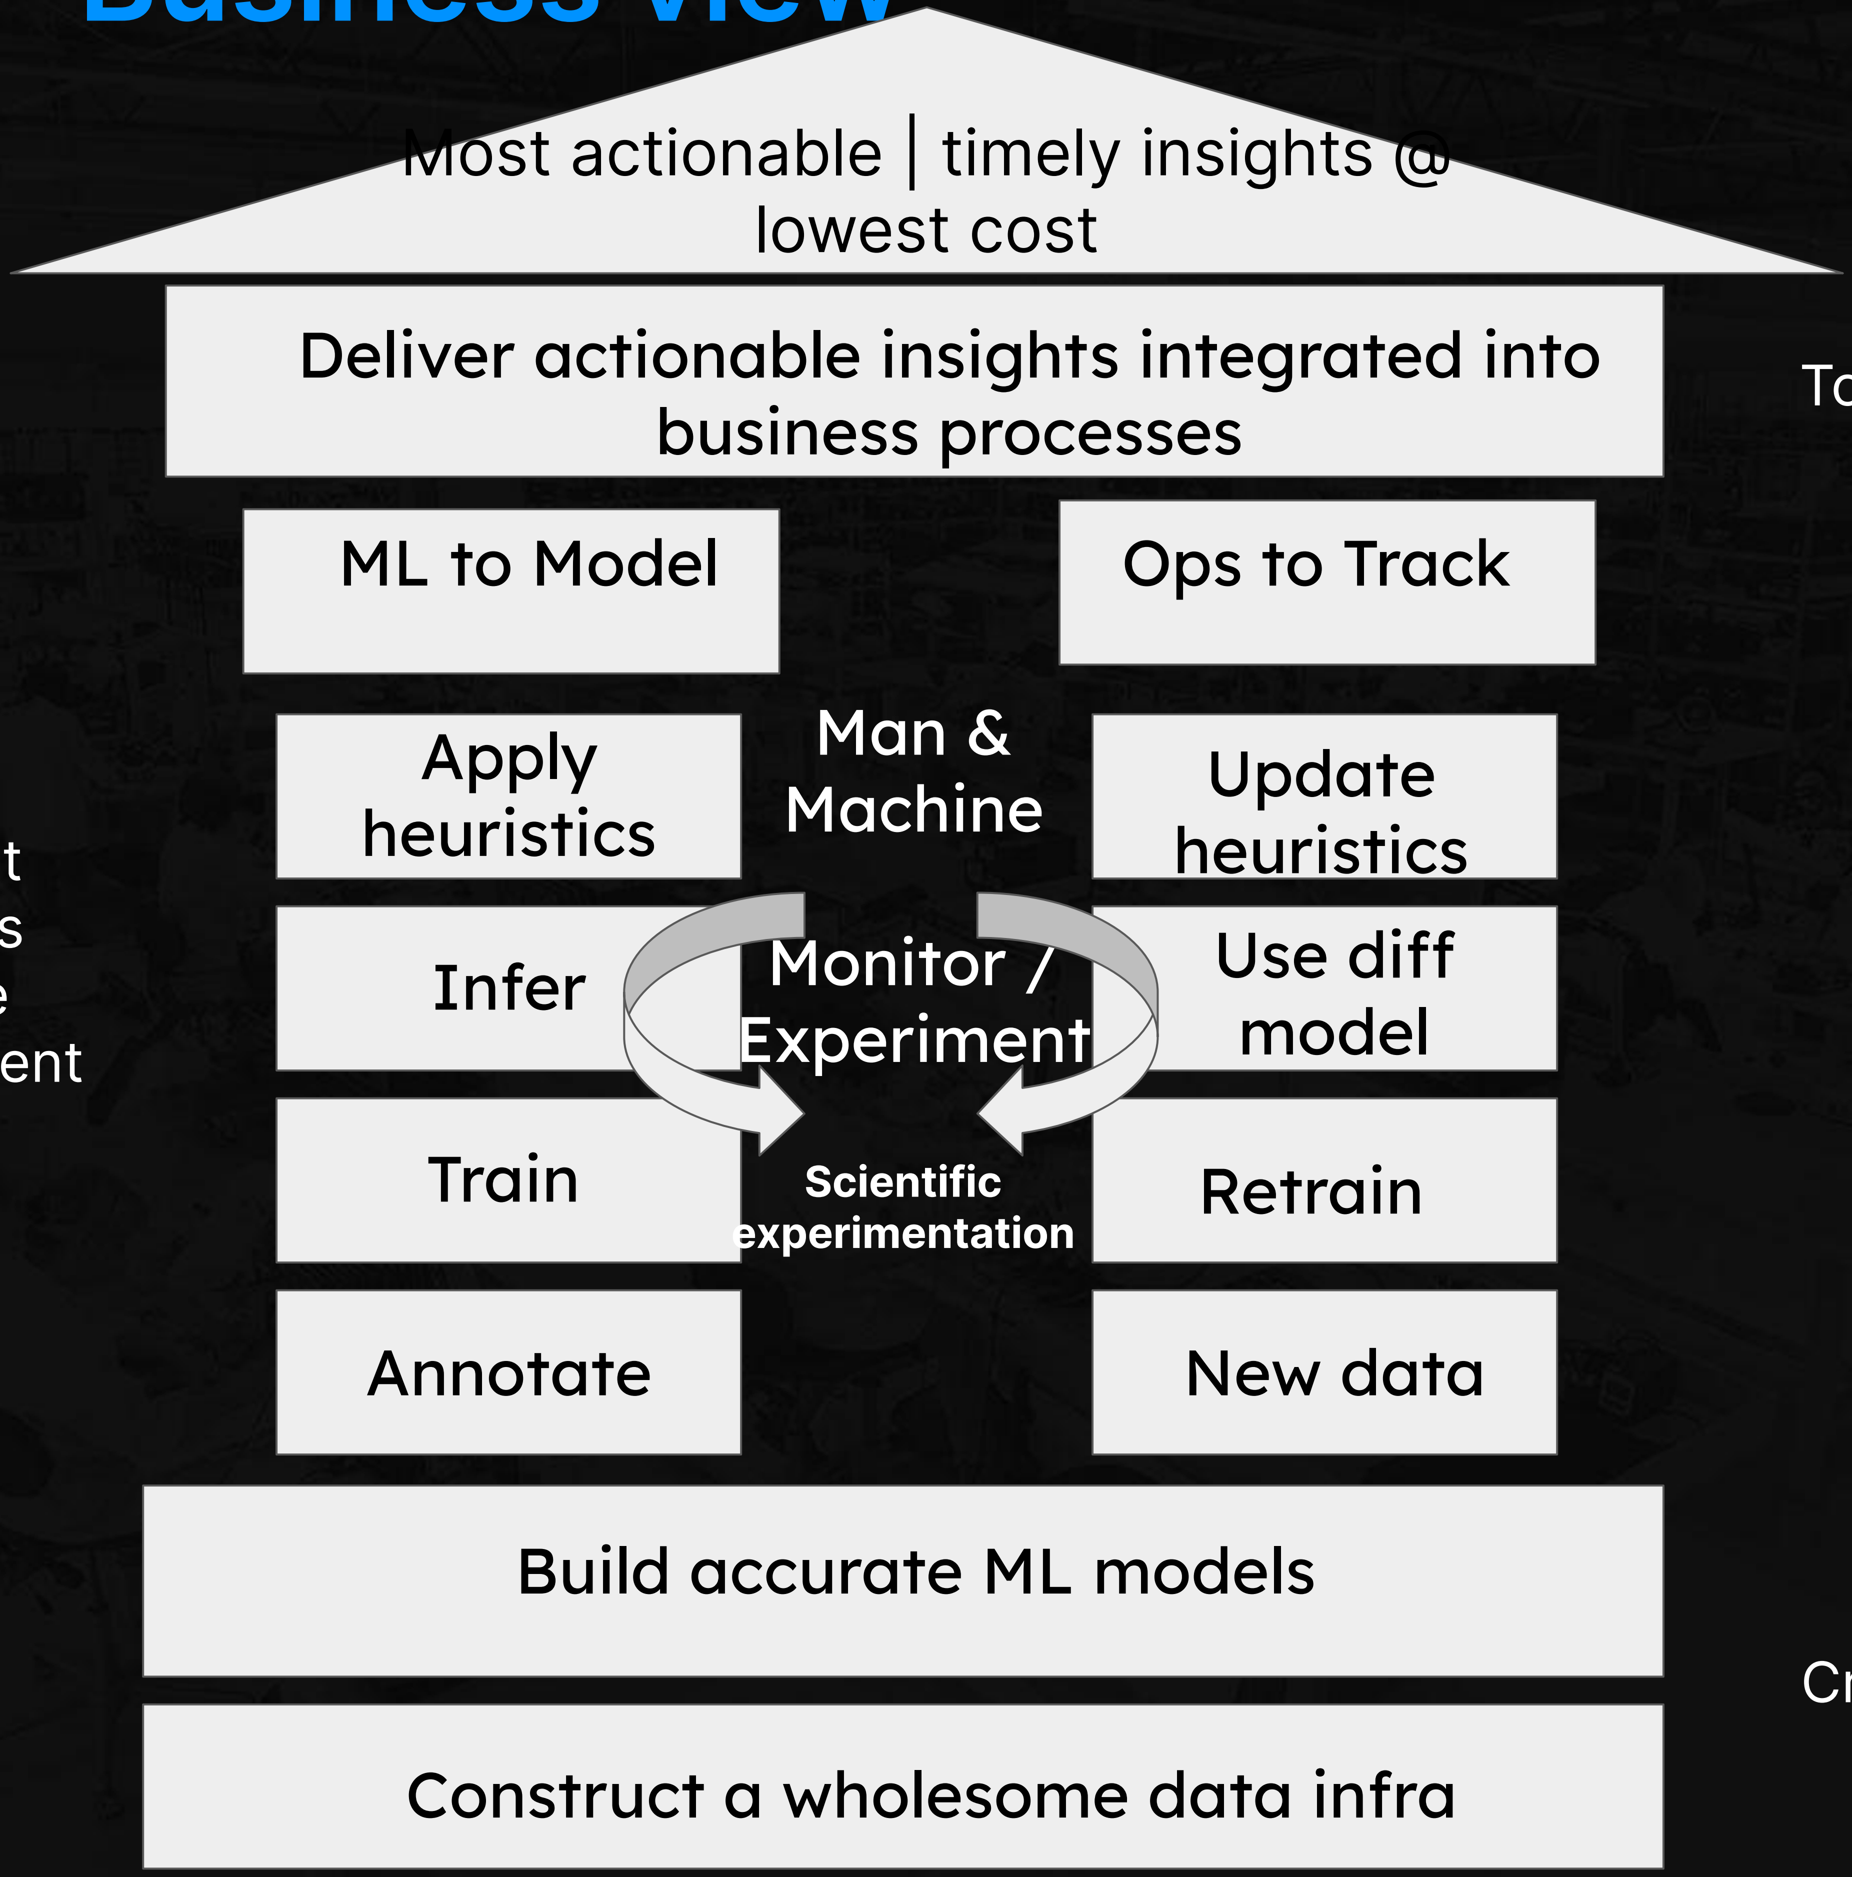

Product
Process
People
Environment

To produce impactful insights

Robust in the face of a dynamic world

Creating quality & timely data (data operations stability)

DRISHTI

MLOps House - Business view

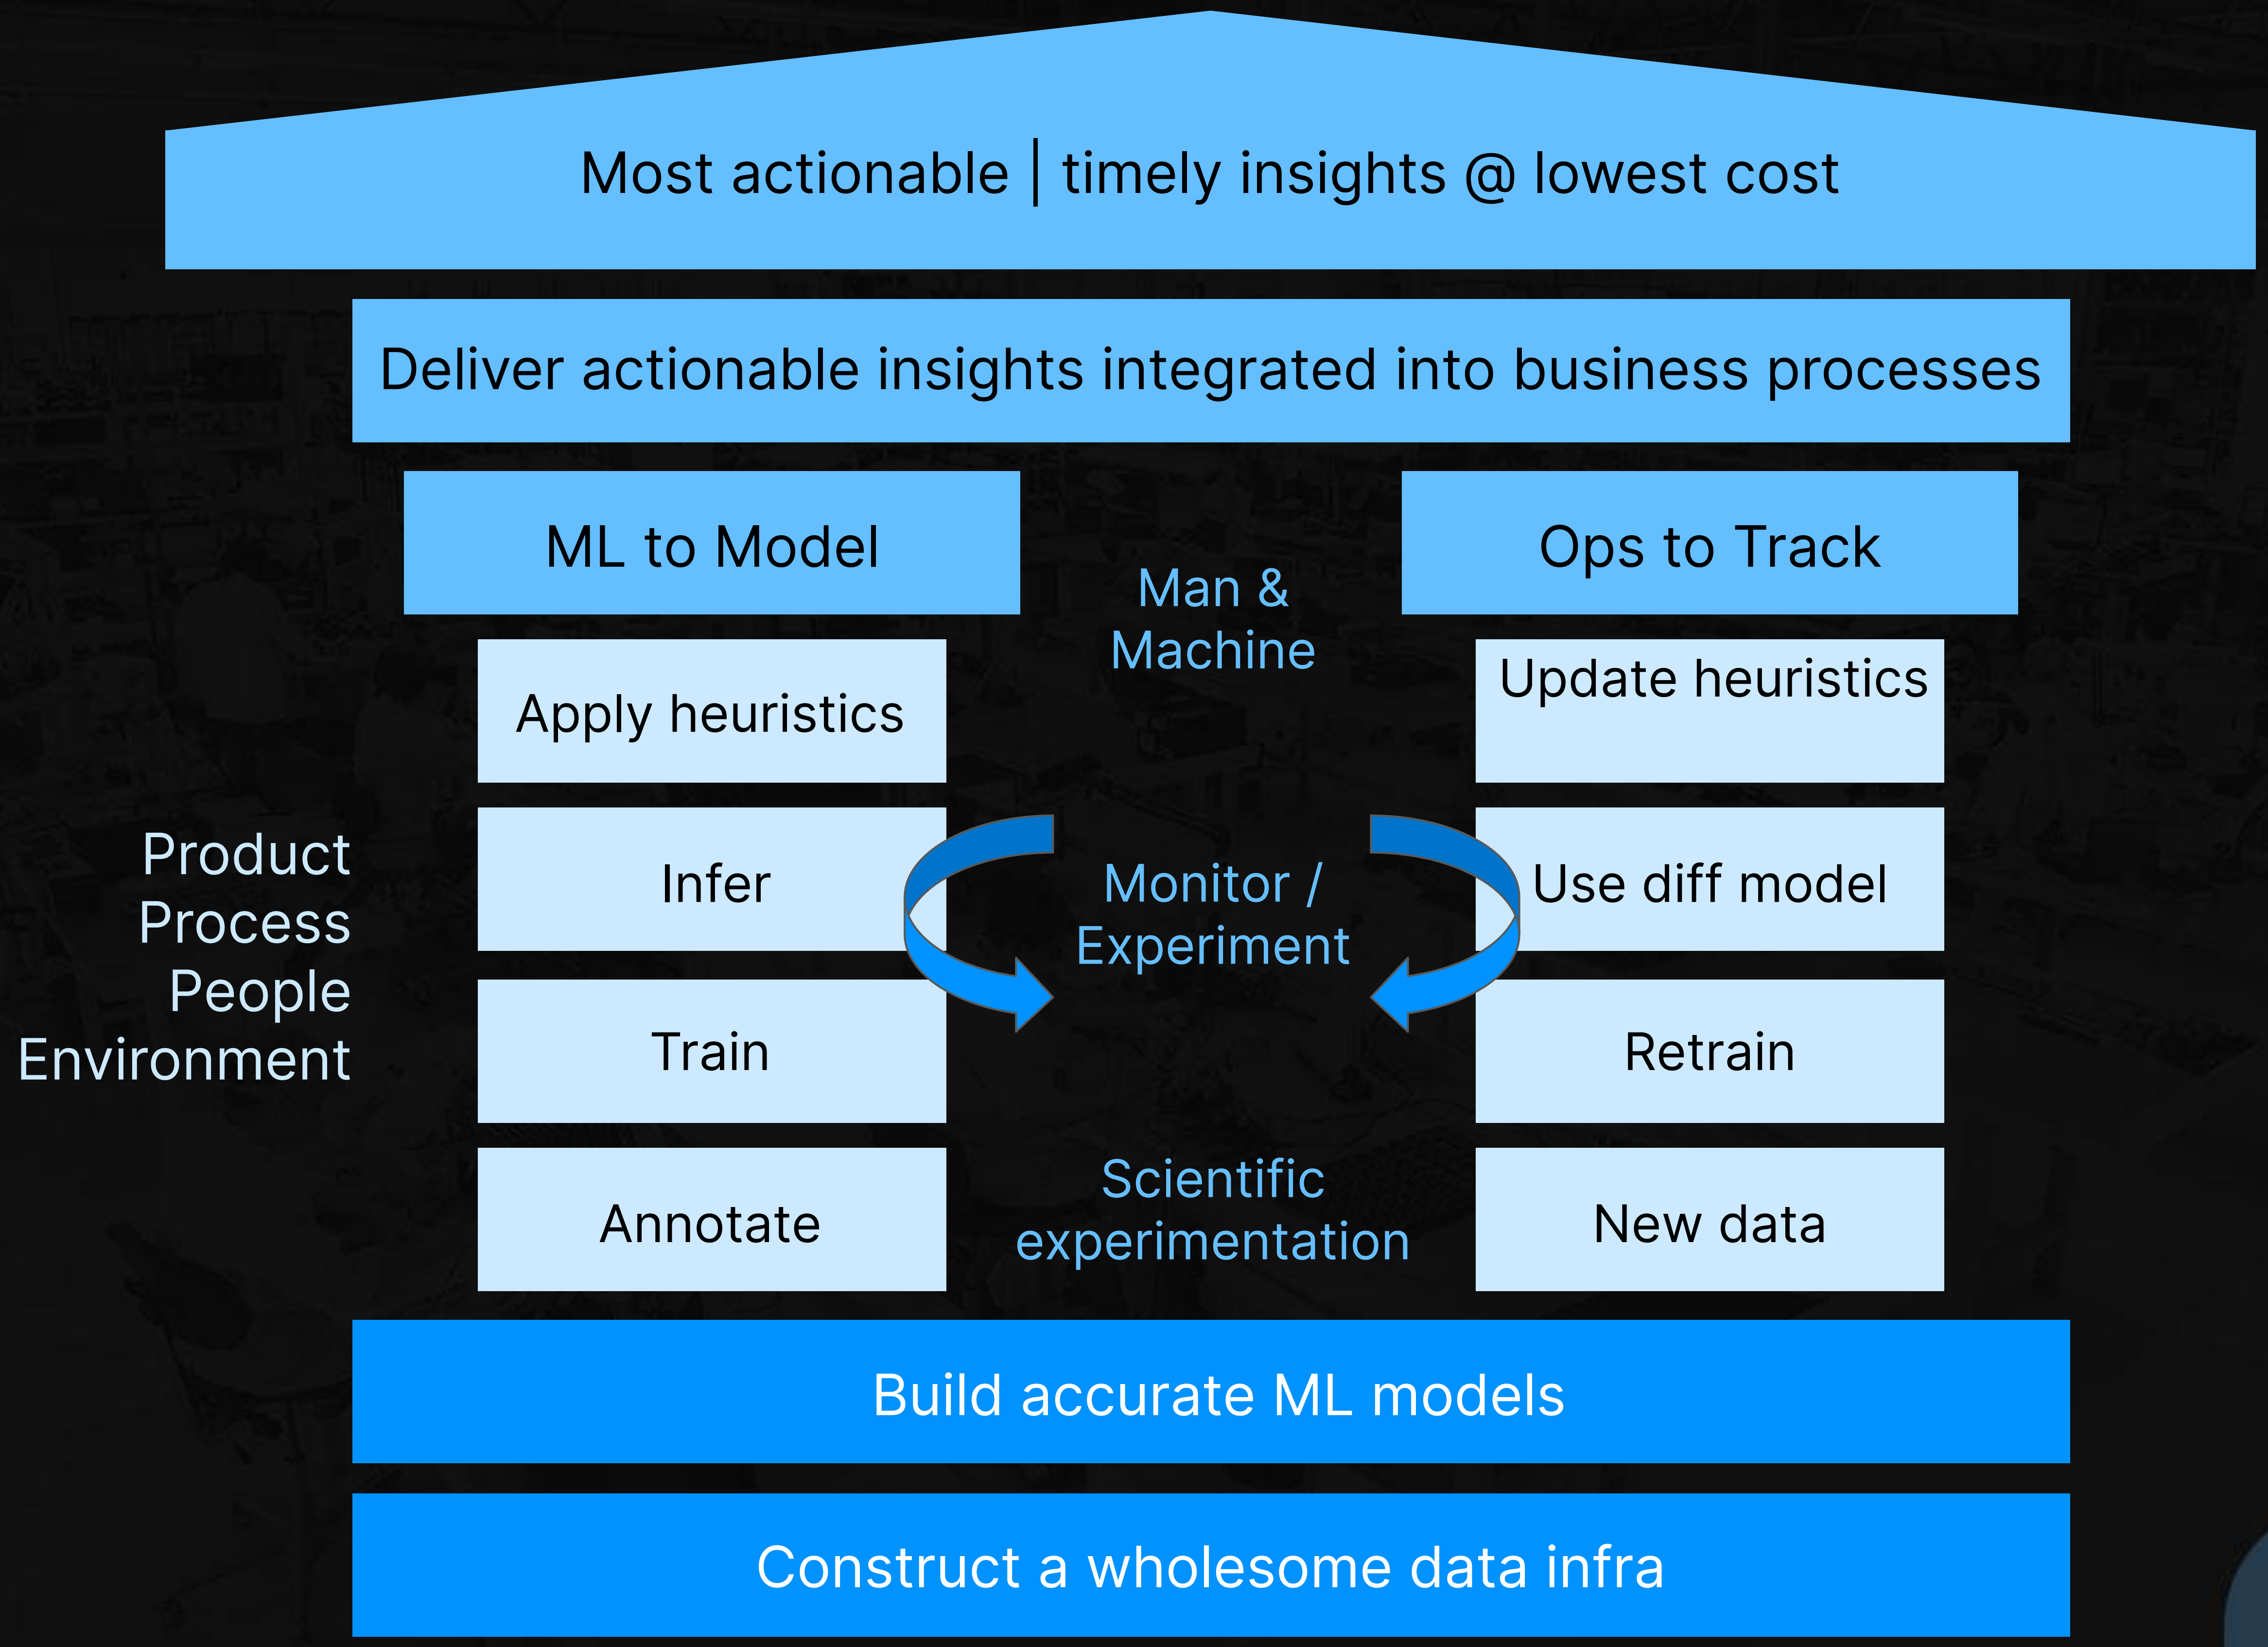

To produce impactful insights

Robust in the face of a dynamic world

Creating quality & timely data (data operations stability)

DRISHTI

MLOps House - Tech/vendor view

Trust and bias?

Goal: Actionable, timely insights @ lowest cost

Deliver actionable insights integrated into business processes

To produce impactful insights

Model the real world

Track the real world

Apply heuristics

Man & Machine

Update heuristics

Product
Process
People
Environment

Infer

Monitor / Experiment

Use diff model

Robust in the face of a dynamic world

Train

Scientific experimentation

Retrain

Annotate

New data

ML models

Creation: PyTorch | TensorFlow | C
Management: A | B

Creating quality & timely data
(data operations stability)

Data

A | B | C

DRISHTI

MLOps House - Tech/vendor view

Trust and bias?

Goal: Actionable, timely insights @ lowest cost

Deliver actionable insights integrated into business processes

To produce impactful insights

Model the real world

Man & Machine

Track the real world

Apply heuristics

Update heuristics

Product
Process
People
Environment

Infer

Monitor / Experiment

Use diff model

Robust in the face of a dynamic world

Train

Retrain

Annotate

Scientific experimentation

New data

ML Models : Creation: PyTorch | TensorFlow | C
Management: A | B

Creating quality & timely data (data operations stability)

Data : A | B | C

DRISHTI

Software 2.0: The building blocks

The ML lifecycle from data ingestion to model deployment

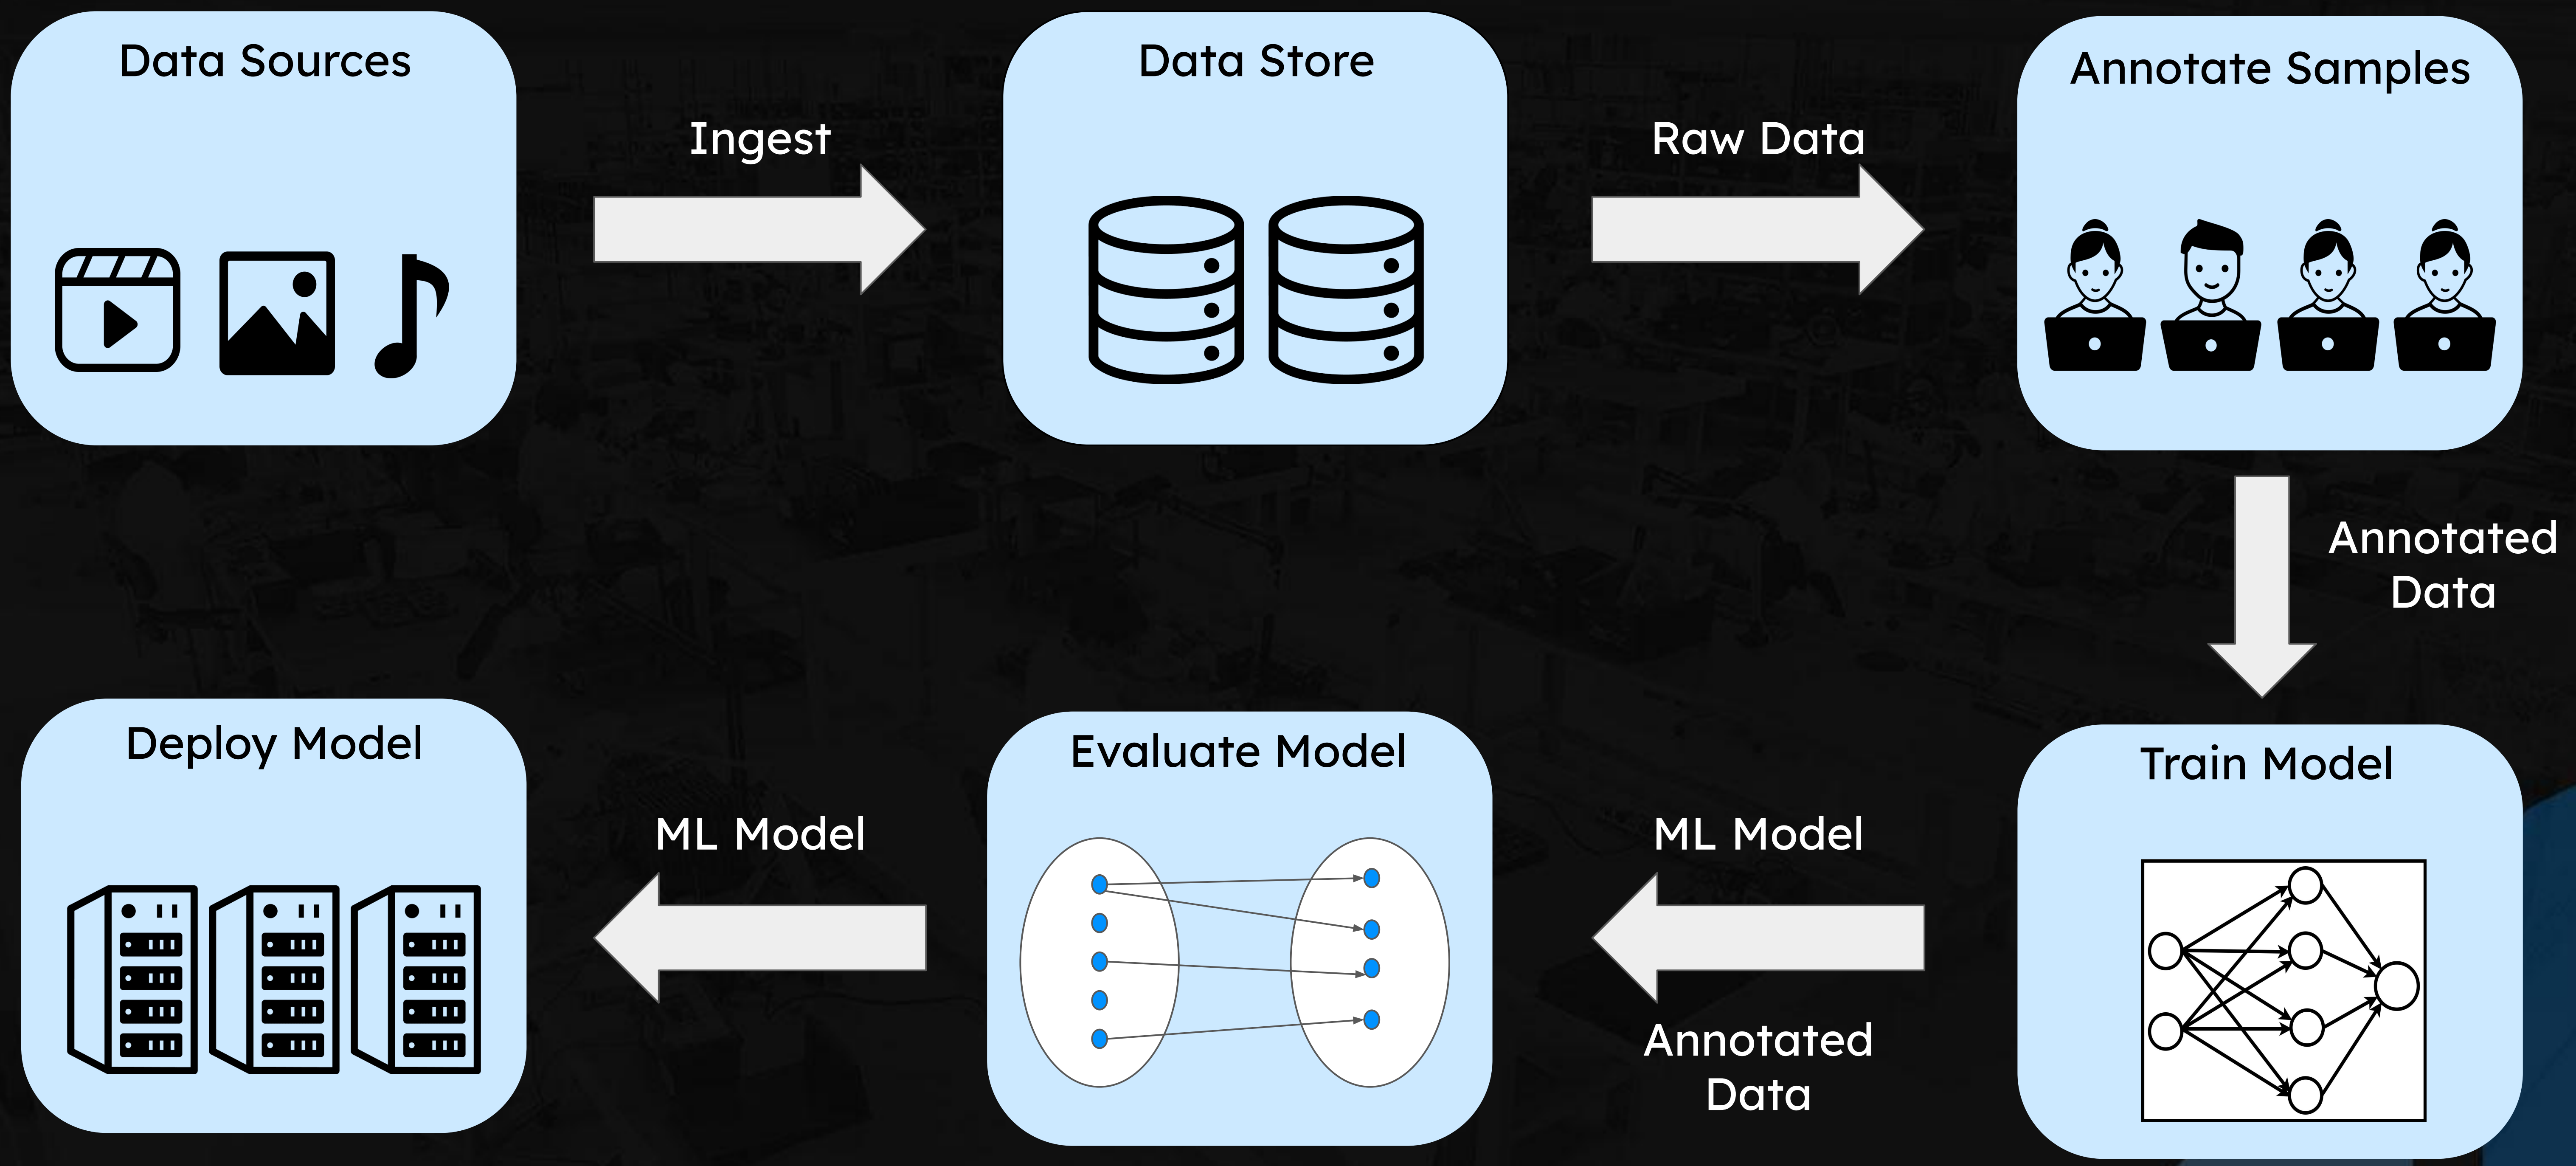

DRISHTI

What is MLOps?

The term MLOps was first coined by Google in their paper on Machine Learning Operations, although it does have roots in software operations.

MLOps is the art and science of bringing machine learning to production

Hybrid MLOps capabilities are defined as those that have some interaction with the cloud while also having some interaction with local computing resources.

DataOps vs MLOps?

MLOps House - Business view

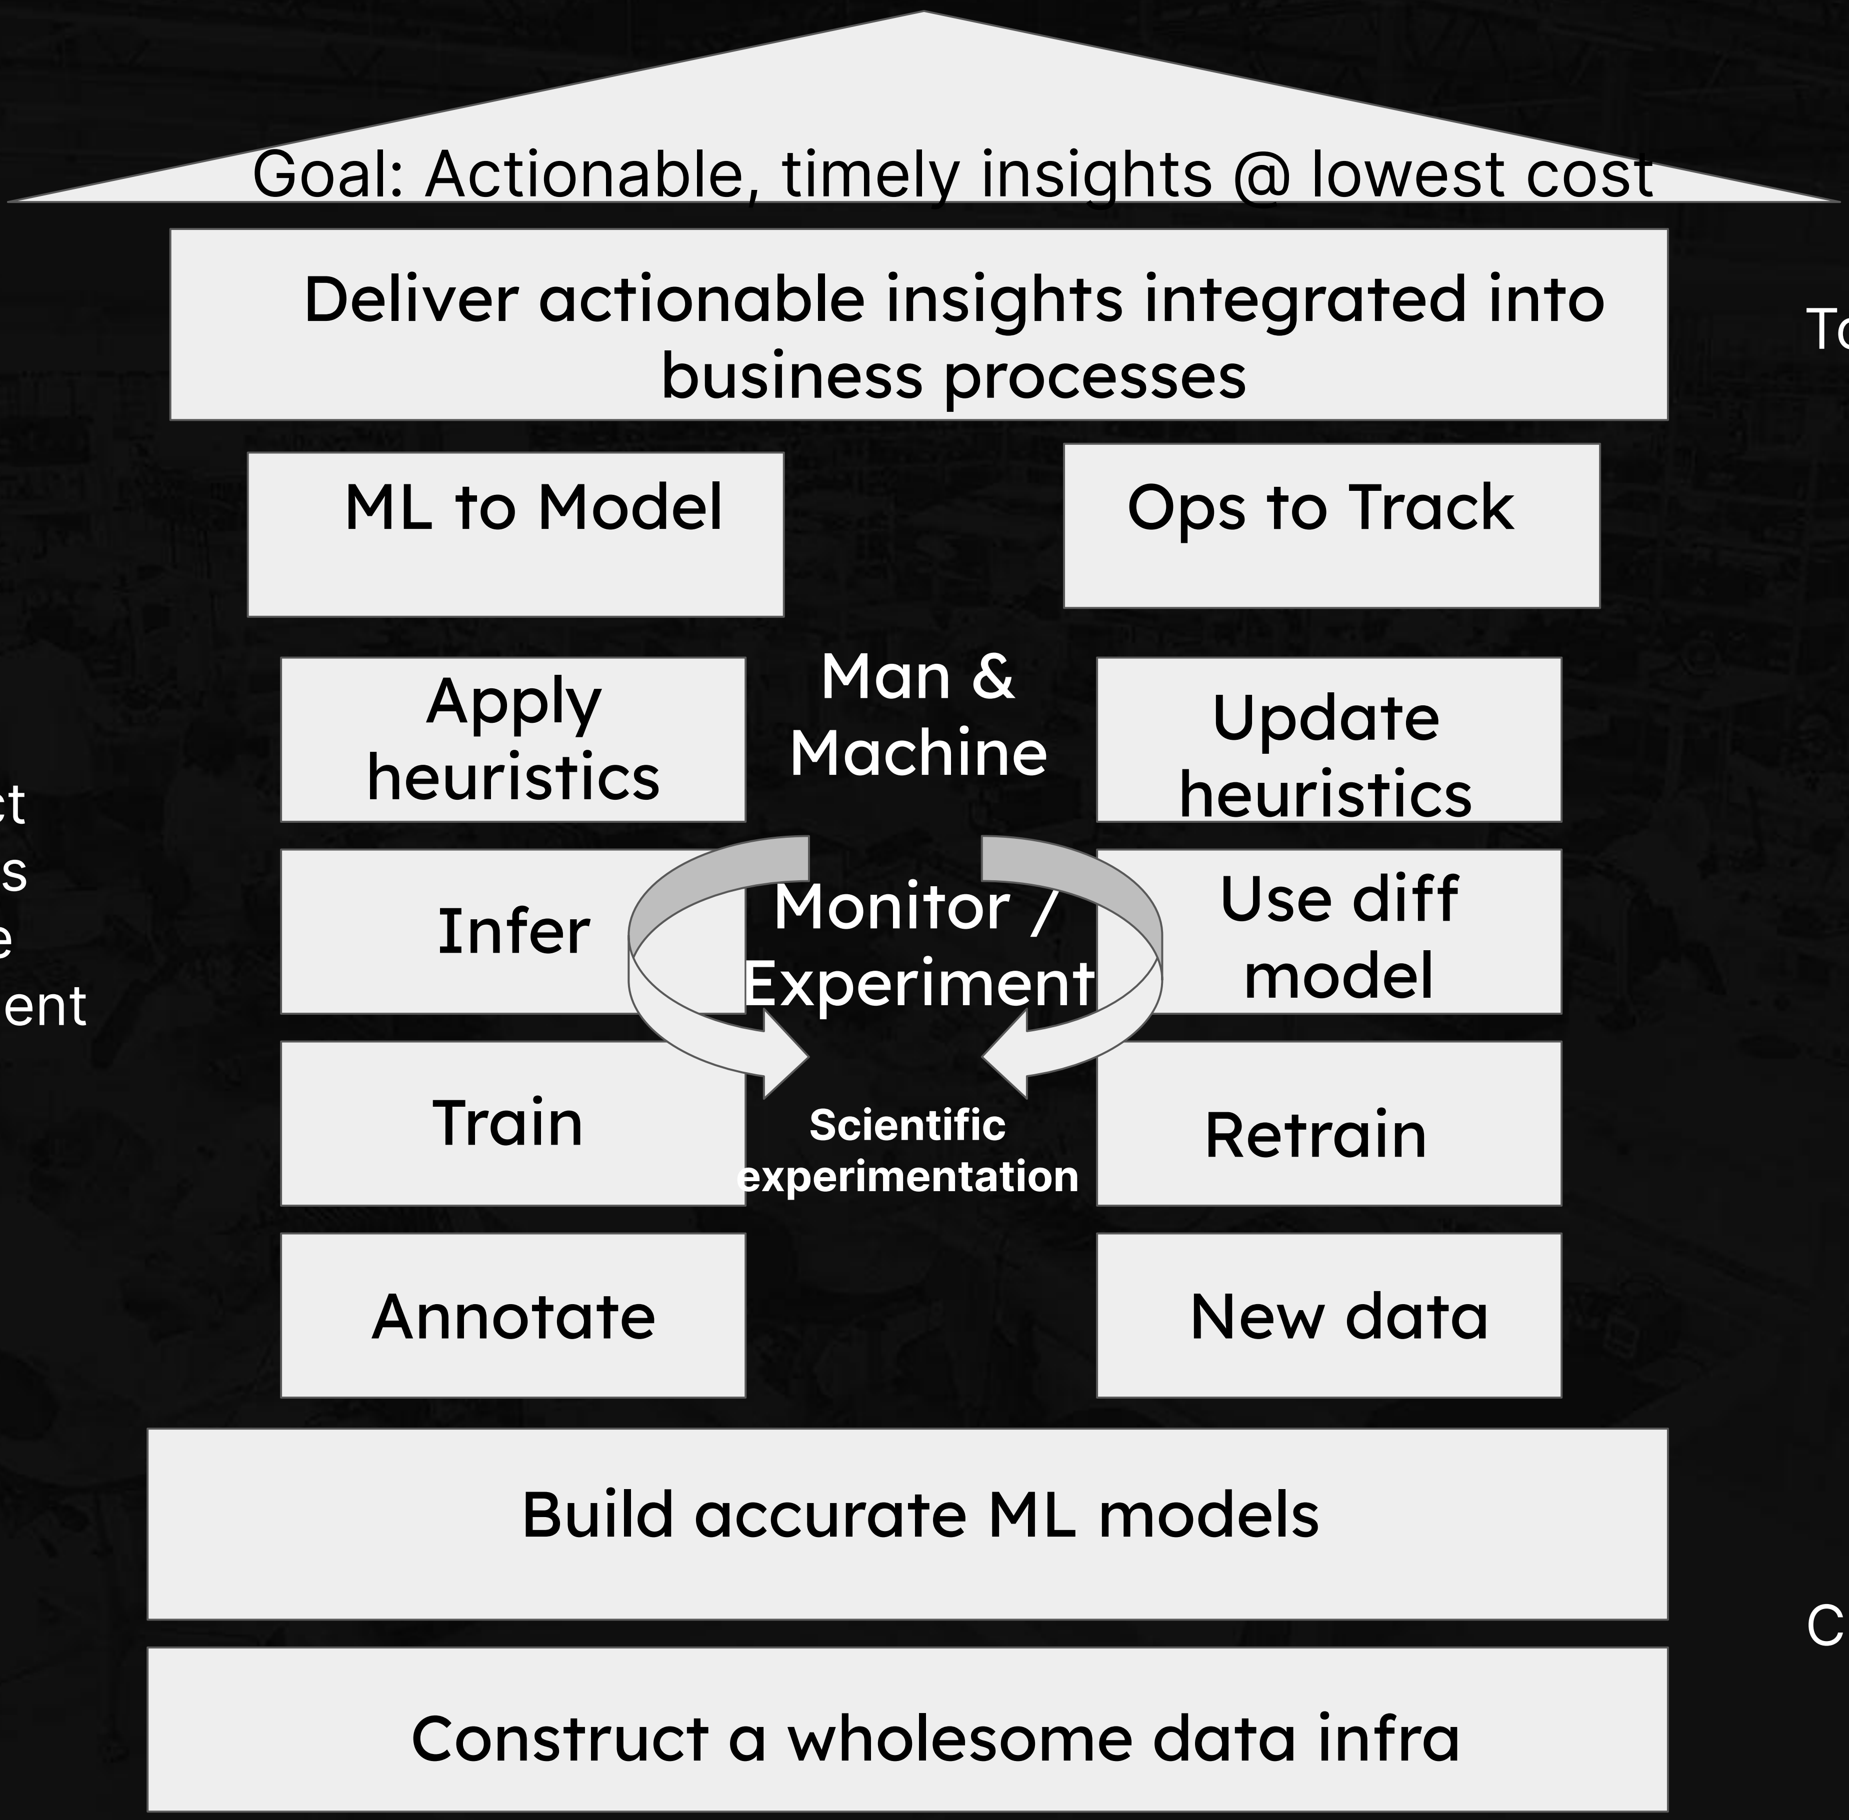

Product
Process
People
Environment

To produce impactful insights

Robust in the face of a dynamic world

Creating quality & timely data
(data operations stability)

DRISHTI

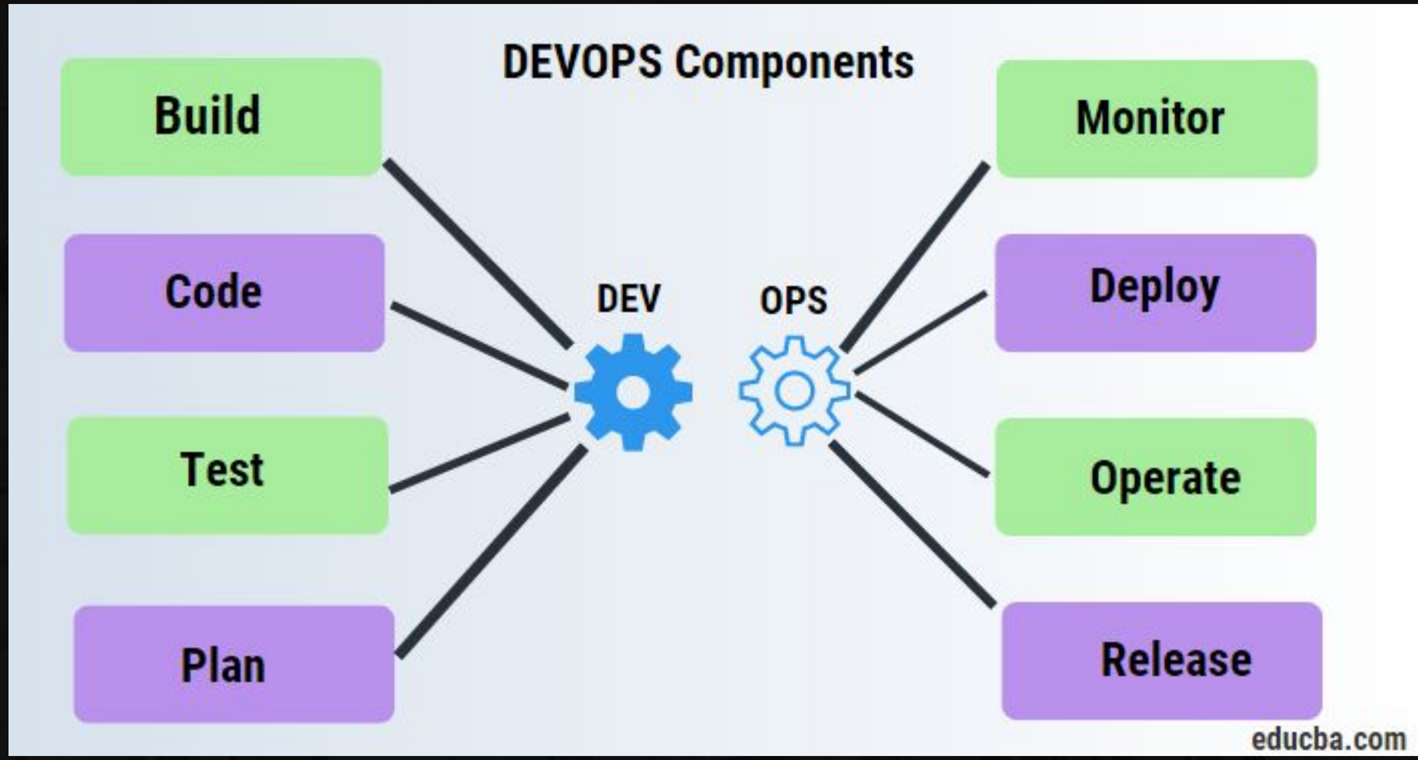
DEVOPS Components
Build
Code
Test
Plan
DEV
OPS
Monitor
Deploy
Operate
Release
educba.com

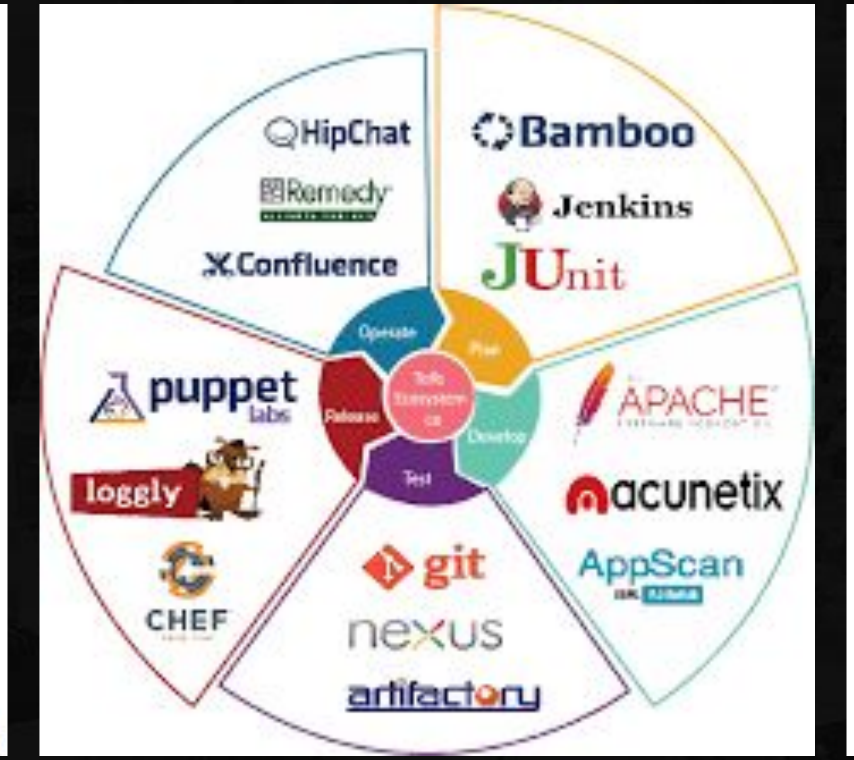
HipChat
Remedy
Confluence
Bamboo
Jenkins
JUnit
puppet labs
loggly
CHEF
APACHE
acunetix
AppScan
git
nexus
artifactory


Development
IT Operations
DevOps
Quality Assurance

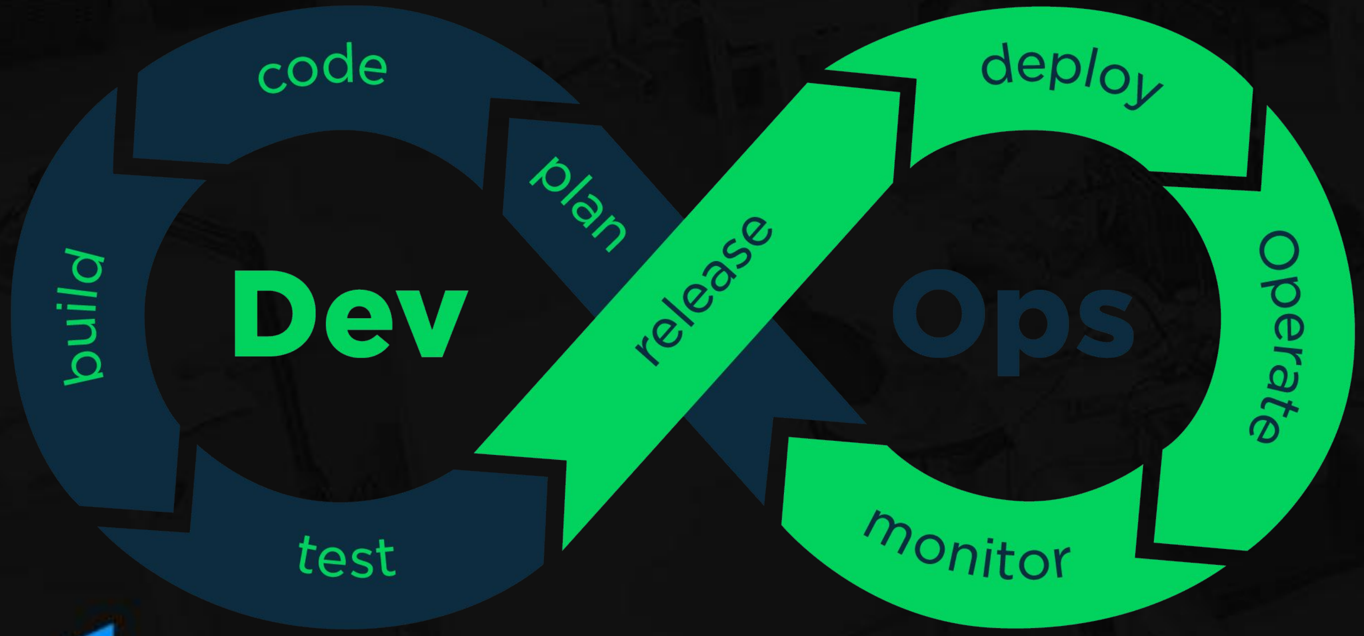
code
plan
build
Dev
test
release
deploy
Ops
Operate
monitor

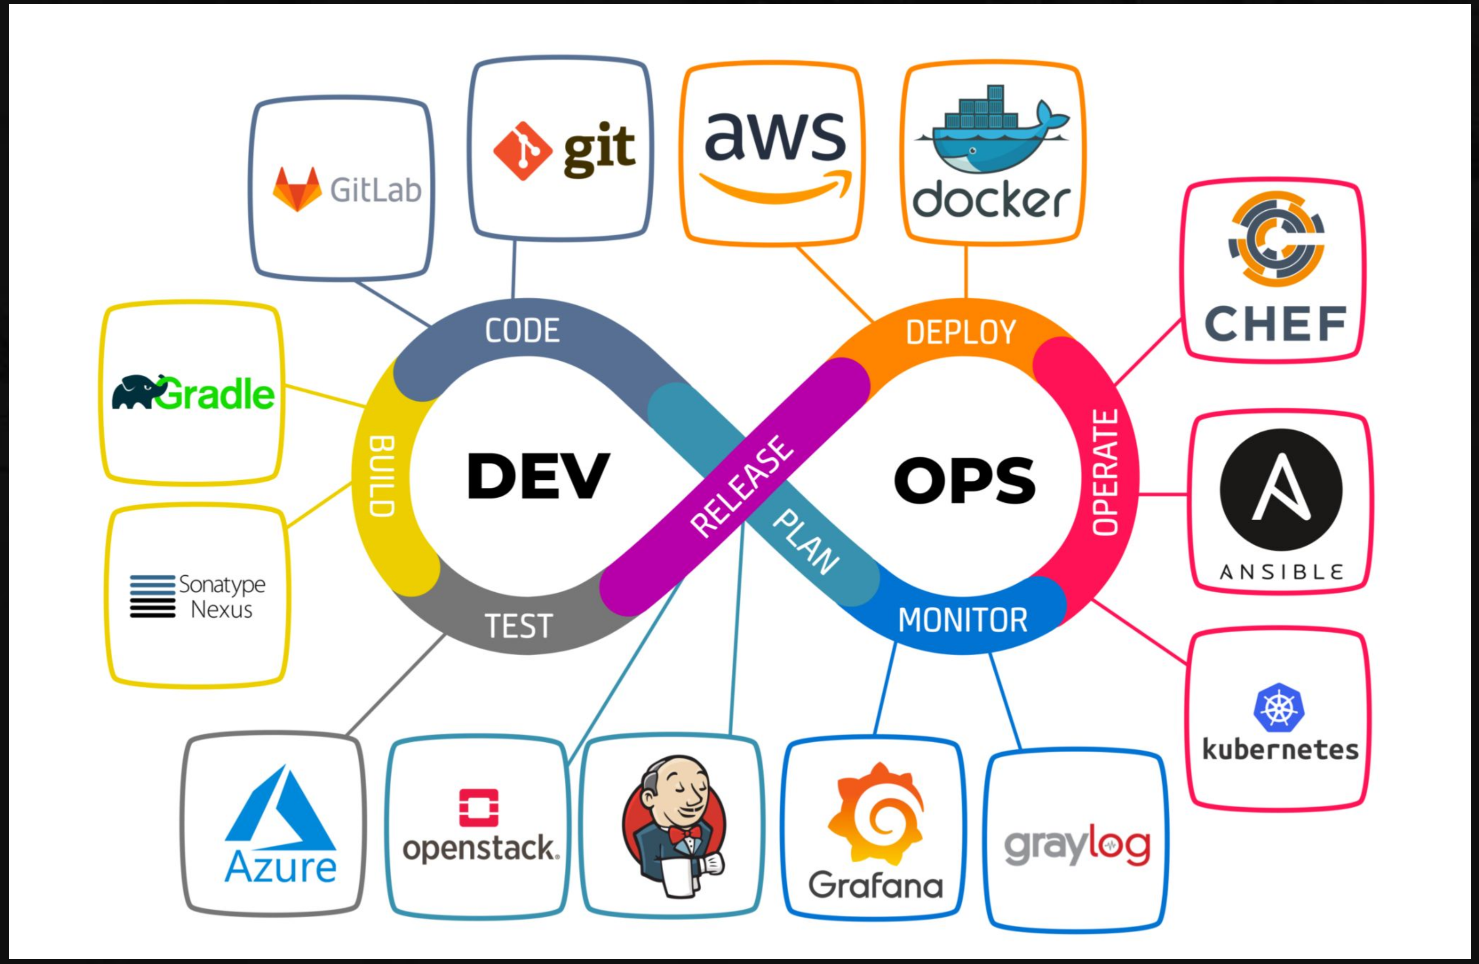
GitLab
git
aws
docker
CHEF
Gradle
Sonatype Nexus
CODE
BUILD
TEST
DEV
RELEASE
PLAN
DEPLOY
OPS
OPERATE
MONITOR
ANSIBLE
kubernetes
Azure
openstack.
Grafana
graylog

"Tools of the Trade"

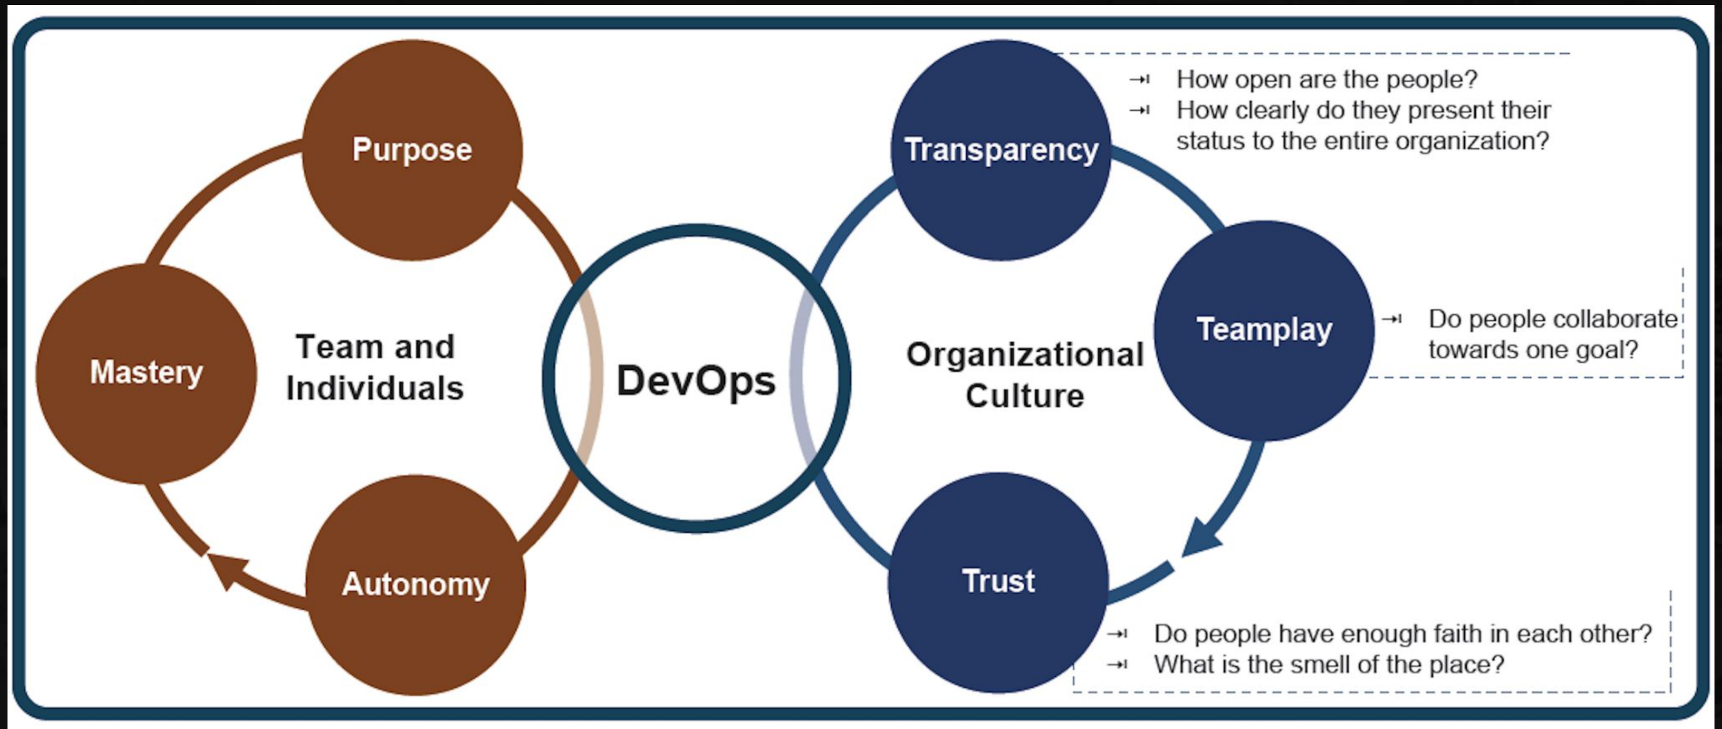

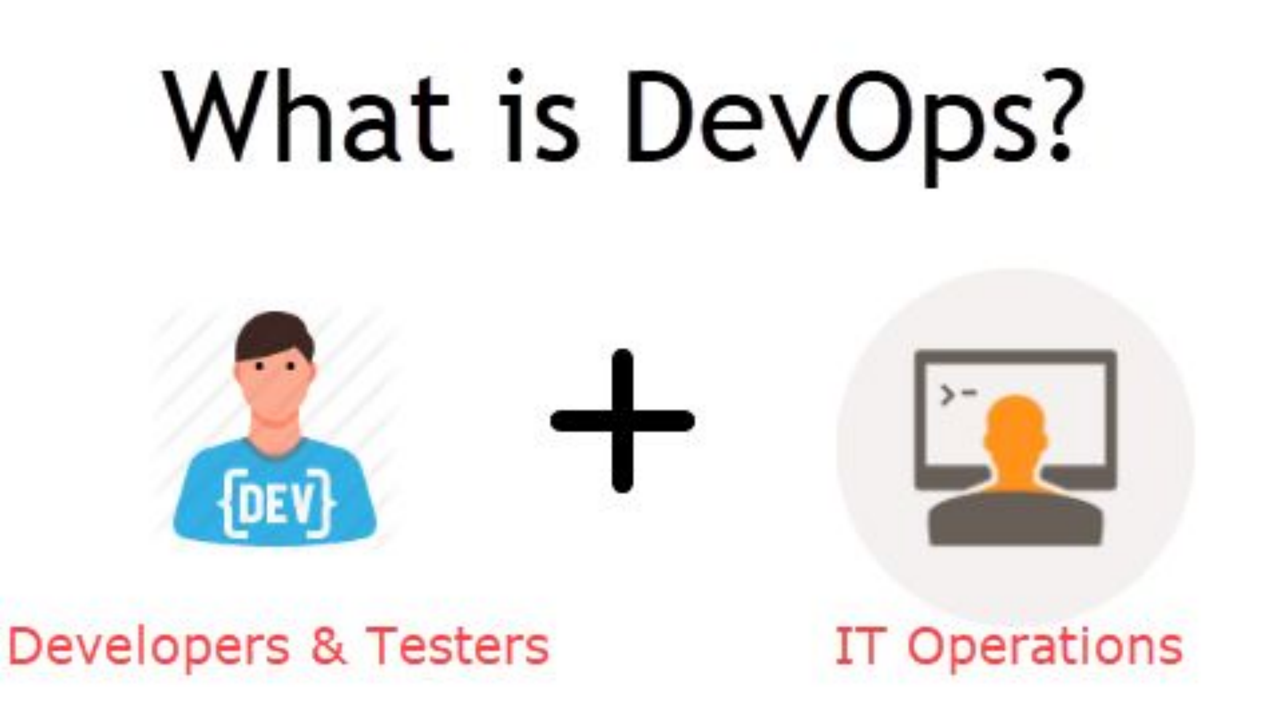

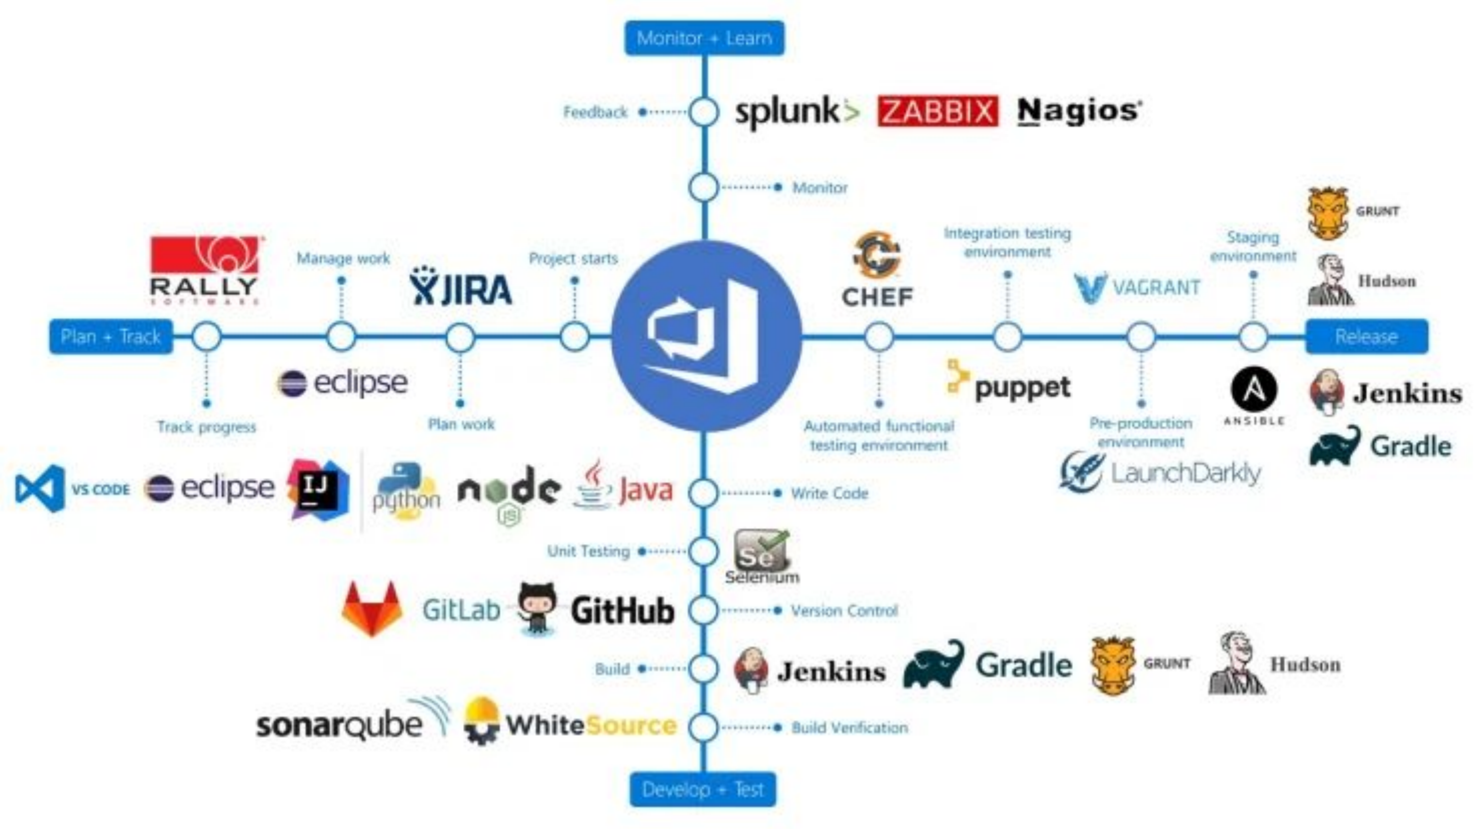

Trust and bias?

The term MLOps was first coined by Google in their paper on Machine Learning Operations, although it does have roots in software operations.

MLOps is the art and science of bringing machine learning to production

Hybrid MLOps capabilities are defined as those that

A day in the life of an AI company (aka, MLOps)

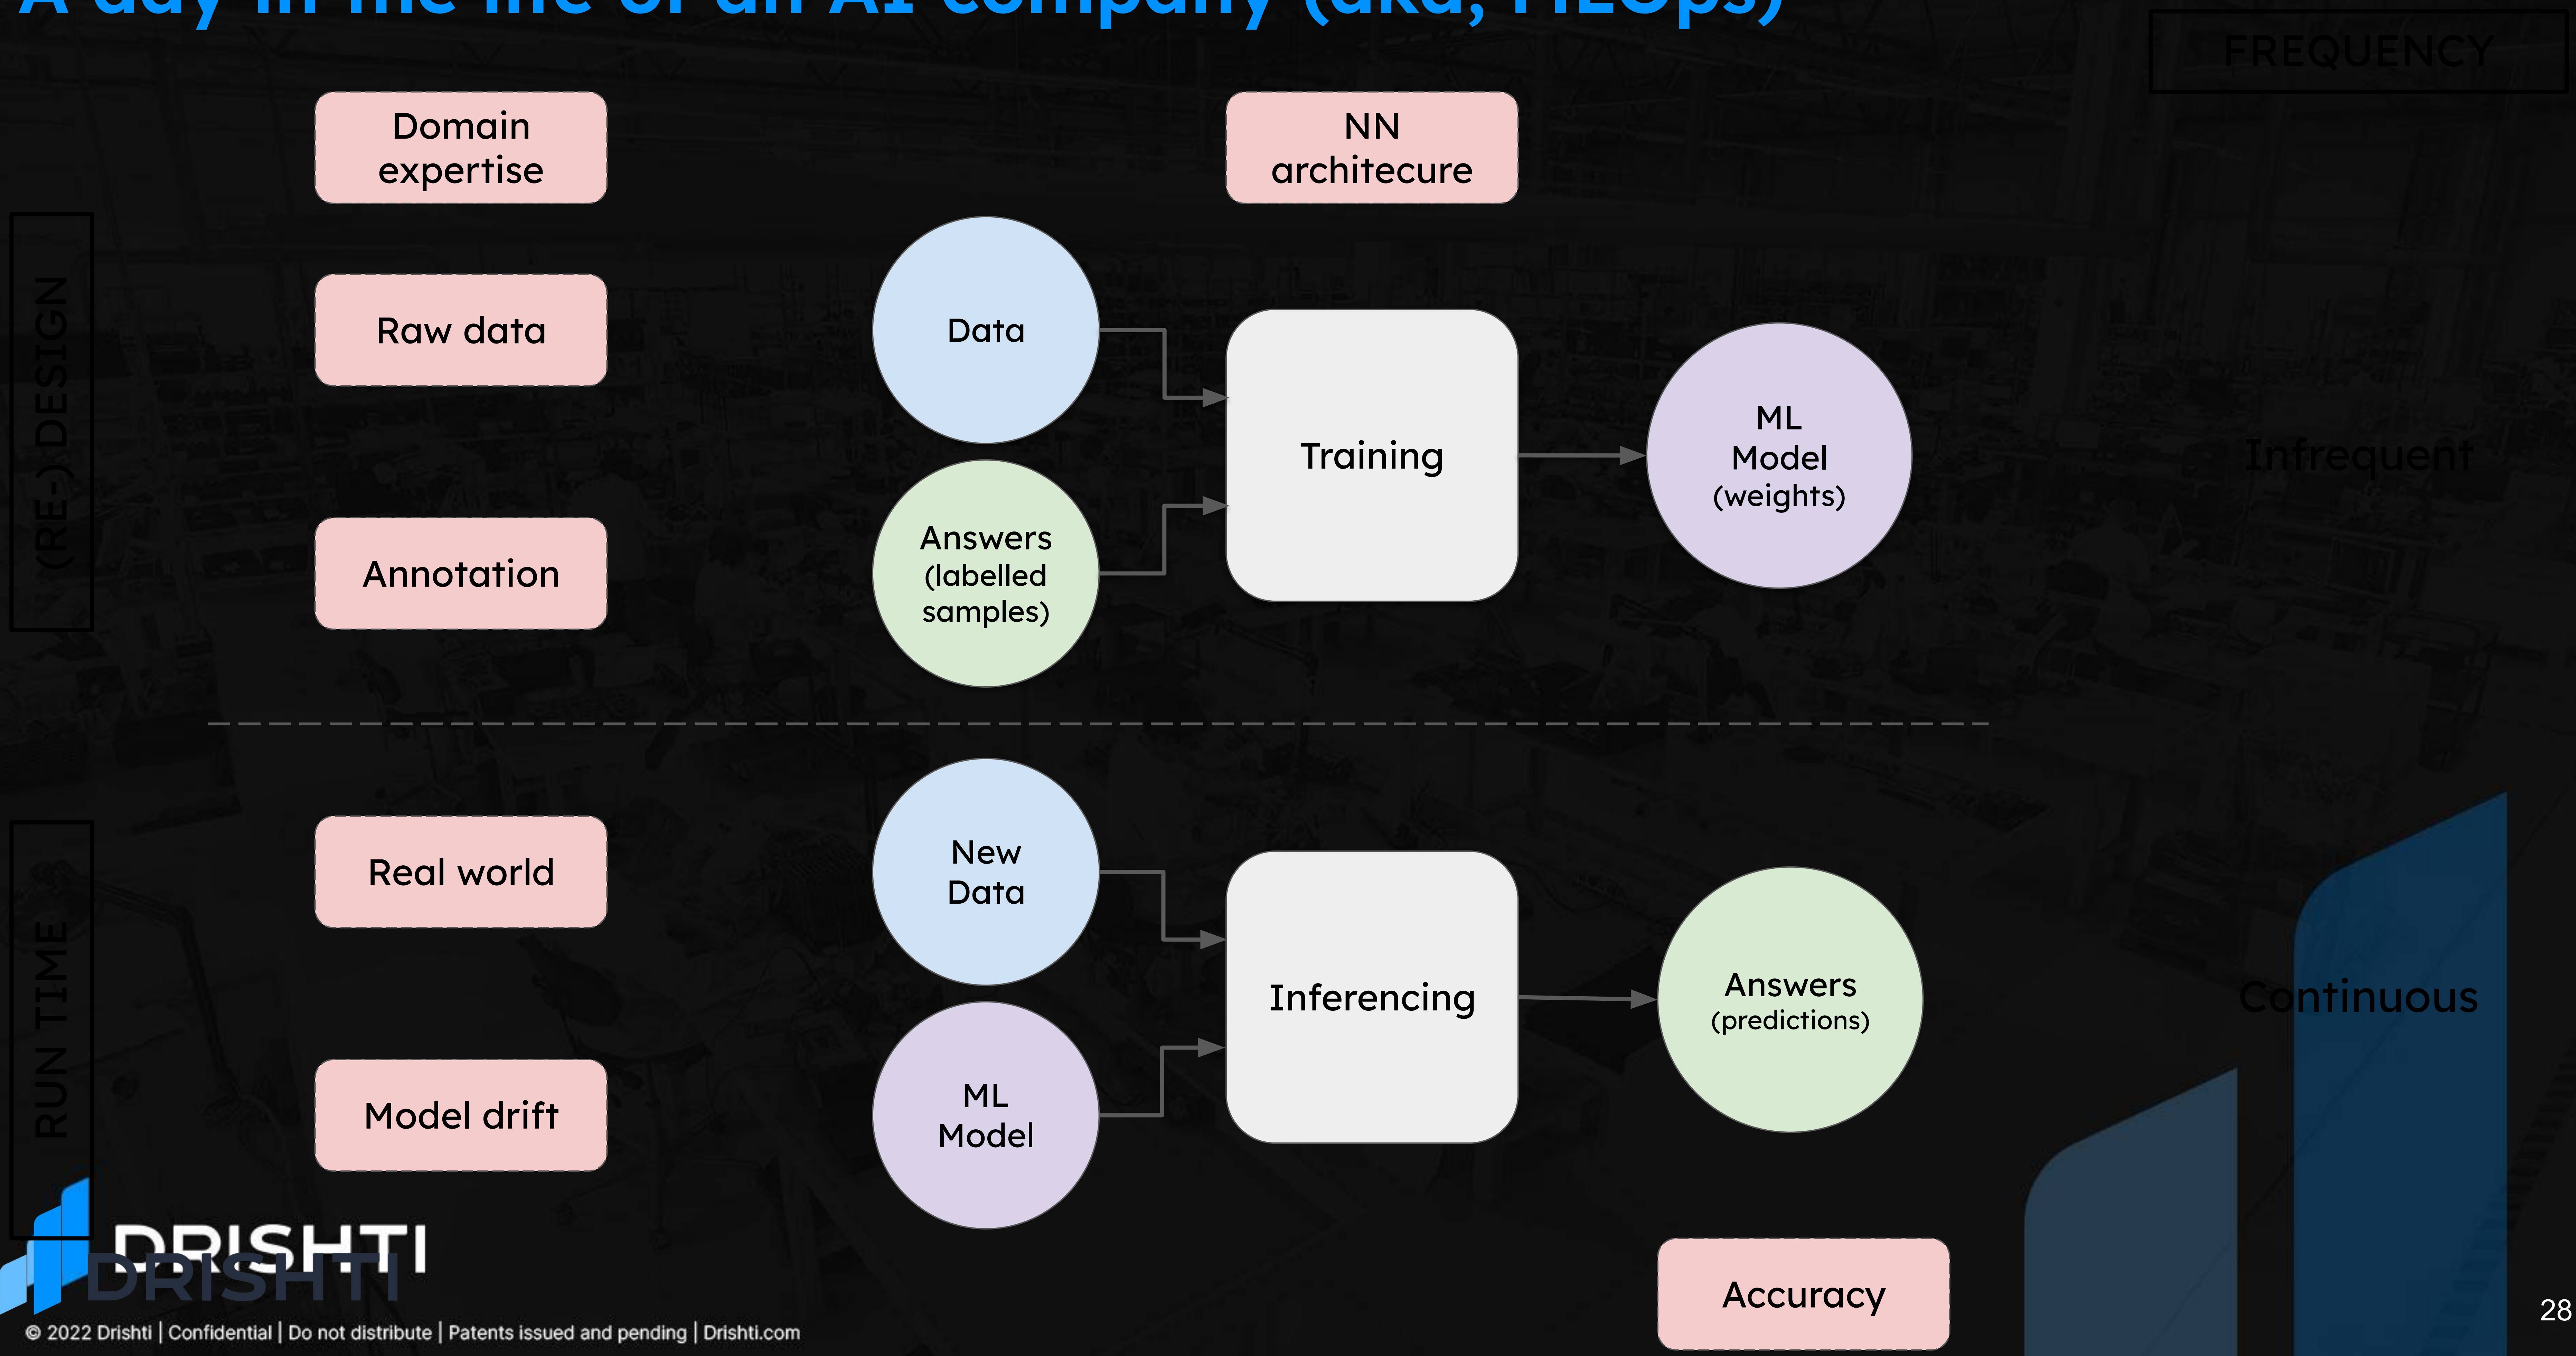

A day in the life of an AI company (aka, MLOps)

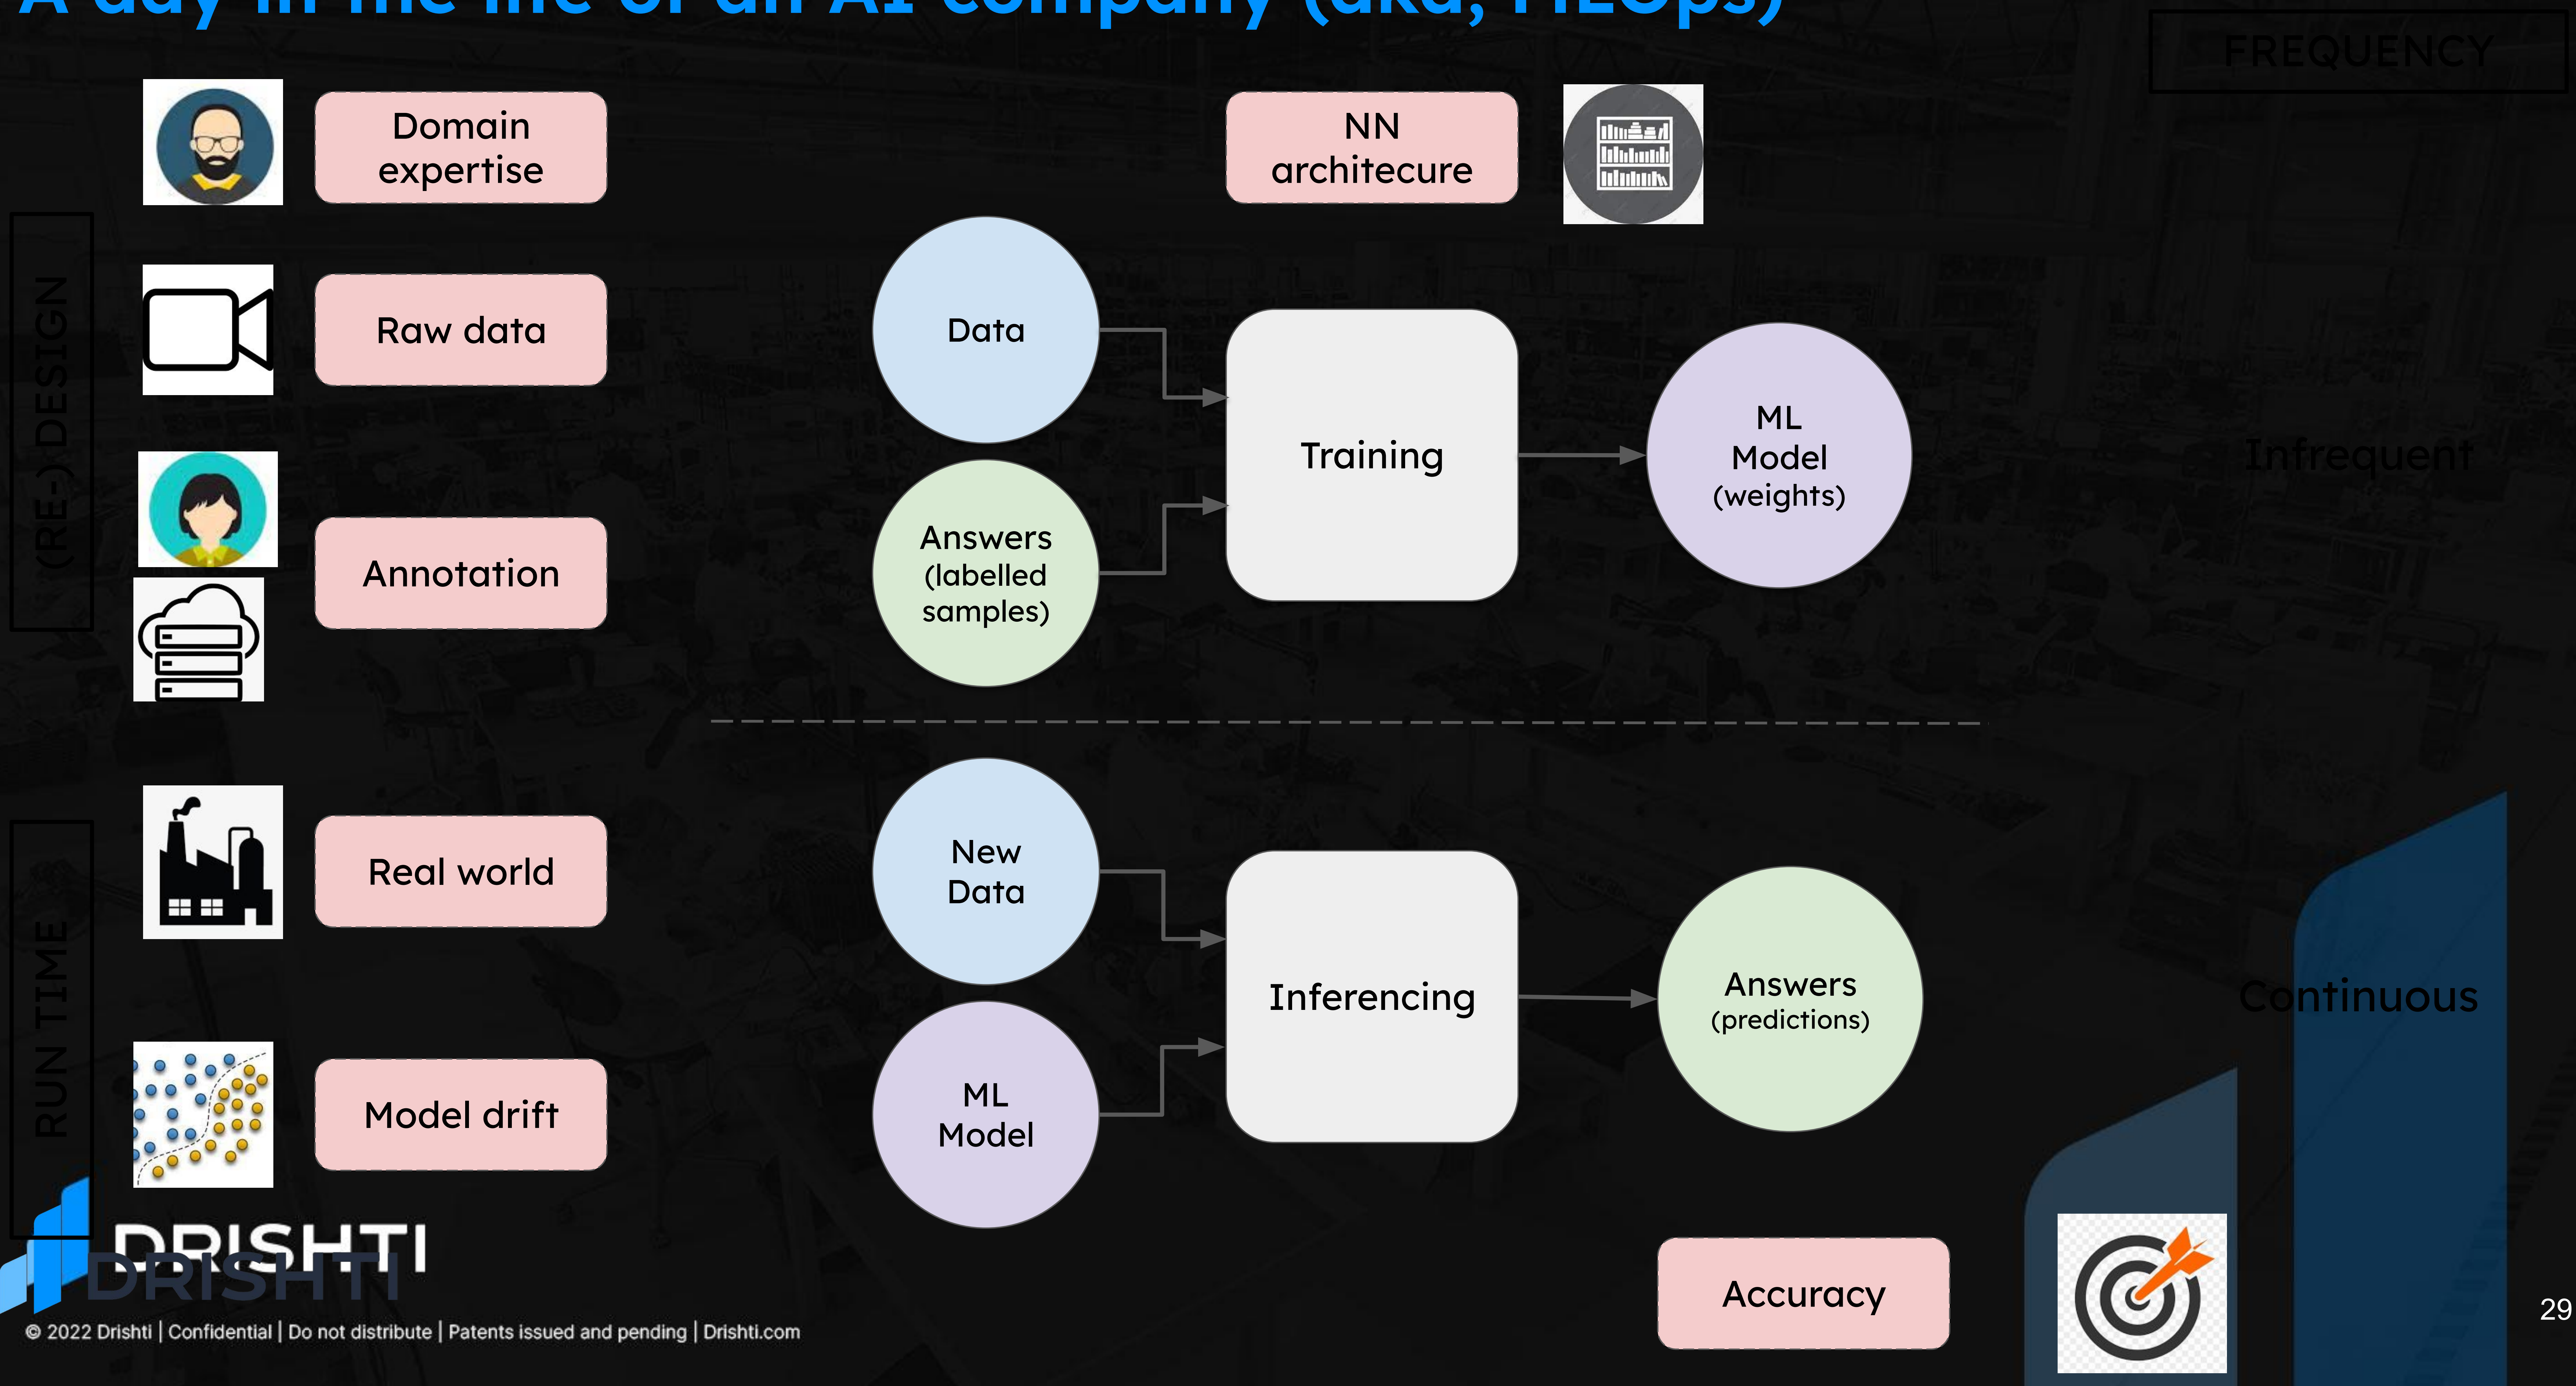

ML Lifecycle

How to train and deploy models

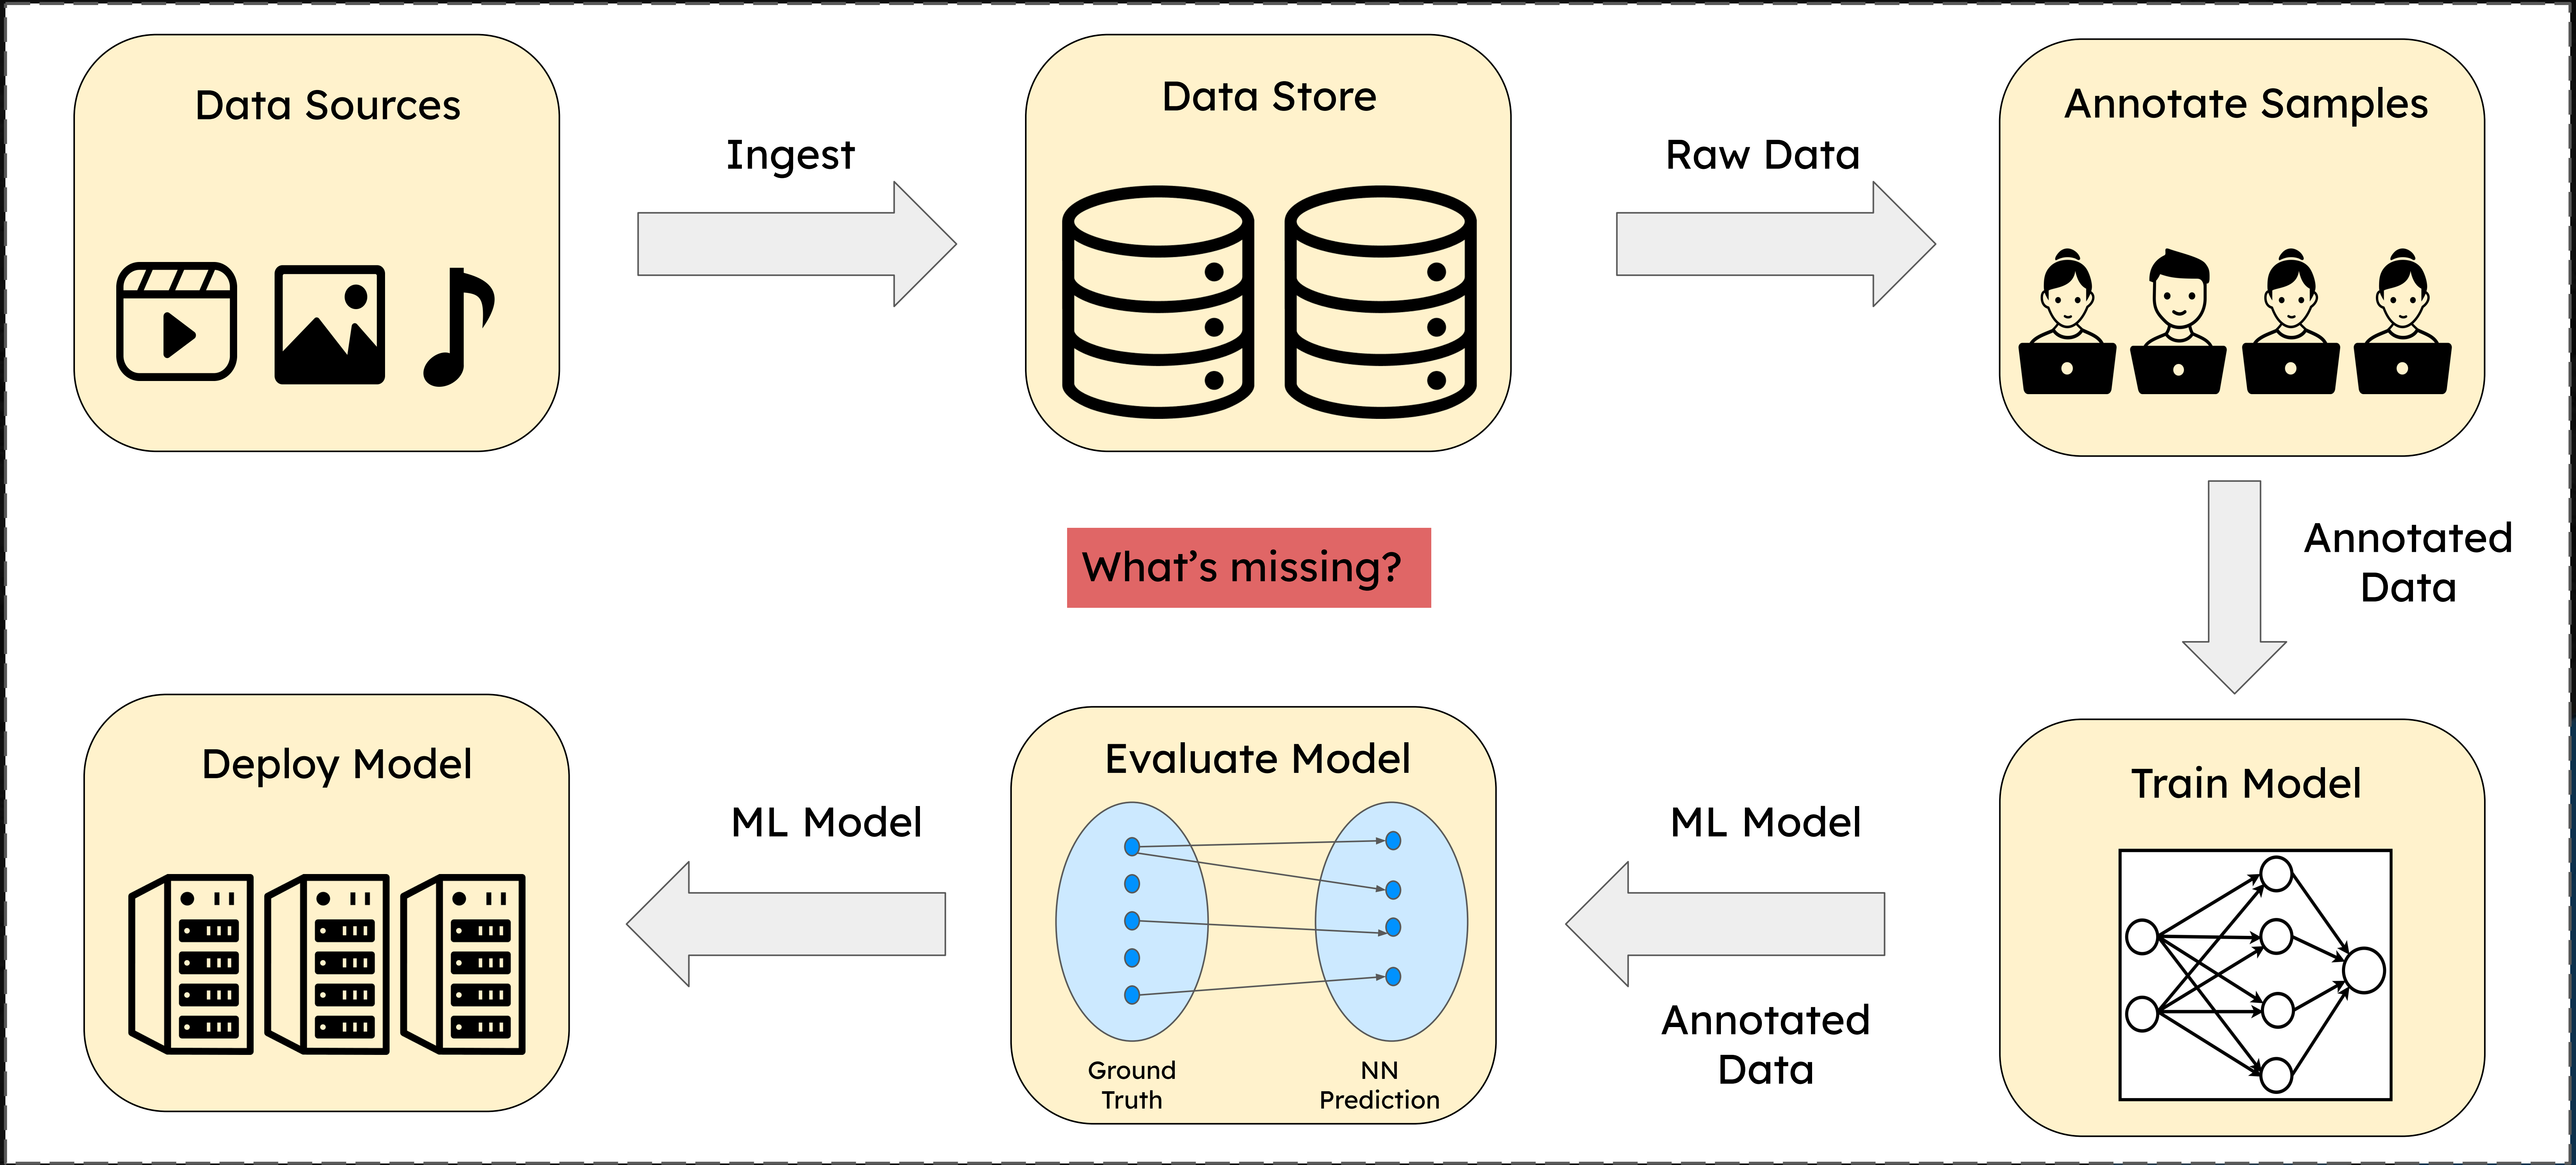

DRISHTI

Quality

CONFIDENTIAL

FREQUENCY

(RE-) DESIGN

Domain expertise

Raw data

Annotation

NN architecure

Code quality
Logic quality

Infrequent

RUN TIME

Real world

Model drift

Data quality
Annotation quality
Neural network accuracy

Continuous

Accuracy

DRISHTI

Team

FREQUENCY

(RE-) DESIGN

Domain expertise

Raw data

Annotation

NN architecture

Code quality
Logic quality

Infrequent

RUN TIME

Real world

Model drift

Data quality
Annotation quality
Neural network accuracy

Continuous

Accuracy

DRISHTI

Workflow

FREQUENCY

(RE-) DESIGN

Domain expertise

Raw data

Annotation

NN architecure

Code quality
Logic quality

Infrequent

RUN TIME

Real world

Model drift

Data quality
Annotation quality
Neural network accuracy

Continuous

Accuracy

DRISHTI

Measurement

FREQUENCY

(RE-) DESIGN

Domain expertise

Raw data

Annotation

NN architecure

Code quality
Logic quality

Infrequent

RUN TIME

Real world

Model drift

Data quality
Annotation quality
Neural network accuracy

Continuous

Accuracy

Different about AI

- Generalization

DRISHTI

What is MLOps? (single slide)

- Software 1.0 + Software 2.0
- Draw parallel to devops

The Foundation of an AI strategy (Graphic for all)

- Impact MLOps had on AI advancement
- MLOps capabilities (covered)
- Operationalizing AI (covered)
- Day in the life (covered)

DRISHTI

Building Proficient MLOps teams (roll into team slides)

- Skills that are necessary
- Lessons learned from DevOps teams (overlap)

ML Team (NN, Annotation, MLQuality)

Operationalizing AI (based off of Drishti - 2 slides) (covered)

- What a practical AI system looks like
- How a company’s AI operations need to functions

DRISHTI

Closing slide - thank you

DRISHTI

APPENDIX

A day in the life of an AI company

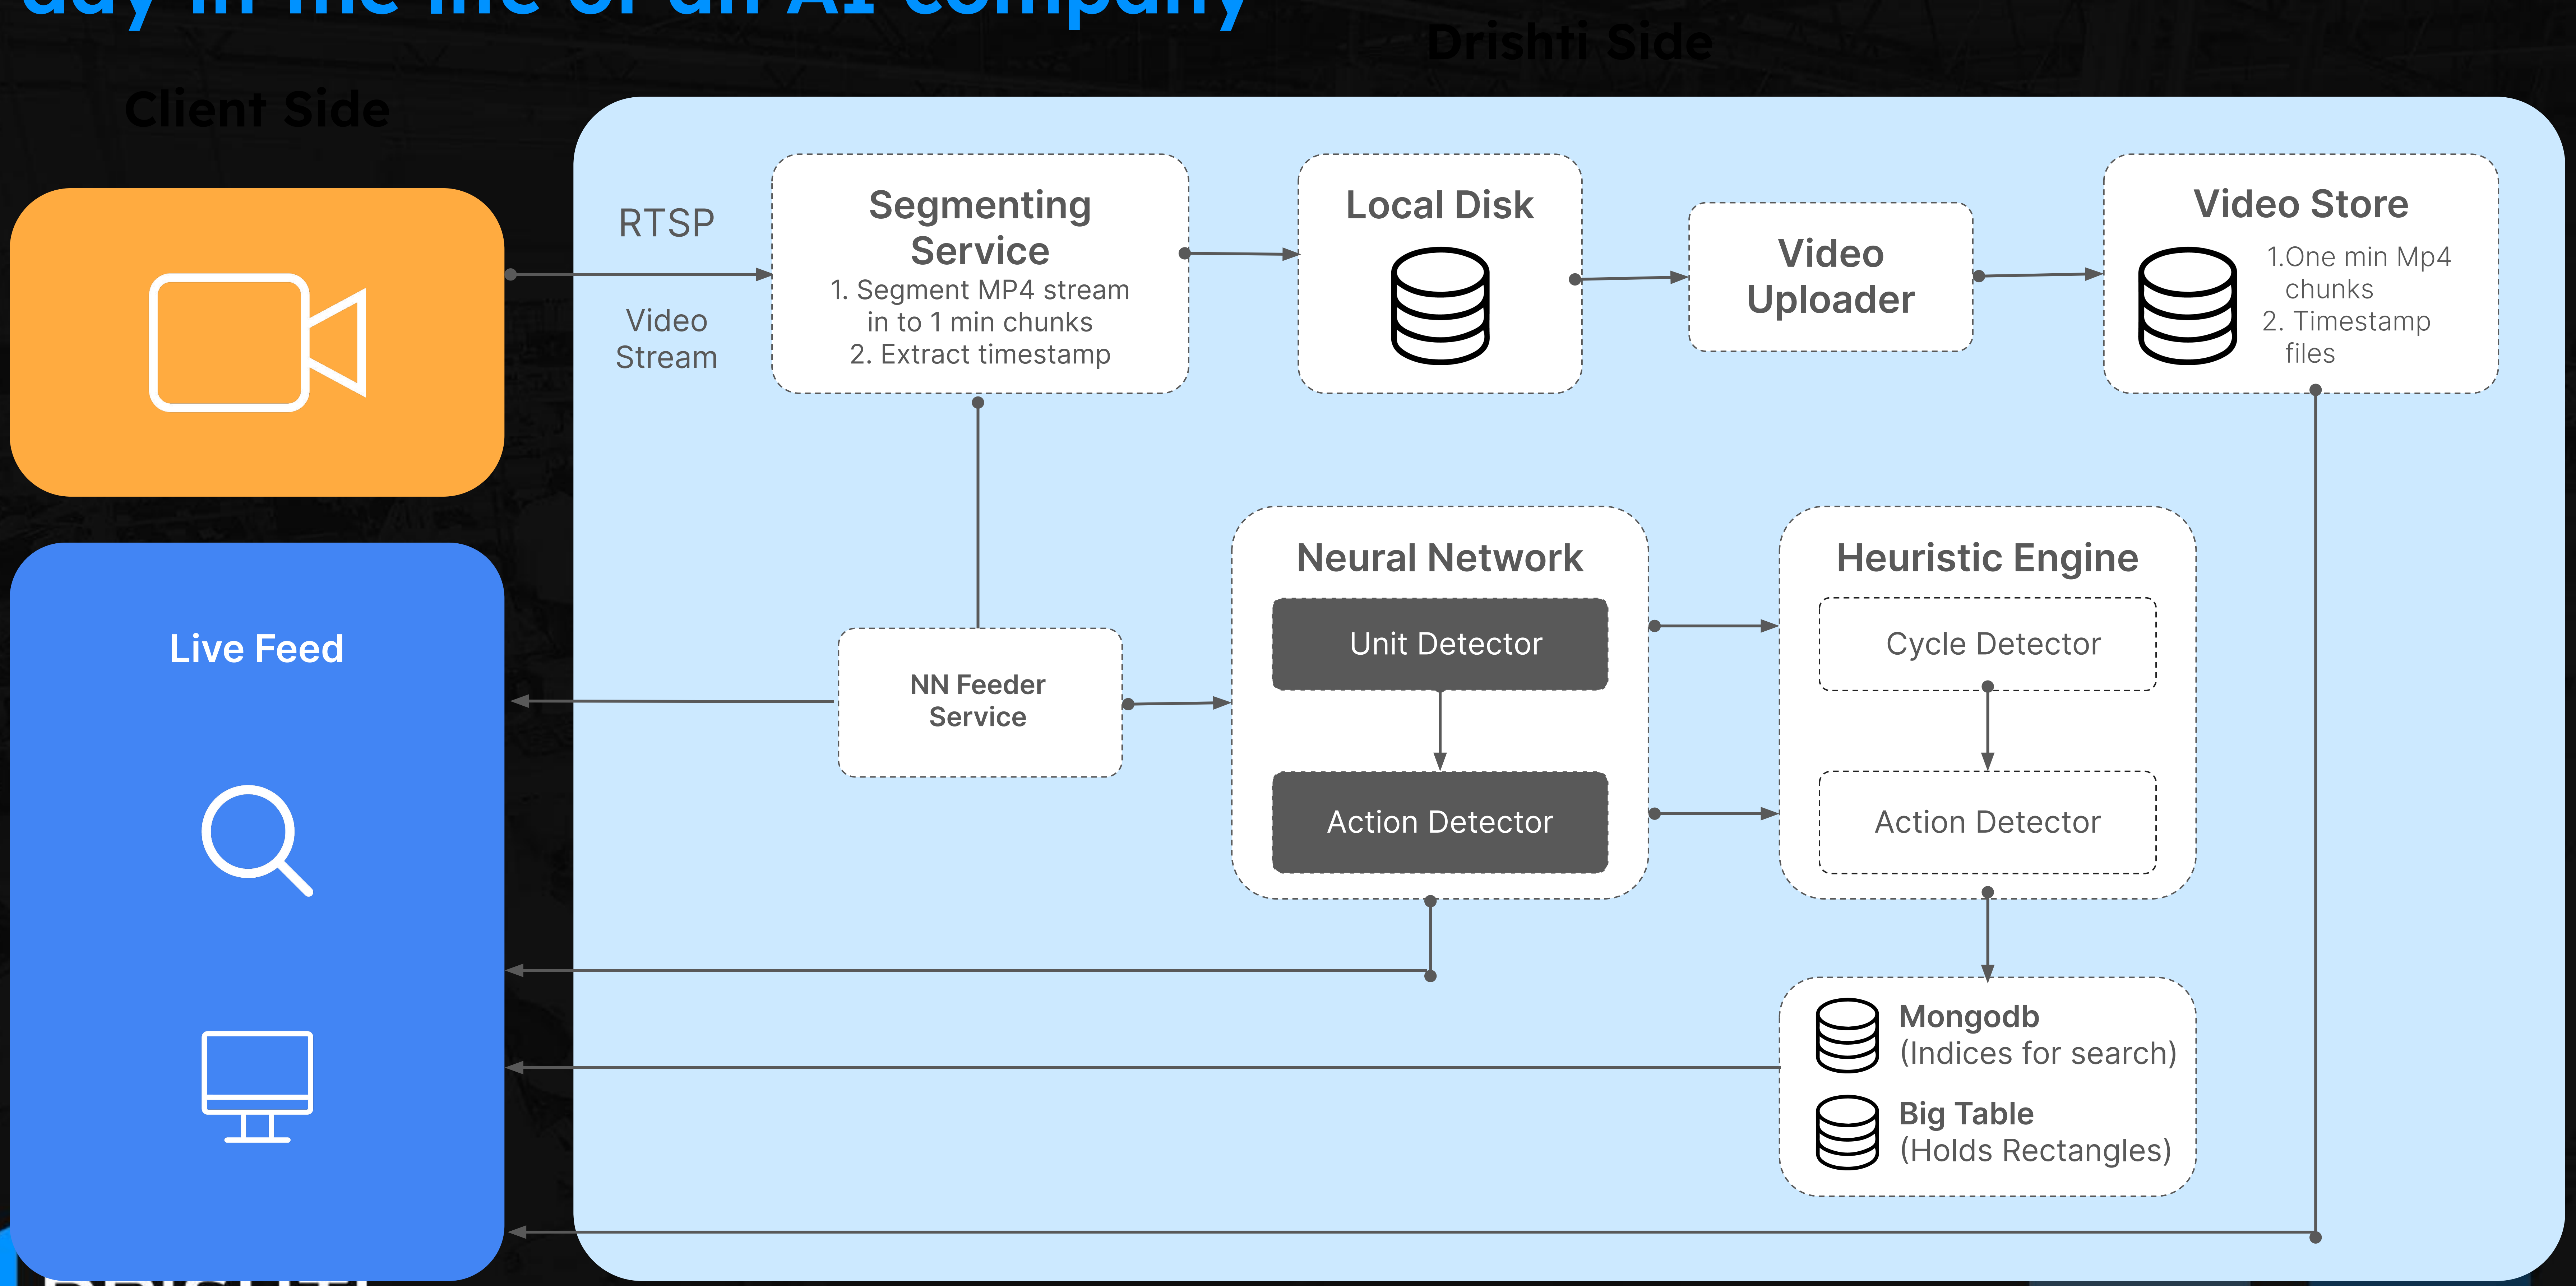

Impact of the real world on the AI system

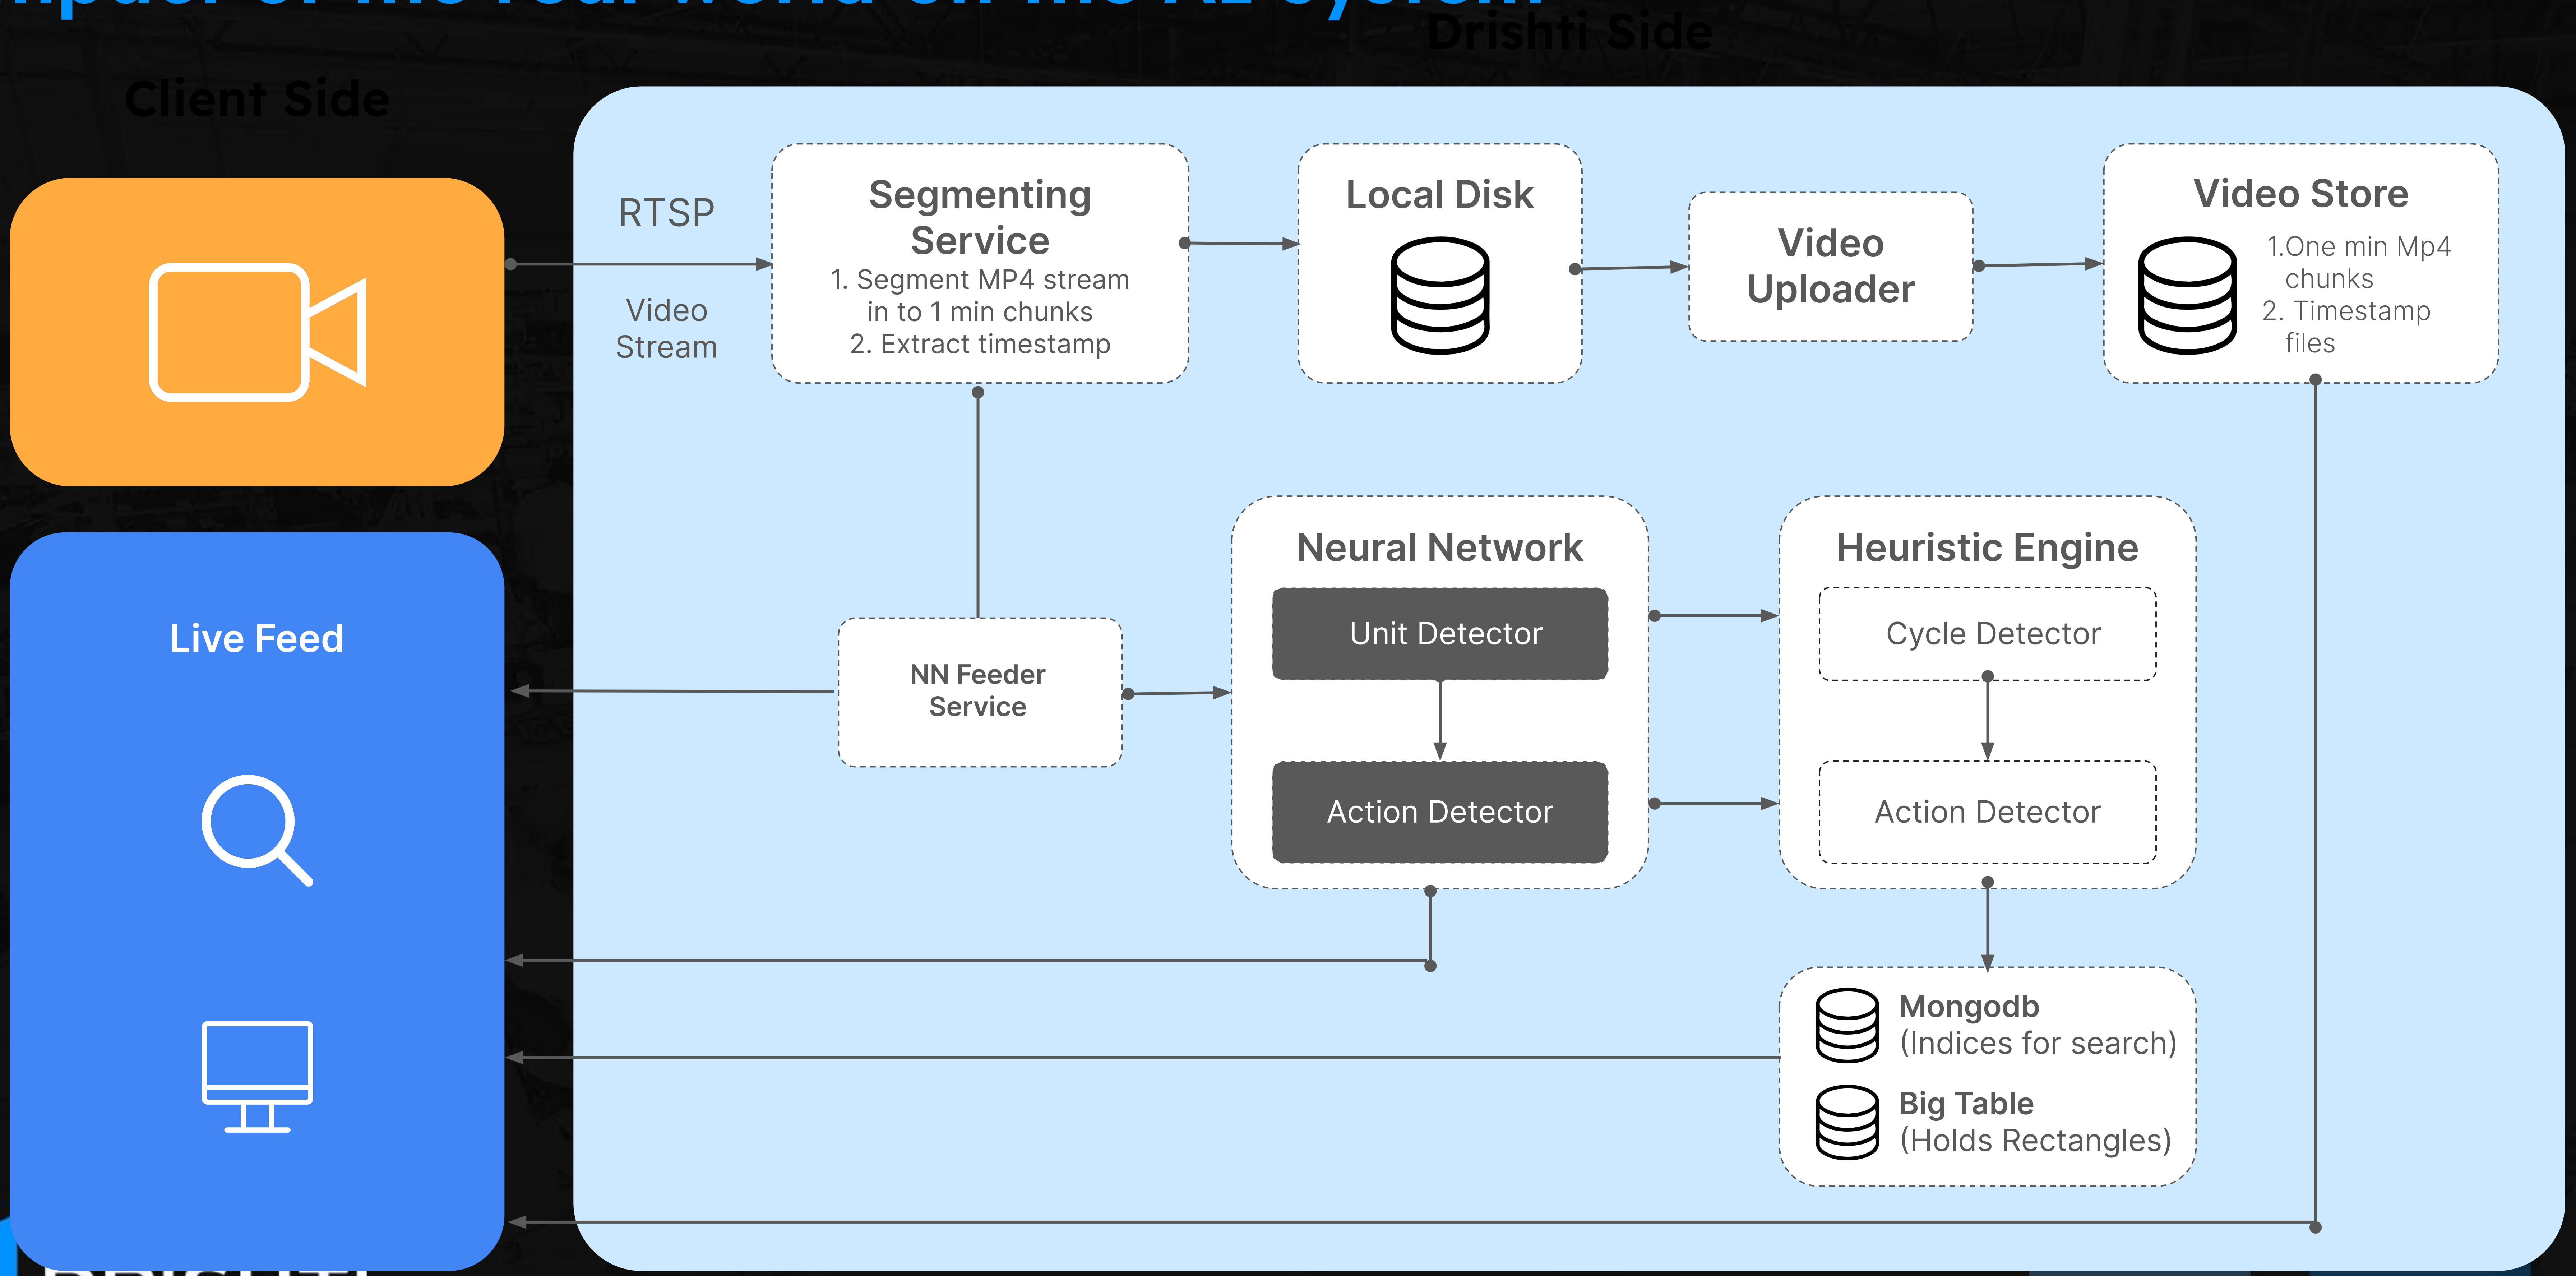

Drishti's AI technology must be robust to...

Data labeling → AMbiguaity, collaboration, tooling
Model accuracy → Physics, ensemble, heuristics
Data drift → Monitoring and retraining
Generalization

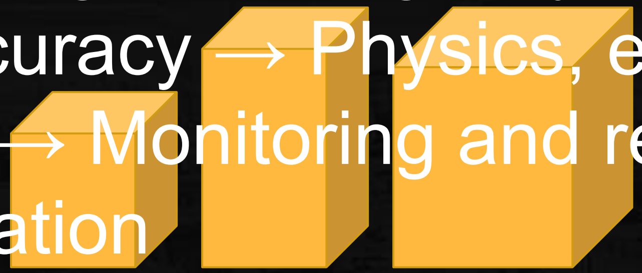

Variation in unit size

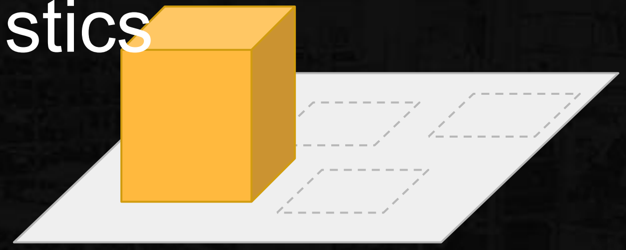

Indeterminate location

Irregular trajectories

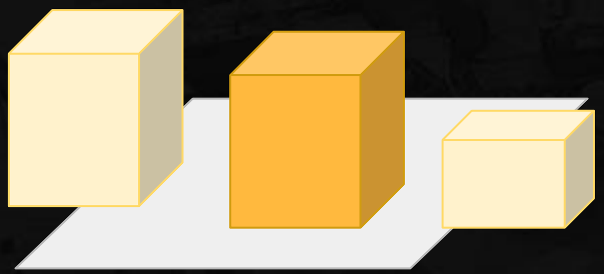

Multiple units in field of view

Camera occlusion (hand, head, tools)

Operator variance (hands, clothes, etc)

Lighting changes

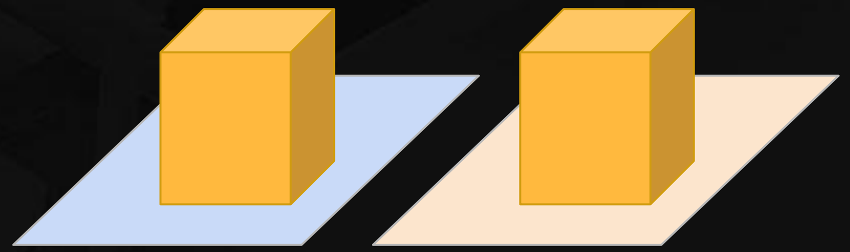

Background variation

DRISHTI

Software 1.0 vs Software 2.0

ML systems are different from typical software systems

Software 1.0
Traditional programming

Data
Program (rules)
→ Computer → Output

Software 2.0
Machine learning

Data
Weights
→ Computer → Program (probability)

Quality | Code quality | Data quality
Logic quality | Annotation quality
NN Accuracy

Software 1.0 vs Software 2.0

ML systems are different from typical software systems

Software 1.0
Traditional programming

Data

Program (rules)

→ Computer → Output

Software 2.0
Machine learning

Data

Weights

→ Computer → Program (probability)

Team structure

Code quality

Data quality

Logic quality

Annotation quality

NN Accuracy

Software 1.0 vs Software 2.0

ML systems are different from typical software systems

Software 1.0
Traditional programming

Data

Program (rules)

Computer

Output

Software 2.0
Machine learning

Data

Weights

Computer

Program (probability)

Workflow

Code quality

Logic quality

Data quality

Annotation quality

NN Accuracy

Software 1.0 vs Software 2.0

ML systems are different from typical software systems

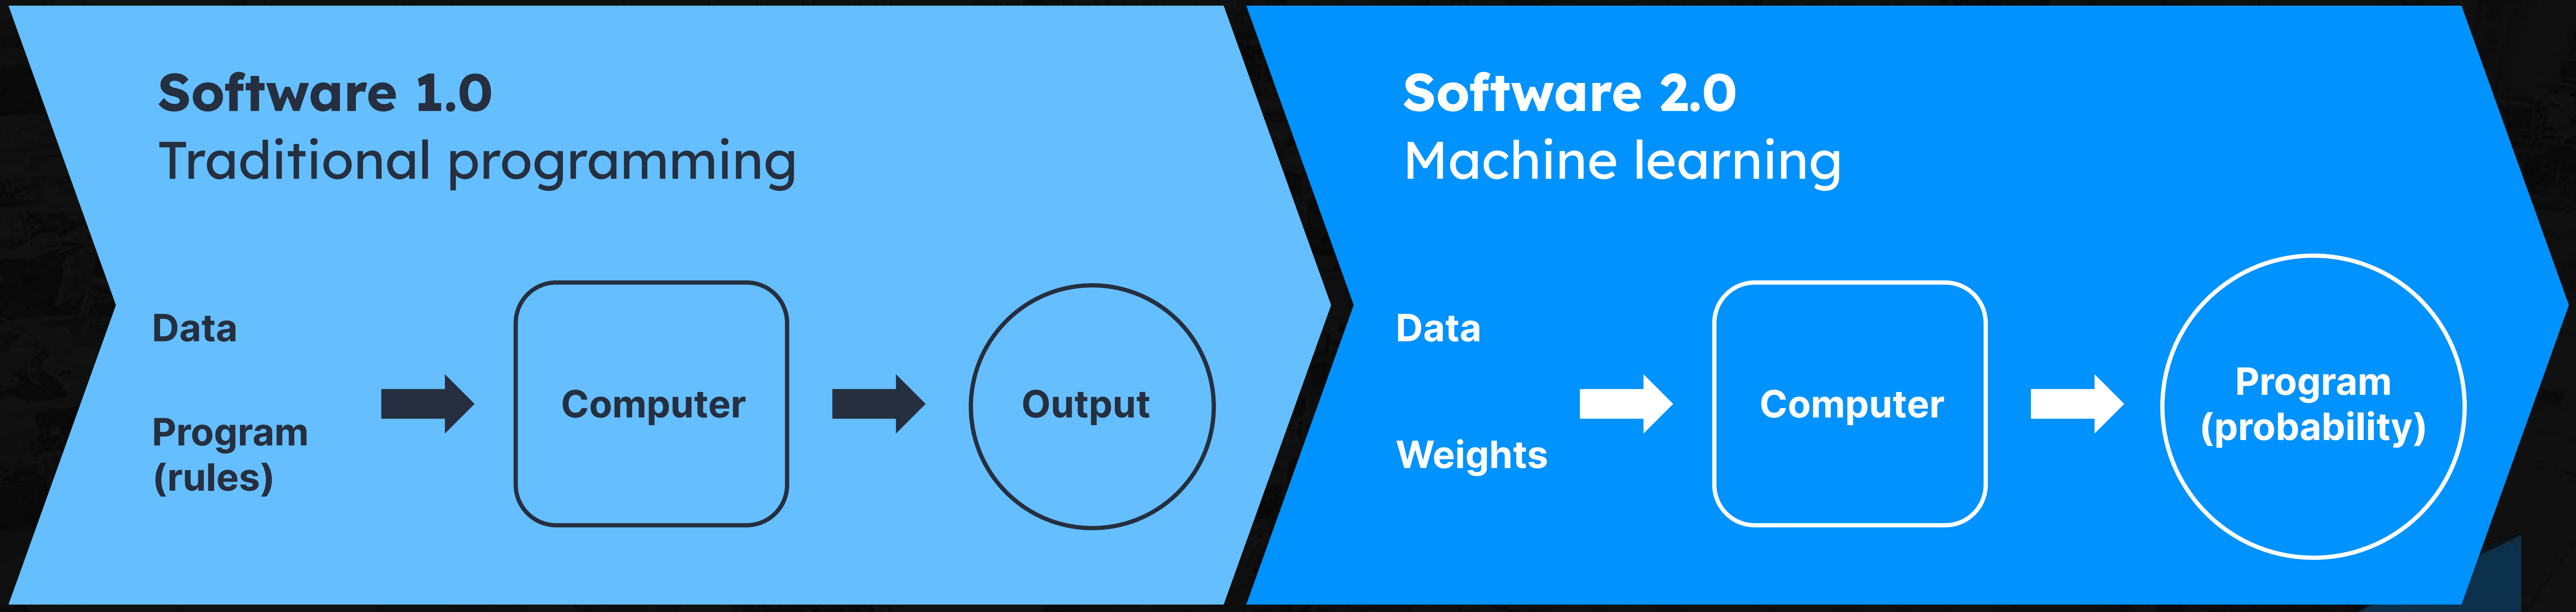

Measurement Code quality Data quality
Logic quality Annotation quality
NN Accuracy

DRISHTI

Software 2.0: It's all about data

ML systems are different from typical software systems

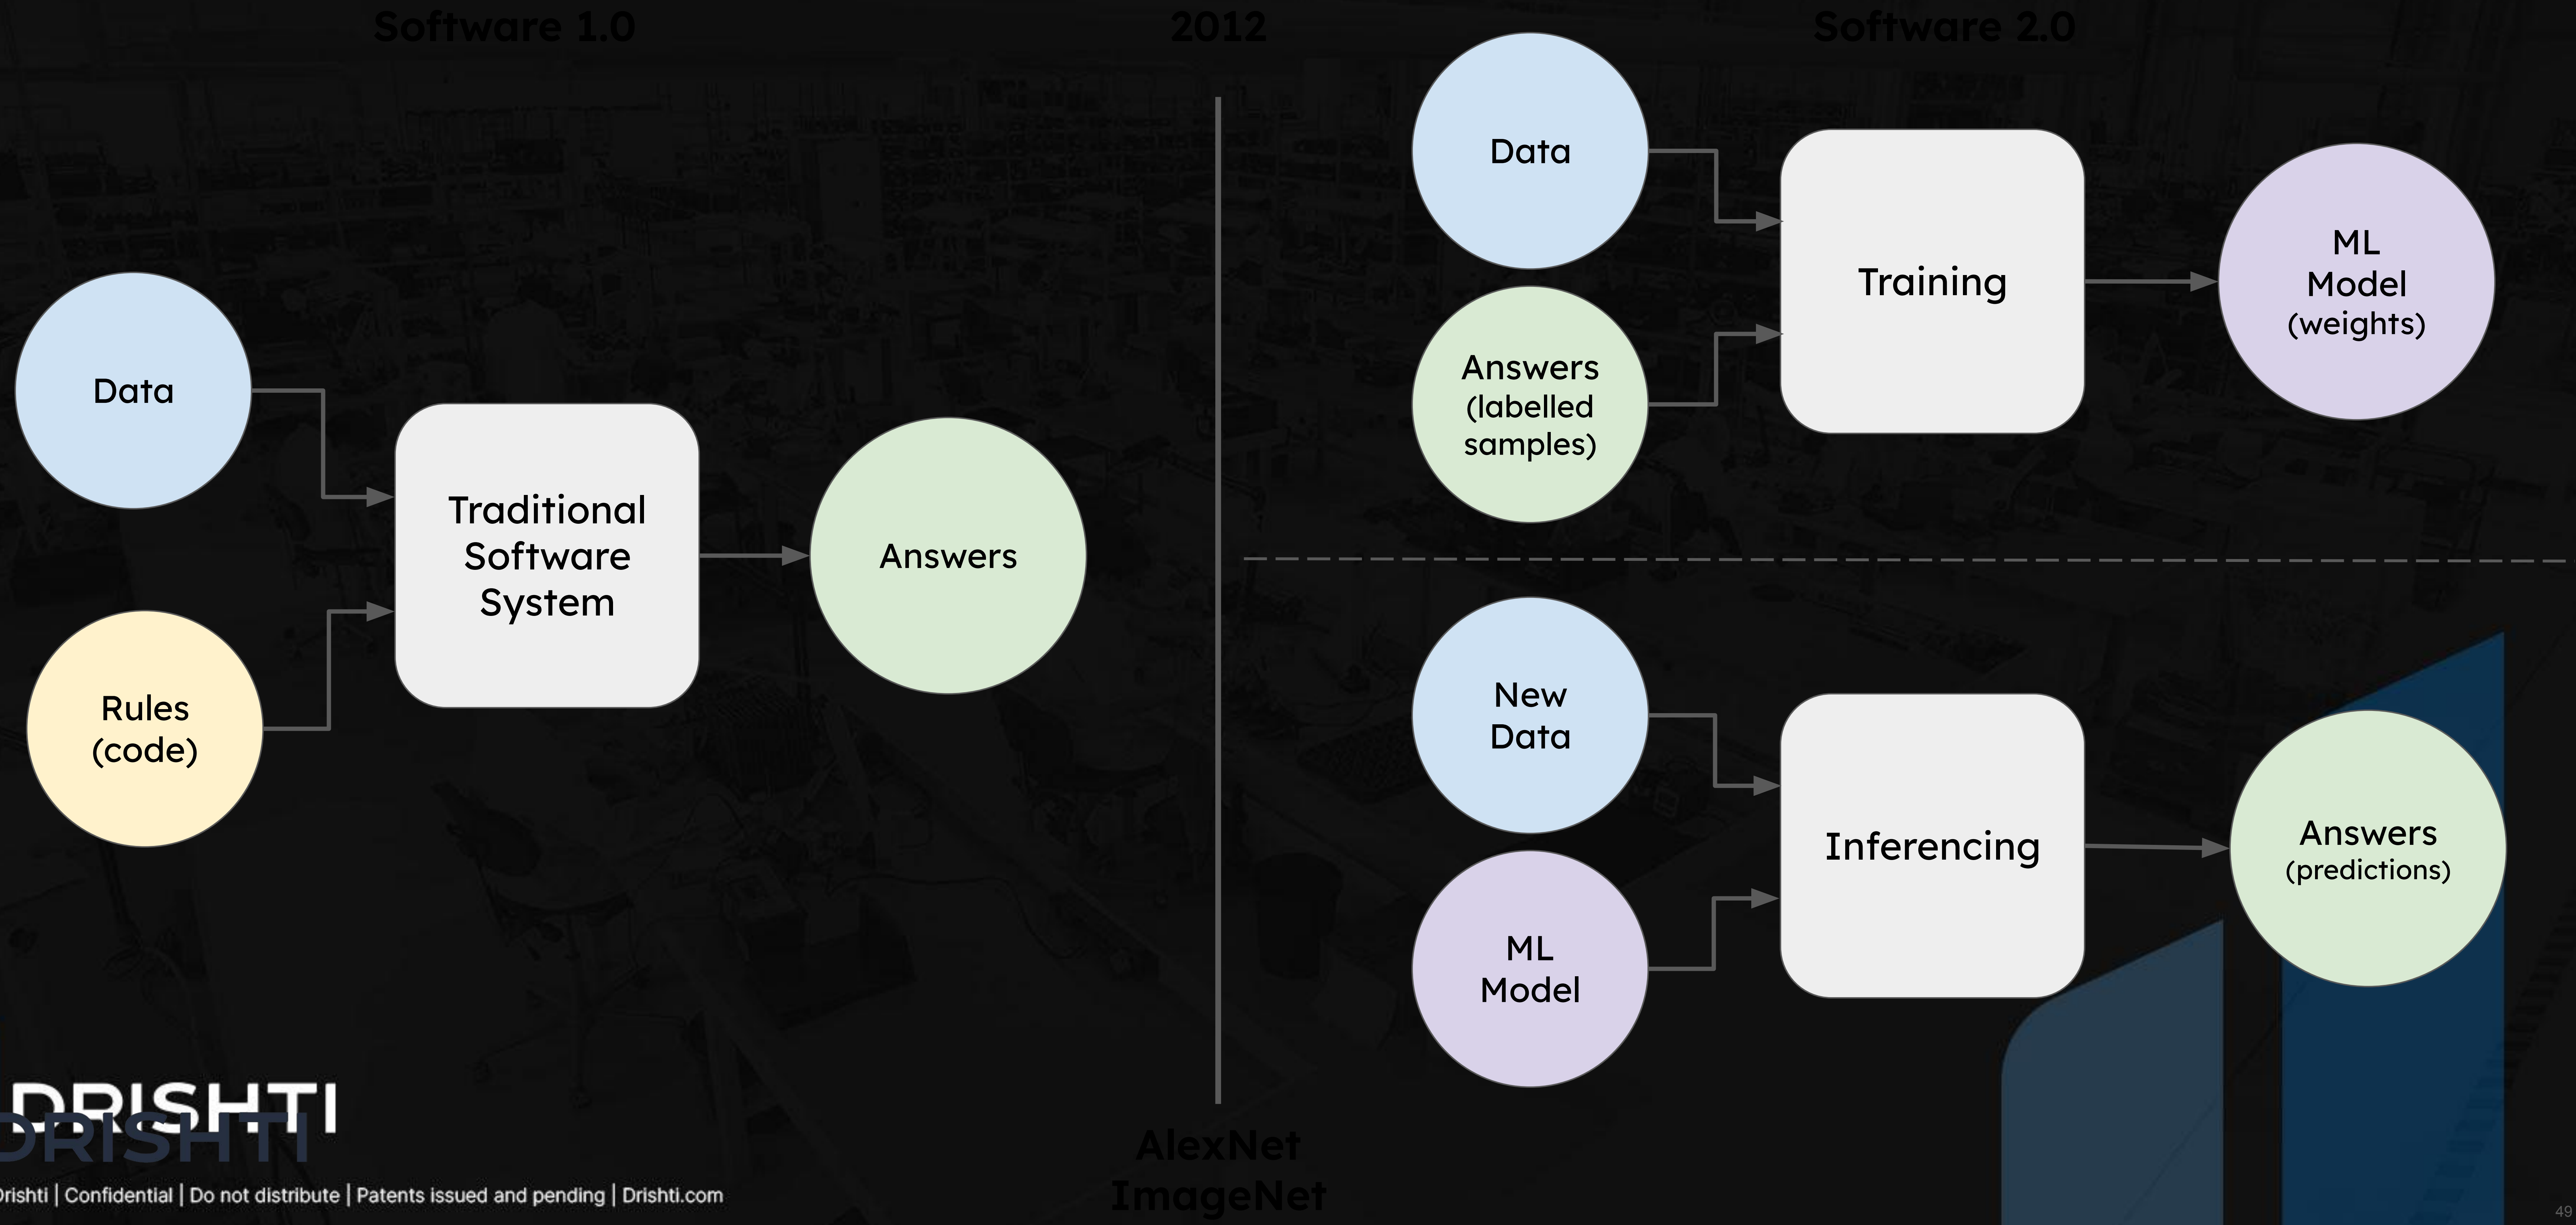

DRISHTI

DRISHTI

MLOps House

Trust and bias?

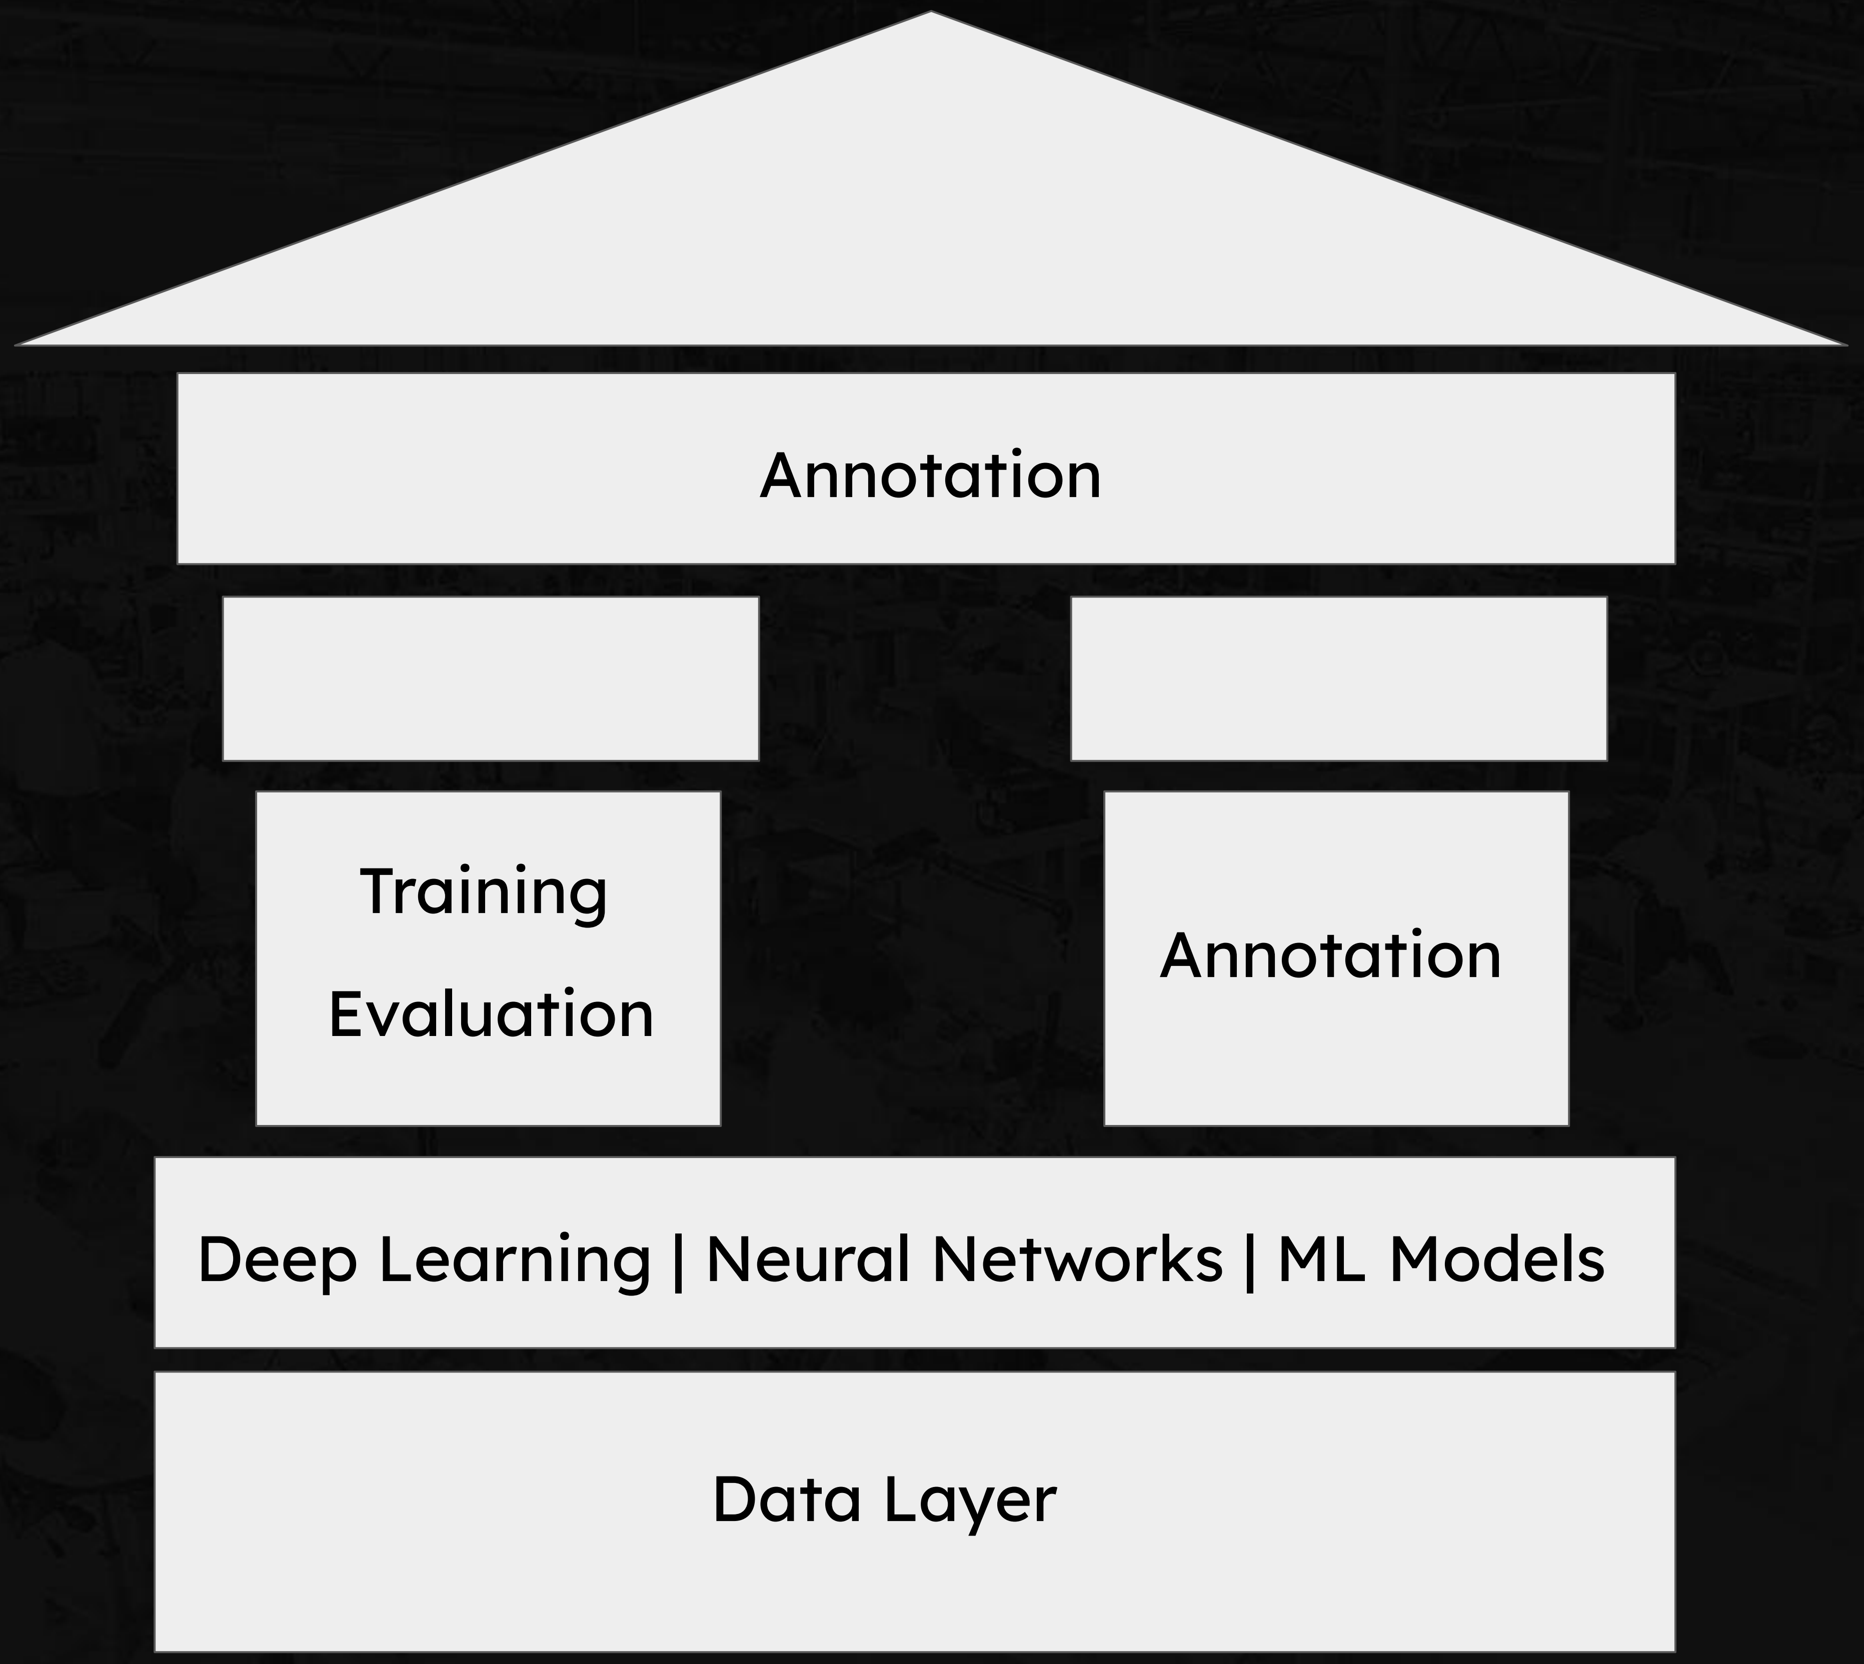

DRISHTI

MLOps House - WIP

Trust and bias?

Goal: Solve the business problem

Presentation | UI

Business Rules | Analytics

Machine	Human
Retraining	HITL
Monitoring	Annotation
Inferencing	Engineering
Training	ML

Deep Learning | Neural Networks | ML Models

Data Layer

Data quality

DRISHTI

MLOps House - Business view

Trust and bias?

Goal: Actionable, timely insights @ lowest cost

To produce impactful insights

Business workflow

Presentation | UI

Detection

Response

Product change
Process change
People change
Environment change

Machine monitoring

Retraining

HITL

Heuristics

Robust in the face of a dynamic world

Integration

<Blah>

Training

Accurate ML Models

Annotation

Inferencing

Creating quality data (data operations stability)

Wholesome Data Layer

DRISHTI

DRISHTI

Ops

AI

lla

Founder & Chairman

Sujay Narumanchi
Founding Engineer

MANUFACTURING LEADERSHIP COUNCIL
2021 Awards for Ford, DENSO + Hella

THE TOYOTA WAY 2ND EDITION
“A revolution in TPS”

Forbes
2020 Forbes AI 50

NVIDIA
2020 Top AI Startup

WORLD ECONOMIC FORUM
2019 Tech Pioneer

A picture is worth a thousand words.

A video is worth a million.

DRISHTI

Brand Colors

Main Color

Hex
#0092FF

Accent Color

Hex
#FF9914

Avilable Variants

Hex
#33A8FF

Hex
#66BEFF

Hex
#99D3FF

Hex
#CCE9FF

Secondary Colors

Hex
#1F4F7D

Hex
#29639B

Hex
#37D08D

Neutral Colors

Hex
#252F3F

Hex
#D4E0F5

Hex
#445460

Hex
#EBF1F8

Hex
#747F88

Hex
#BCC3C9

Gradient Colors

Hex
#0092FF
#29639B

Please make this into a sample gradient

Brand Font

abcdefghijklmo
pqrstuvwxyz
0123456789

abcdefghijklmo
pqrstuvwxyz
0123456789

Lexend Family is Drishti's official font for headings & sub headings. Headlines should always be written in sentence case, not title case.

Inter Family is Drishti's official font for copy & paragraphs.

Important Note: You can add both fonts to any Google Doc by selecting "More fonts"