Machine Learning and Data Science: Opportunities and Cautions in Science and Engineering

Discussions from initial efforts at Ford

Aug 10 2021

Devesh Upadhyay
Core AI-ML Methods and Quantum Computing
Ford Research

Focus

The intention of ML$_4$I is to foster and illustrate the adoption of ML methods for practical industrial outcomes. The forum will consist of a robust and open dialog between industry, research institutions, vendors, and academia to strengthen the technology transfer of ML methods to industrial needs.

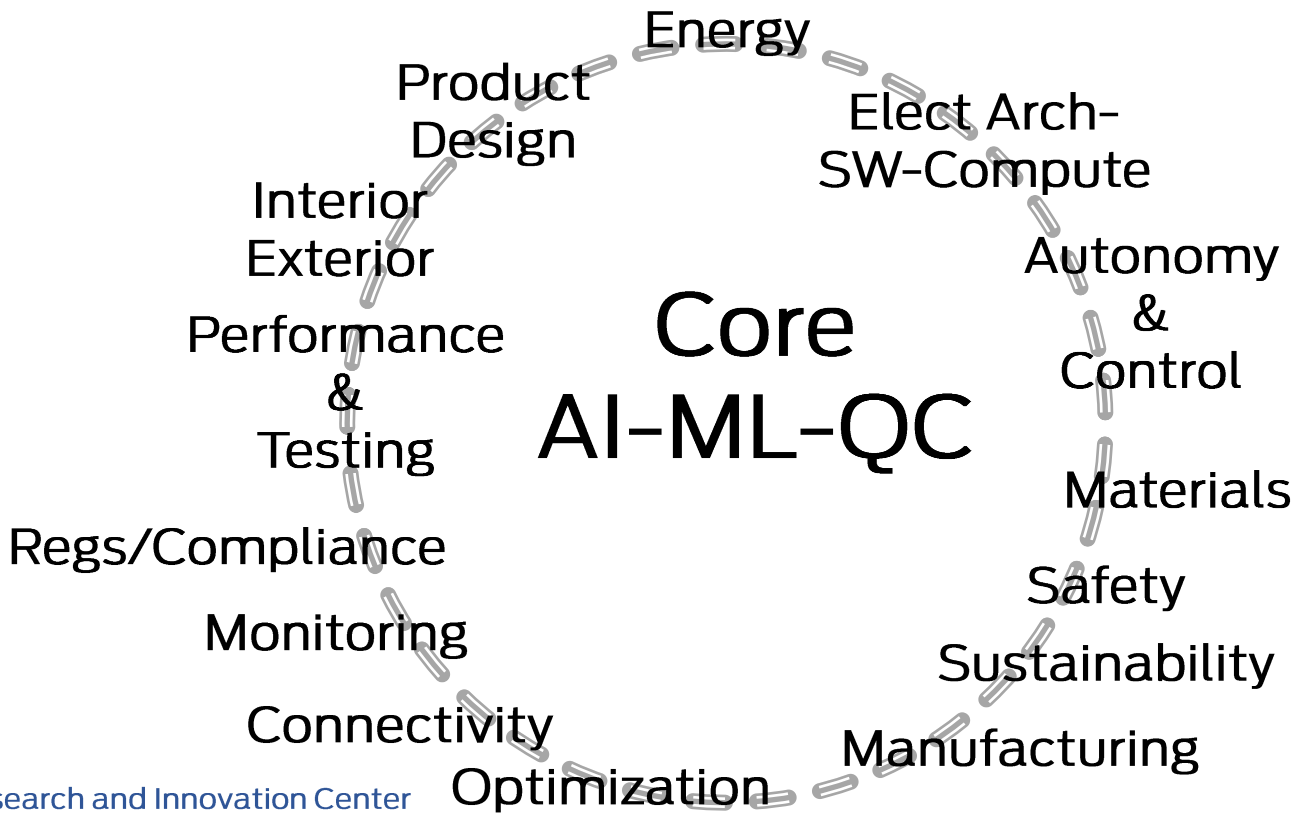

Discussion today

We provide high lights of the problem domains, specific to the Automotive Industry, and discuss impacts and experiences of ML and data-sci .

1. *Changing Landscape in Industry : Science, Engineering, Manufacturing*
   - ❑ We discuss the Automotive Industry
     - ❑ We do not discuss Autonomous driving
2. *Adoption of ML and Data Science in Engineering/Manufacturing*
   - ❑ *Enablers , Pressures and Opportunities*
3. *Use cases … some examples from industry (Ford)*
   - ❑ *A mix of examples to show case the span of ML-Impact*
4. *Outlook, Opportunities , Challenges*
   - ❑ *Data*
   - ❑ *Robustness*
   - ❑ *Causality*
   - ❑ *Cautions*

Changing Landscape in Science and Engineering

- We are in "The Fourth Paradigm"
- Science -- Engineering -- Manufacturing in industry
  - Old vs New
- Manufacturing – Industry 4.0
- Labs of the Future
  - High throughput research
  - Looped Intelligence
  - Continuous testing and Calibration

Data-Models-Compute … progression to the 4th paradigm

Friedmann's Cosmic equations built upon Einstein's Field Equations and generalizes the expansion of the universe → physics-math

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G\rho}{3} + \frac{\Lambda c^2}{3} - \frac{kc^2}{a^2}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right) + \frac{\Lambda c^2}{3}$$

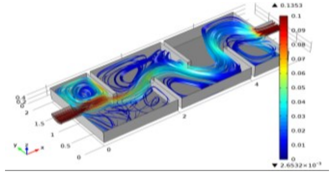

Navier Stokes … cannot be solved except for degenerate cases and very simple geometries → DNS, LES

Adapted from:
The Fourth Paradigm, Data-Intensive Scientific discovery
Tony Hey et.al

Empirical Science
Described natural phenomena but lacked extensions beyond observation

Theoretical constructs
Mathematical framework allowed generalizations based on Laws

Numerical Solvers
*Progress through Numerical (computational) solutions to complex problems**

4th Paradigm
Large data
Large dimensions
Large Compute
New Methods -- ML

Copernicus, Keppler, Galileo,…

Discovery based on observations.

1000+

Maxwell, Einstein, Friedmann

1650+
Newton

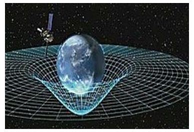

Turing, Von Neumann, Moore's Law

1940+
ENIaC

2000+

- GPU's
- SOC's
- FPGA's
- QC,…

Have observations (**data**)
Needed new instruments
Needed theories for generalizations

Developed theory
Needed solvers for more complex problems

Developed Numerical solvers
Needed better compute and better numerical methods for speed

Have new compute
Developed new methods (**data**)
Move to newer compute - QC

Science-Engineering & Manufacturing

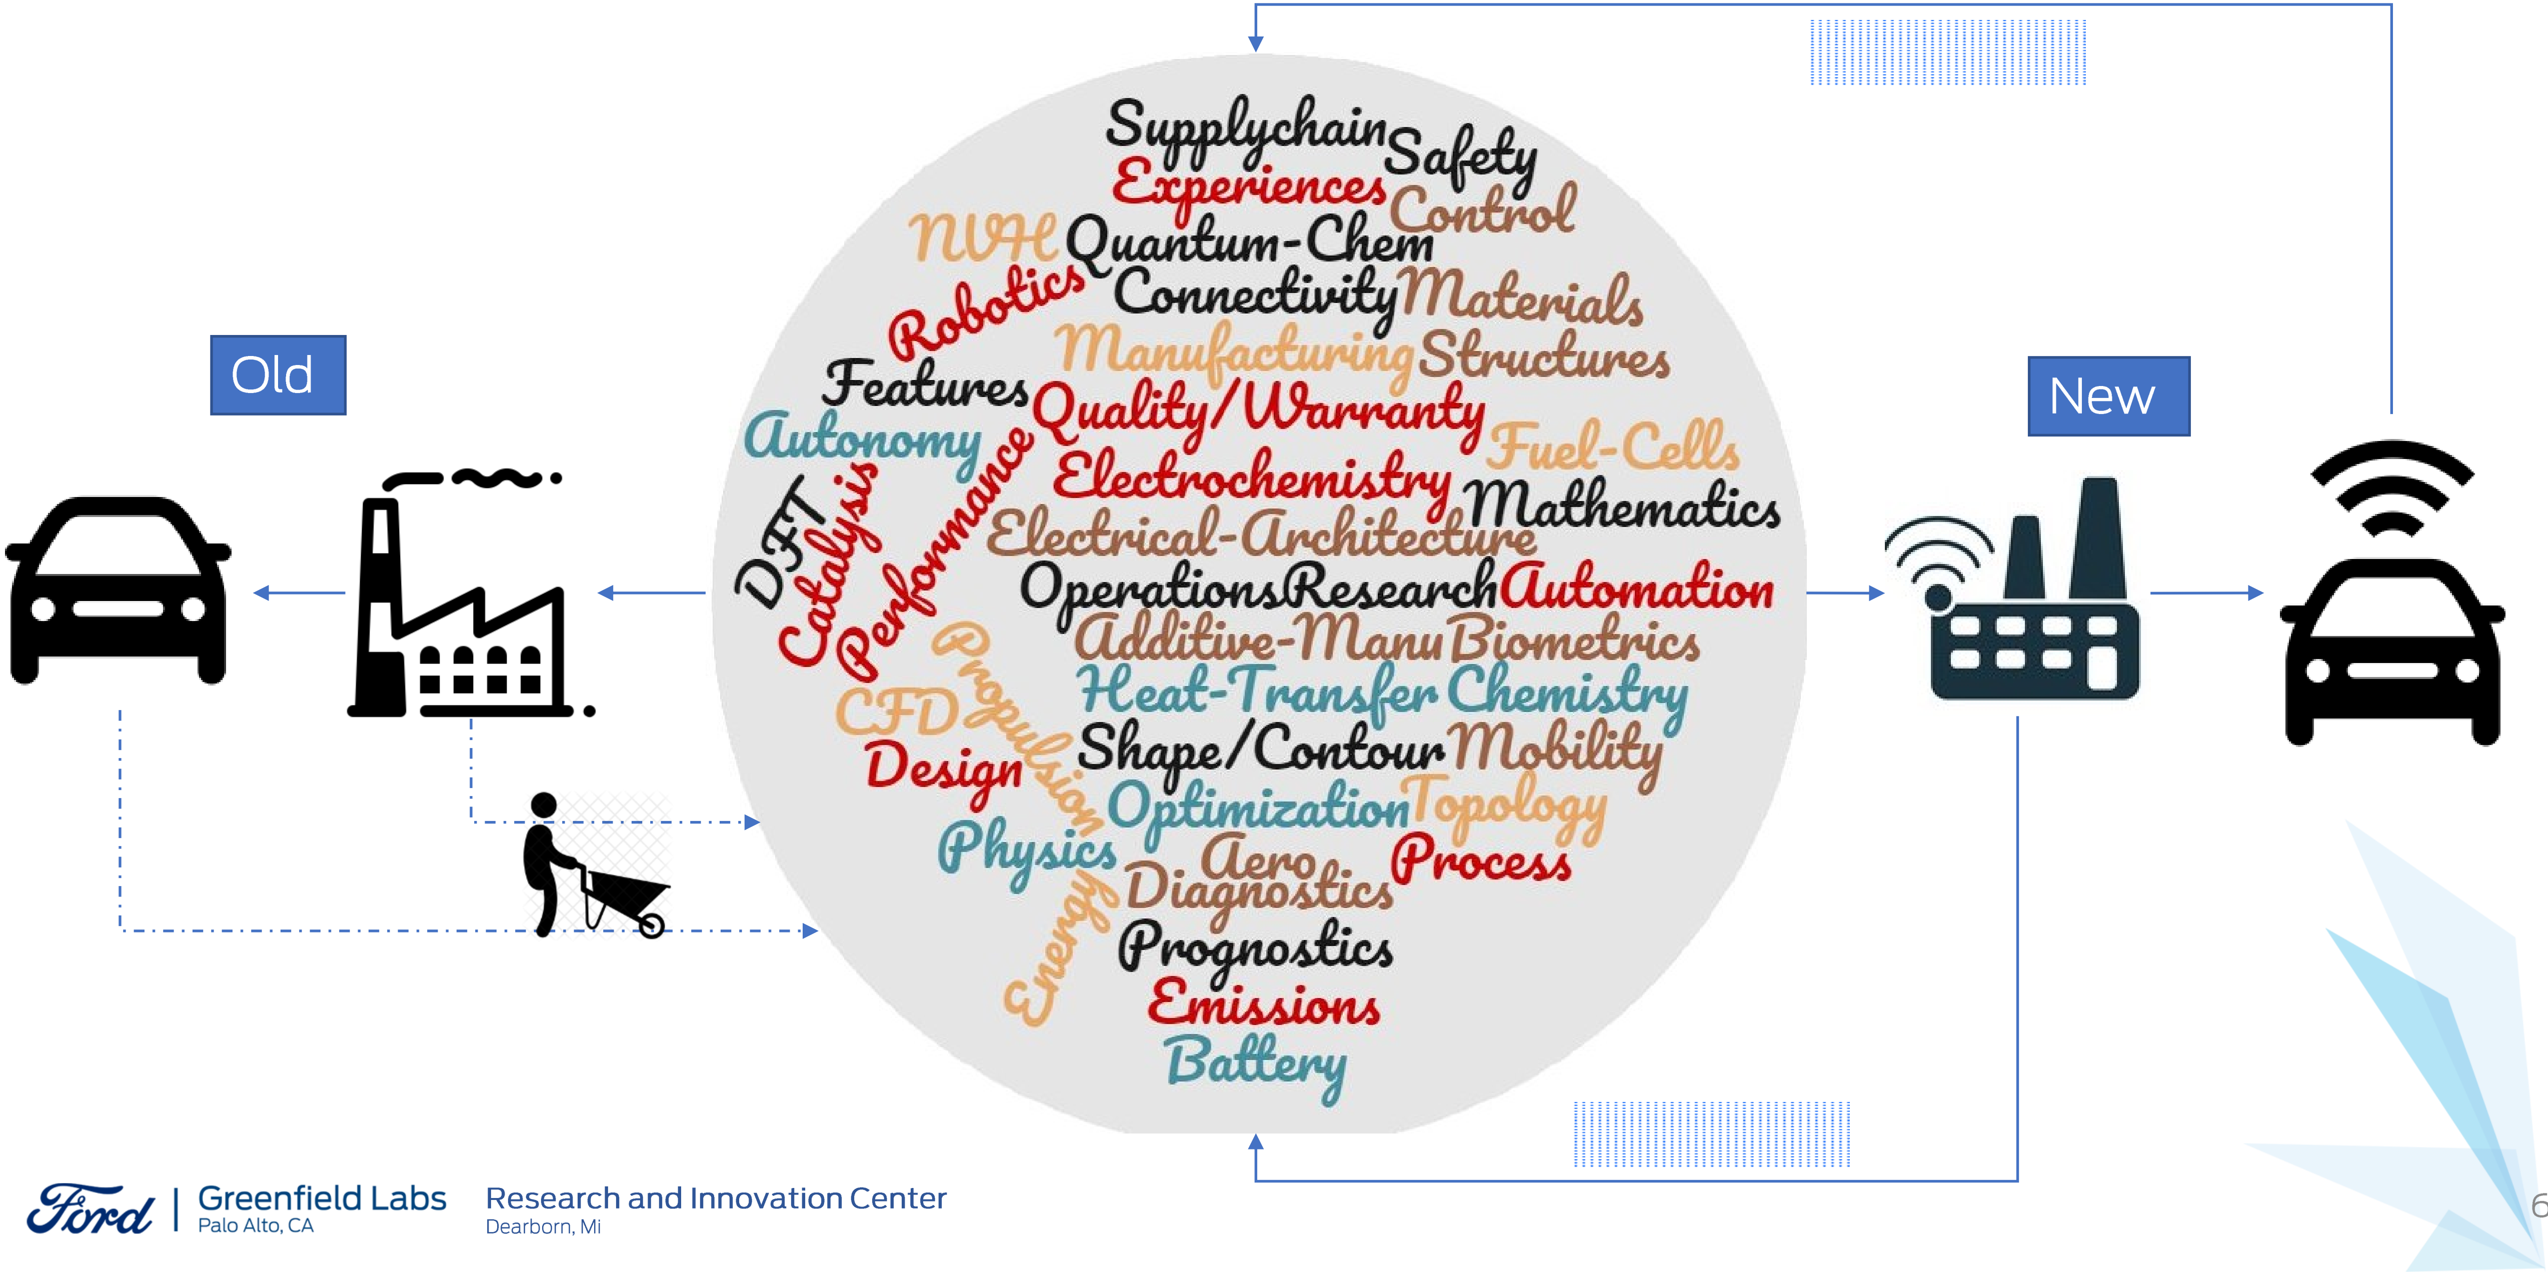

Manufacturing: Industry 1.0 → 4.0

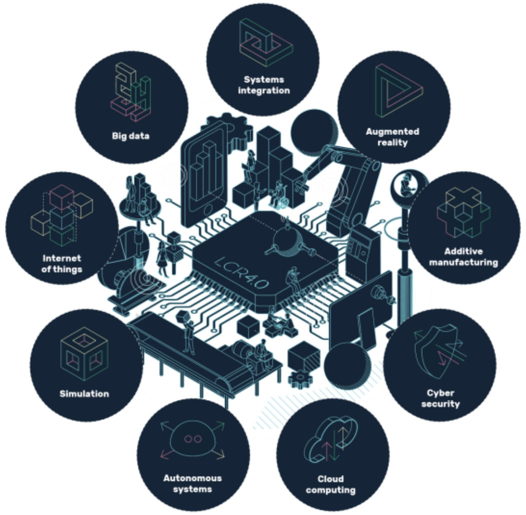

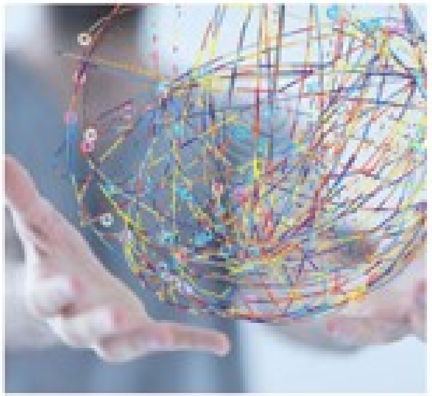

– INNOVATION CYCLE –
MONTHS TO YEARS

– INNOVATION CYCLE –
WEEKS TO MONTHS

– INNOVATION CYCLE –
HOURS TO WEEKS

– INNOVATION CYCLE –
MINUTES TO DAYS

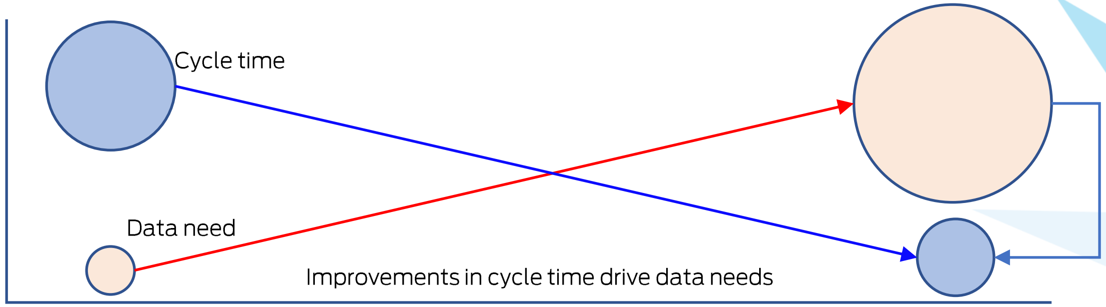

Accelerated discovery with Labs of the Future

LOOPED INTELLIGENCE

An autonomous ~~chemistry~~ laboratory runs experimental cycles intended to yield useful products ~~molecules~~. In the cycle, artificial intelligence models the experiment and designs a product ~~compound~~, robotic equipment runs the synthesis, and AI evaluates the output; researchers {and product designers} interpret the data and adjust experimental/design models or the goal definition as needed

Must have Large scale optimization

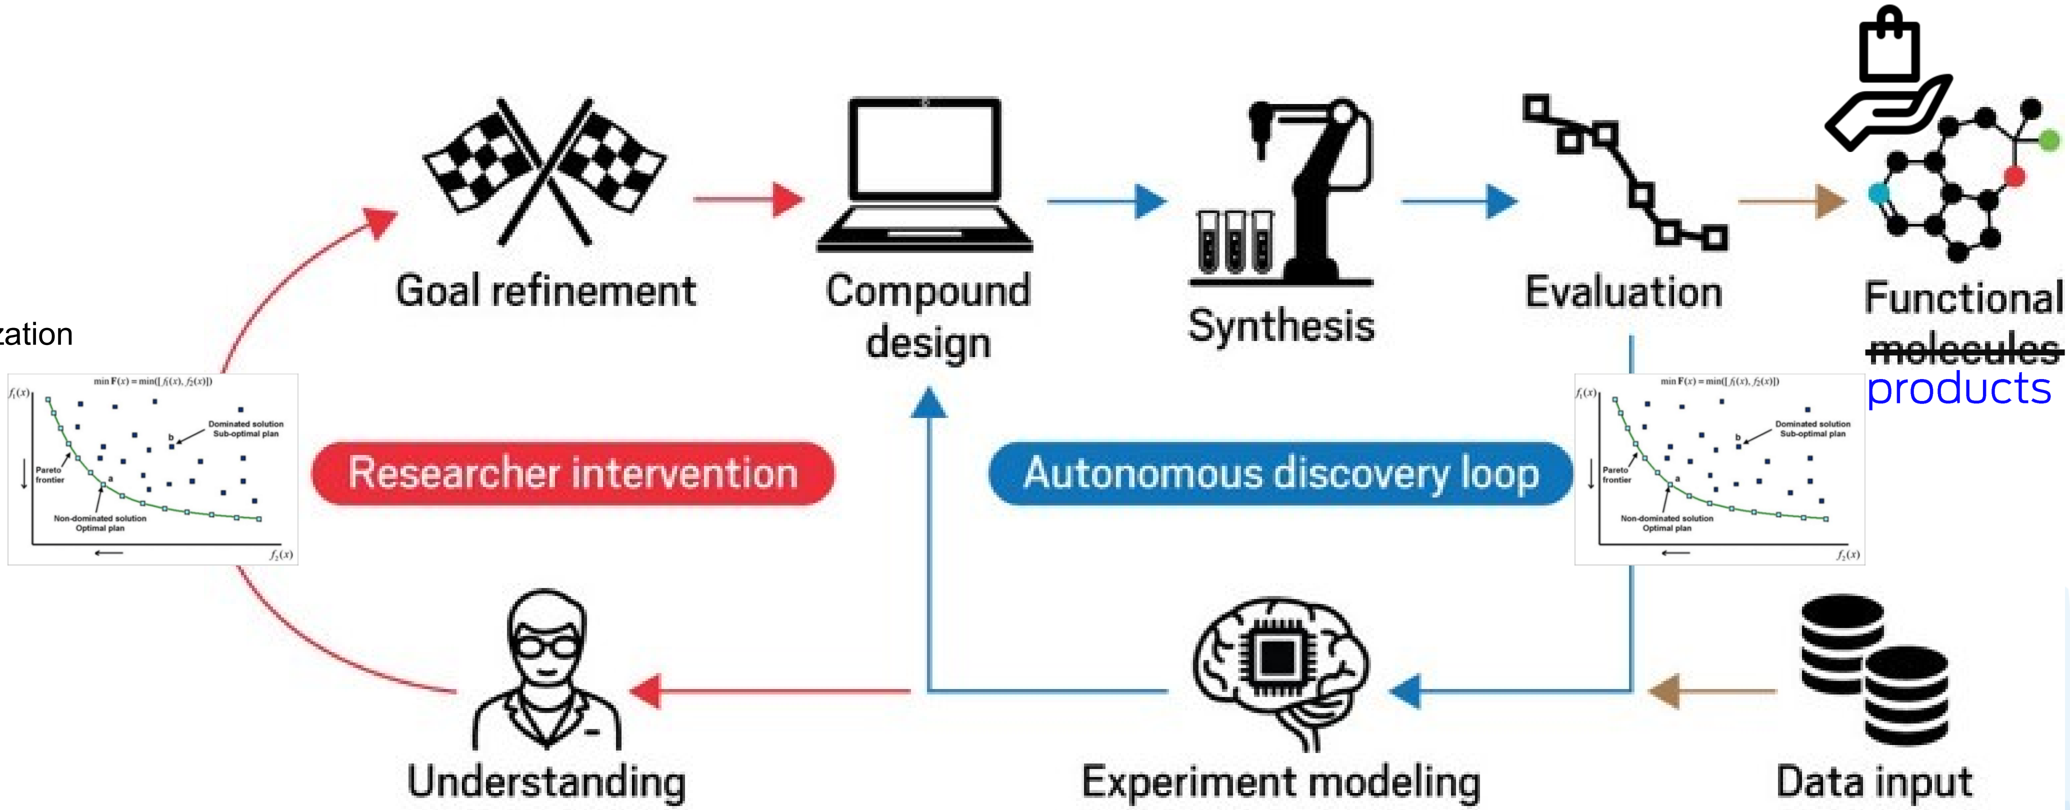

Adapted from: Connor Coley/Massachusetts Institute of Technology

Conny enables continuous testing and calibration

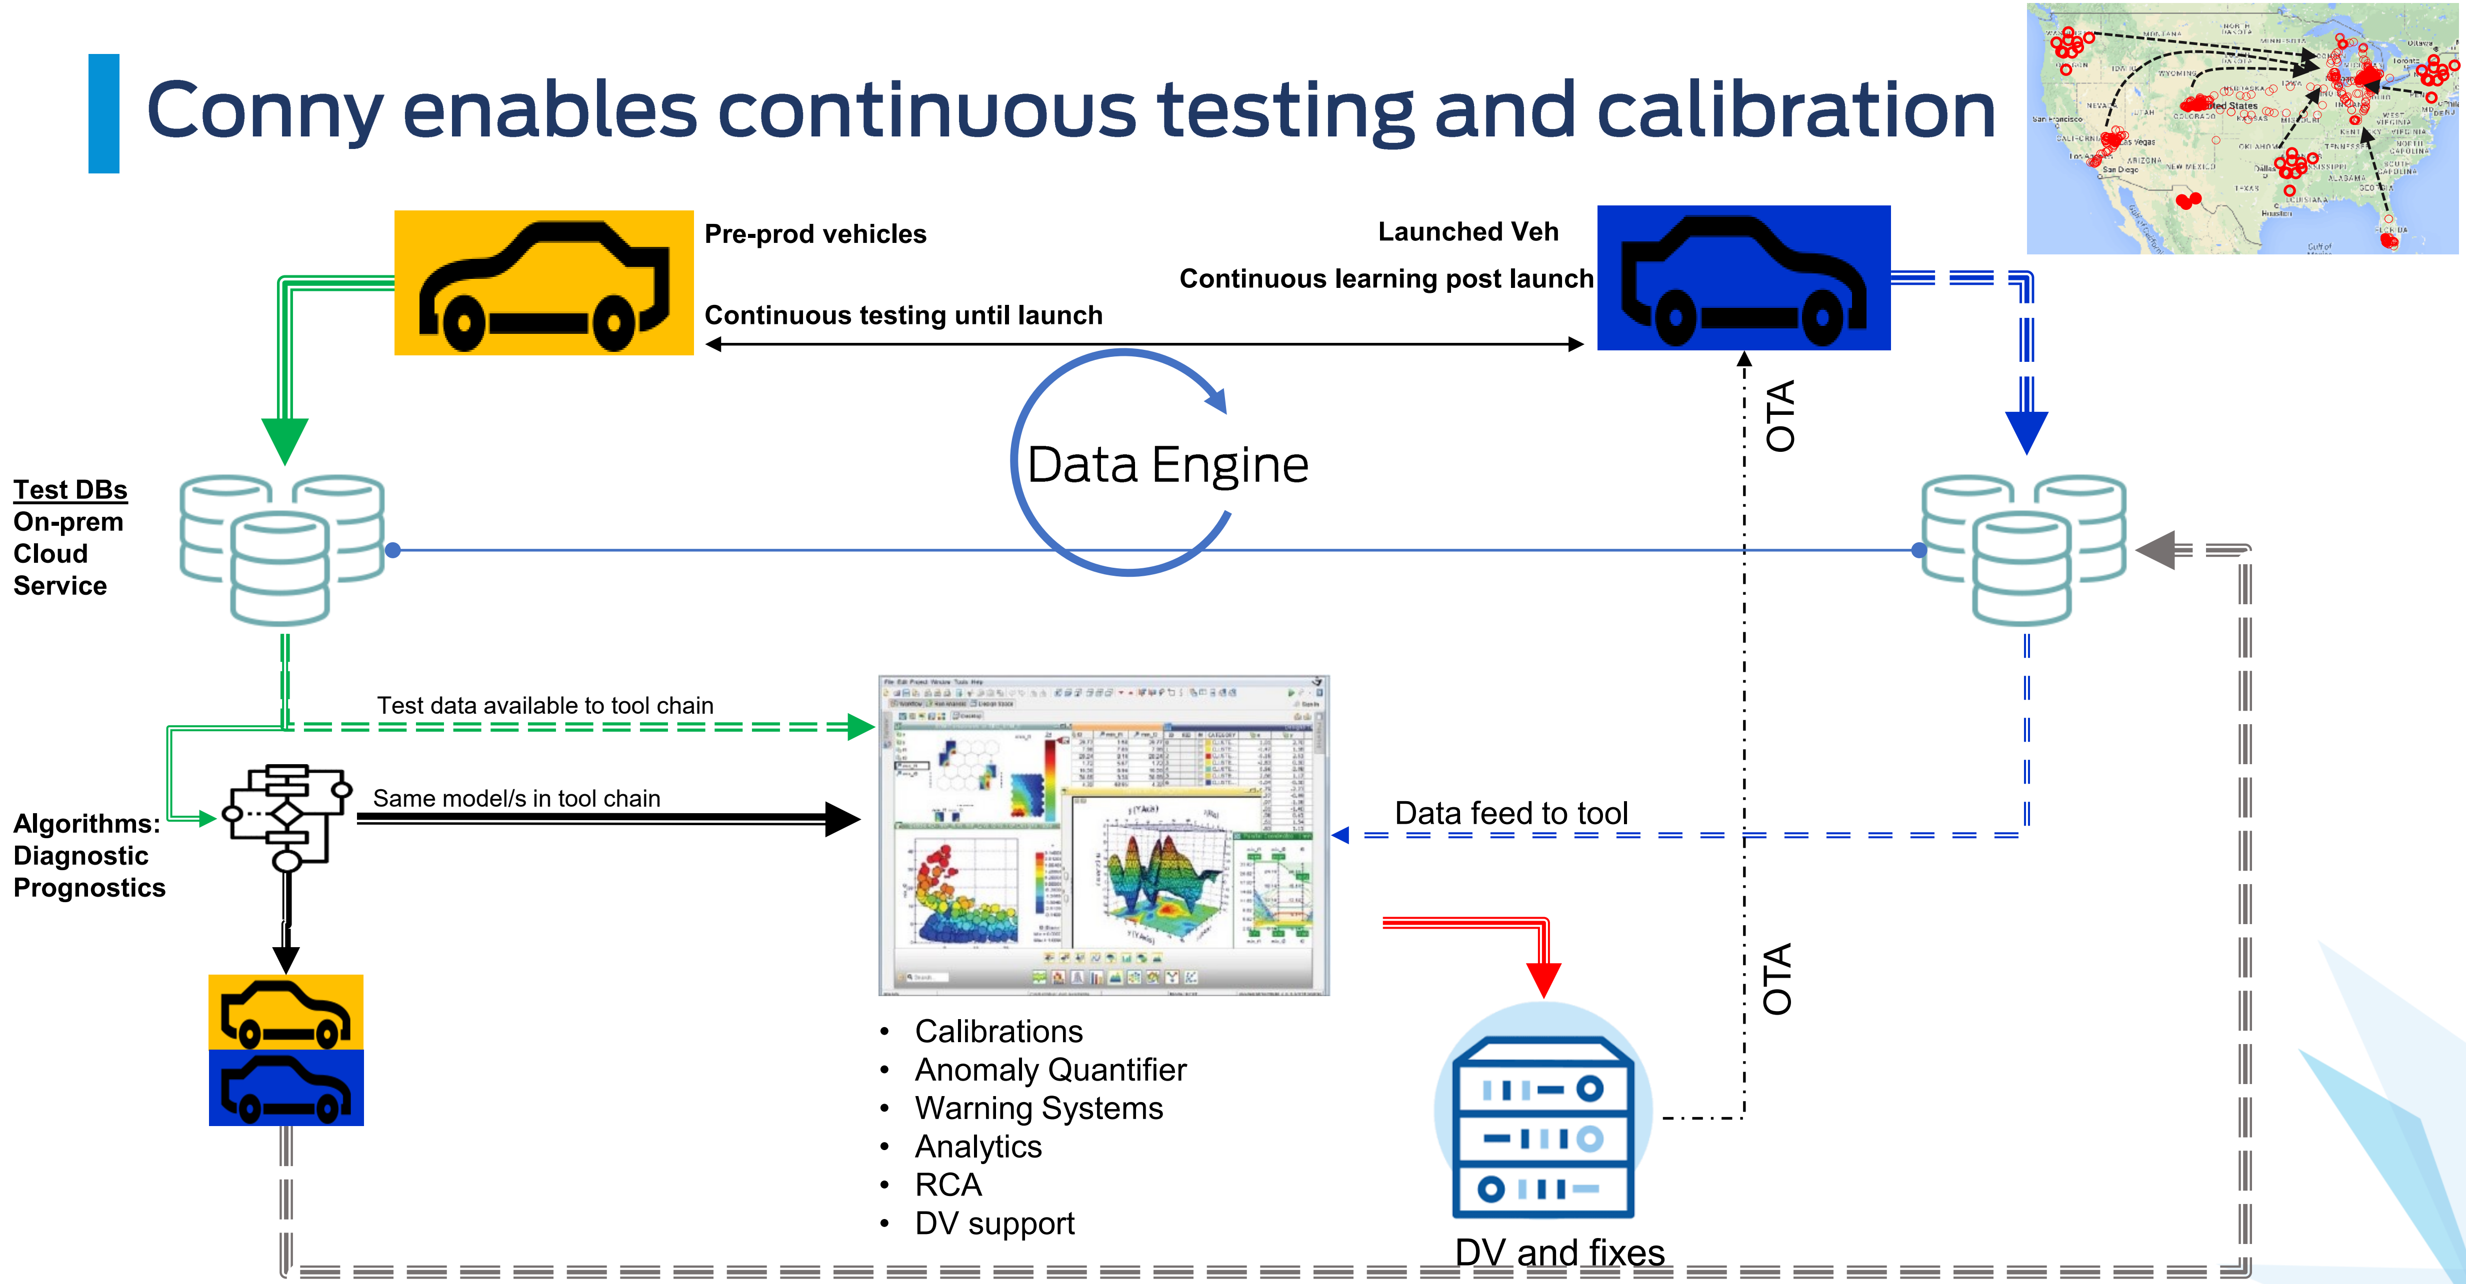

Adoption of ML and Data Science in Industry

3 key enablers for adoption

New sensing in vehicles and manufacturing

- *Perception, IoT , IIoT (5G)*
- *Audio, Language*

Demands/Pressures for adoption

- *Need for high throughput research → fast discovery*
- *Fresh look at high dimensional and/or complex problem domains*
  - *Need for high fidelity but fast surrogate models*
- *Automation → Autonomy, Control*
  - *Vehicles*
  - *Mobility*
  - *Manufacturing*
  - *Robotics*
  - ⋮
- *Asset Monitoring*
  - *Vehicles/Fleets*
  - *Manufacturing*

Emerging technologies and a changing landscape

3- Pillars: 1. Connectivity + New Sensing, 2. Algorithms, 3. Compute

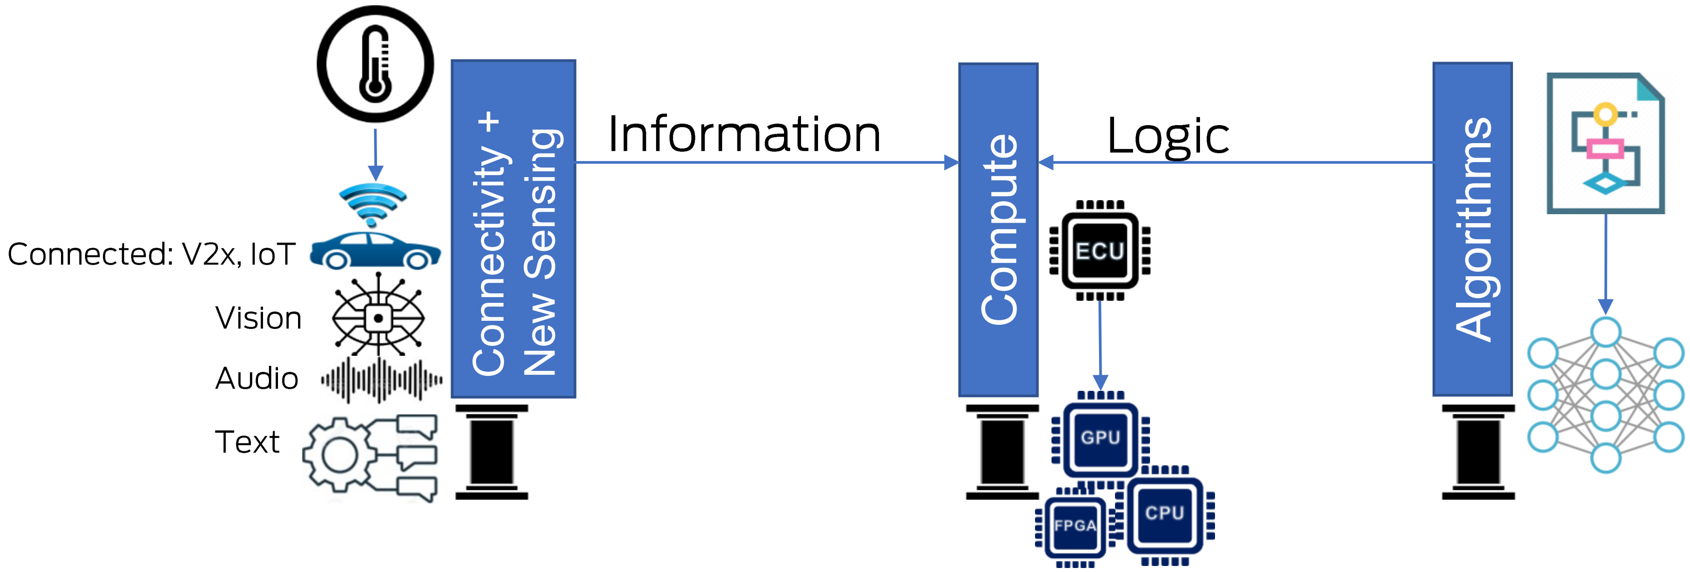

Factory IIoT

- Independent single- veh control
- Human in the loop
- Homogenous on board compute
- Onboard sensing

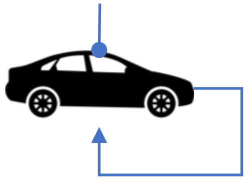

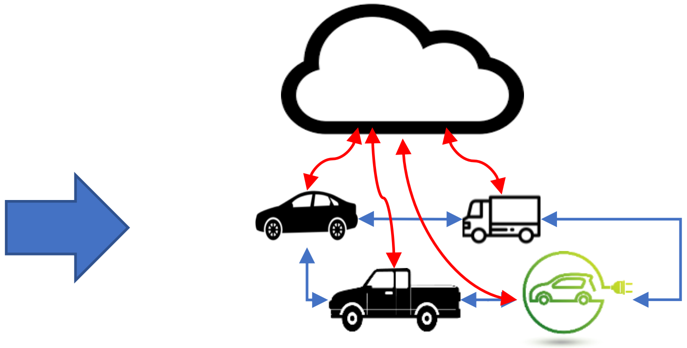

- Dependent multi- veh control
- Human may/may not be in the loop
- Heterogenous compute
- Onboard + V2x+ Edge+ Cloud+ new sensing

Sensing → new information sources are indirect sensing

Indirect sensing (unstructured) rich in Information that must be extracted requiring an additional layer of pre-processing and information encoding in data (labeling)

Sensing is expanding

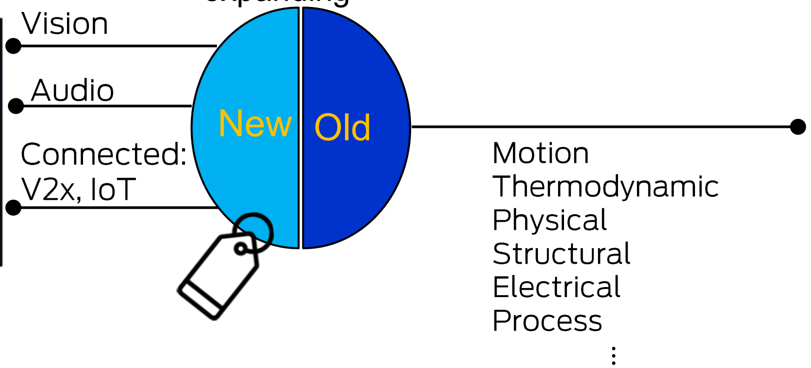

Information is directly encoded in the sensed parameter and is available for reconstruction via a known direct transform.

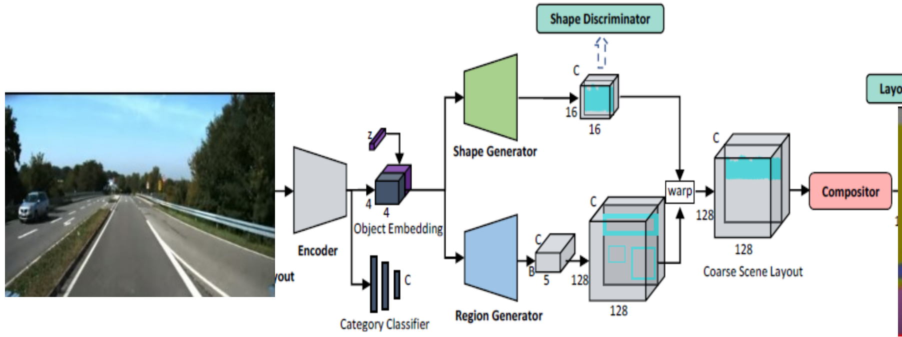

Context $=\psi(\hat{u}_i)$

Scene = $\varphi(u_i)$

Slow traffic → control actions ?

Vision in Manufacturing

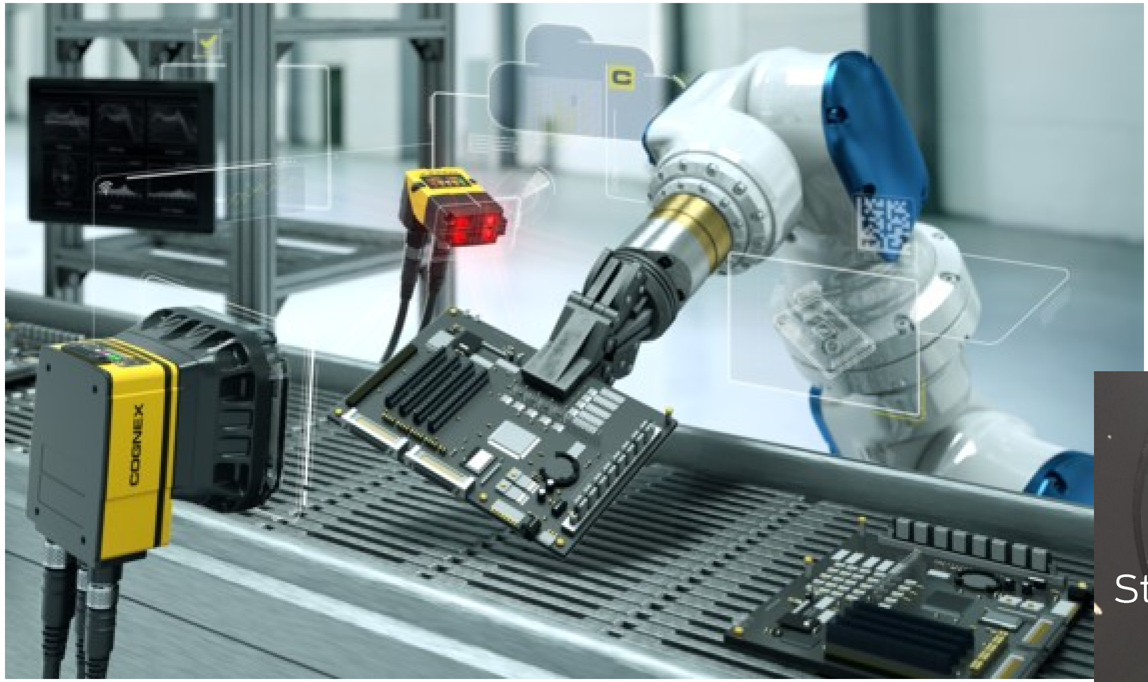

In situ error-proofing

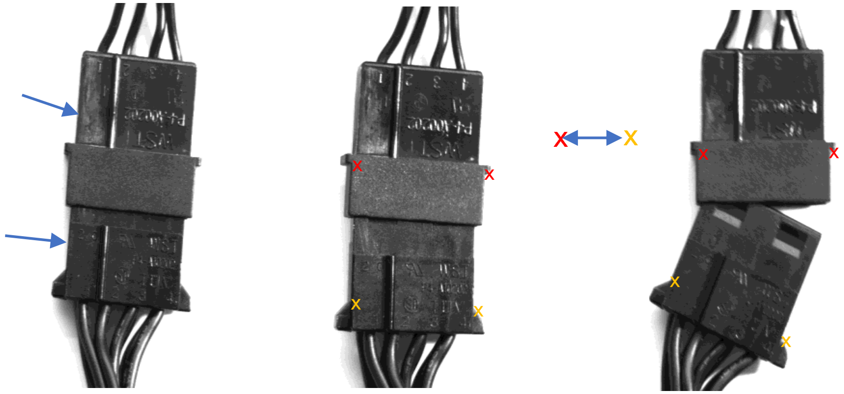

Feature Distance

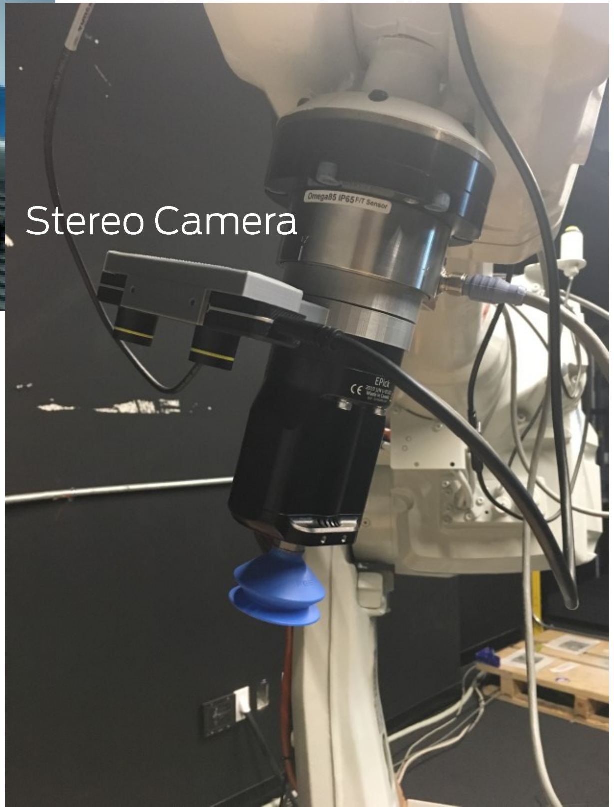

keypoints

Connectivity and AI improve decisions under uncertainty through continuous Learning and exploiting preview

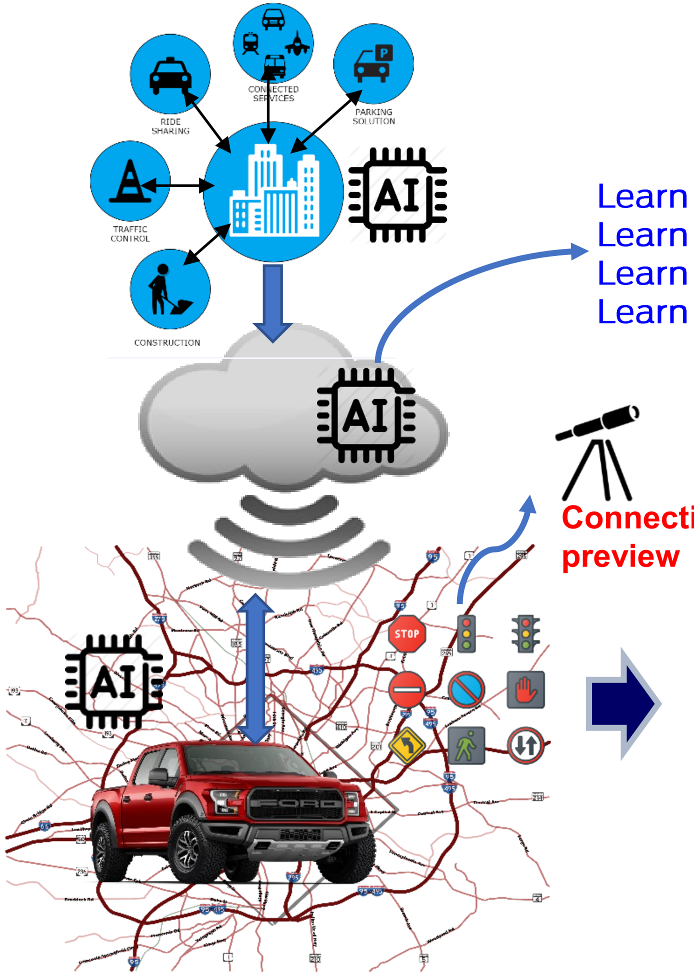

Learn traffic patterns
Learn route preferences
Learn Destinations
Learn drive quality between nodes
⋮

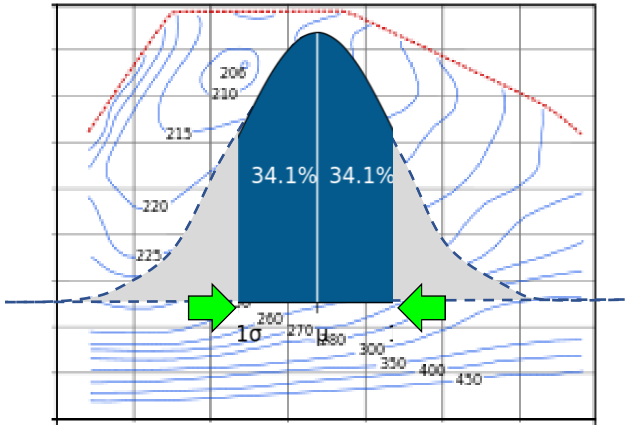

AI+Connectivity helps better manage uncertainty

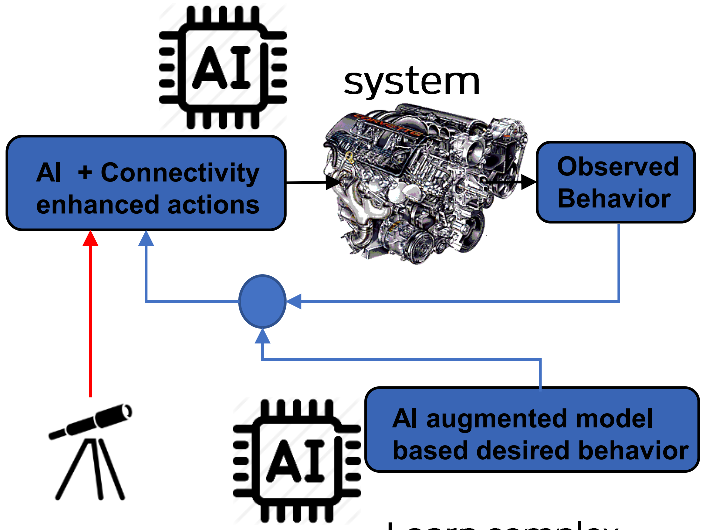

Learn complex system behavior/s

Connectivity with preview gives insight into stochastic events (traffic incidents).

Combination of AI and Connectivity allows better management of uncertainty with improved Risk Reward trade-off.

Demands/Pressures

- ❑ Fear of being left behind
- ❑ Industry must re-orient/adapt and choose carefully
- ❑ Industry must understand its data, adopt new tools, upskill and know where to apply ML.
- ❑ OTS solutions are typically NOT available for engineering problems → must create these
    - ❑ Drives the demand for Engineers and Scientists to be proficient in ML
- ❑ There are many sellers of “false AI dreams” … must carefully navigate between fact and fiction. Companies/Start-ups with multi-mil$ valuations are selling “stuff” that used to be commonplace. Hence, in house expertise is needed to separate fact from fiction.

High throughput research: targeted discovery vs chance

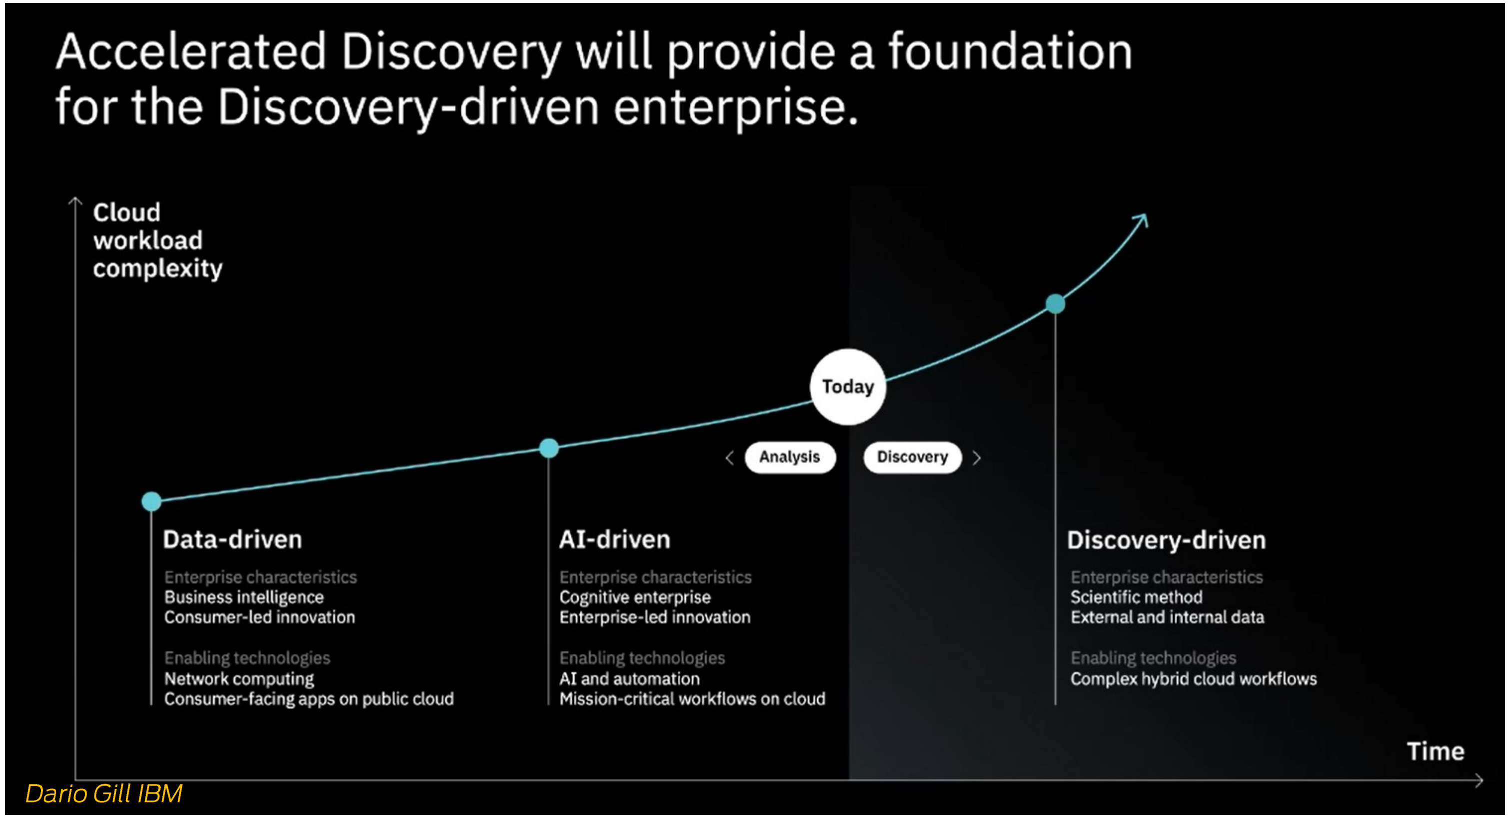

High dimensional, complex problem domains

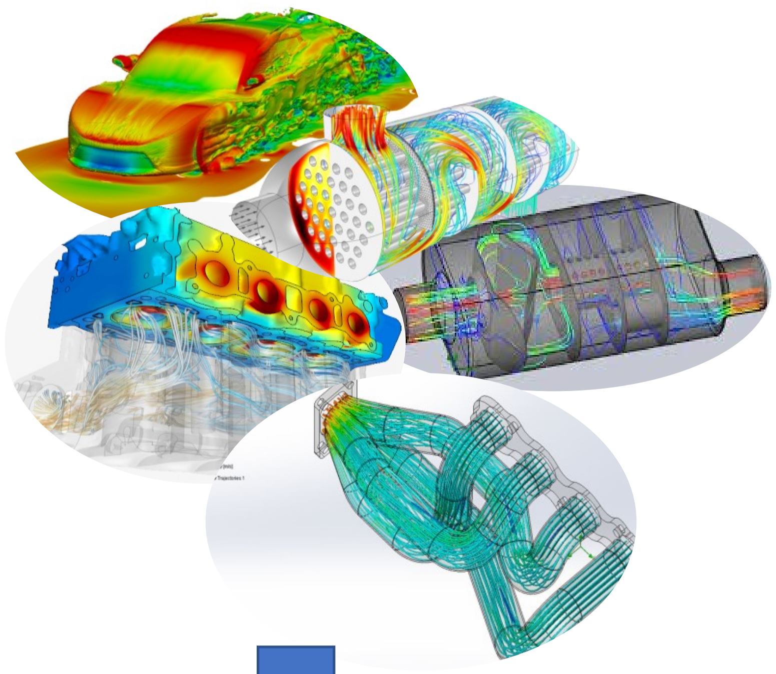

Sim time can take "days" for a single run per case
Curse of dimensionality → formally … computational cost increases with dimensionality of problem
Ex. Finite difference discretization with N points per dimension and d dimensions requires compute over N^d points

3-D PDE with 1000 points per dimension → $1000^3 \times 64$ bits ≈ 8GB RAM → doable on a high-end laptop
4-D PDE with 1000 points per dimension → $1000^4 \times 64$ bits ≈ 8TB Needs specialized nodes on HPC
5-D … needs new compute
Images … very high dimensional , MNIST 28x28 → 784 dimensional input

Standard solvers →Slow Simulations

PINNs → Less data, Interpretable, speed-up

Data Driven Models, Surrogates, ROMS, PINNS, Hybrids

High dimensional, complex problem domains: Li-Ion batteries

Negative Electrode

$$\epsilon_n \frac{\partial c_n}{\partial t} = \frac{\partial}{\partial x}[D_n \frac{\partial c_n}{\partial x}] + a_n(1-t_+)j_n$$

$$-\sigma_n \frac{\partial \Phi_{1,n}}{\partial x} - \kappa_n \frac{\partial \Phi_{2,n}}{\partial x} + \frac{2\kappa_n RT}{F}(1-t_+)\frac{\partial \ln c_n}{\partial x} = I(t)$$

$$\frac{\partial}{\partial x}\left[\sigma_n \frac{\partial \Phi_{1,n}}{\partial x}\right] = a_n F j_n$$

$$\frac{\partial c_n^s}{\partial t} = \frac{1}{r^2}\frac{\partial}{\partial r}\left[r^2 D_n^s \frac{\partial c_n^s}{\partial r}\right]$$

$$\rho_n C_{p,n}\frac{dT_n}{dt} = \frac{\partial}{\partial x}\left[\lambda_n \frac{\partial T_n}{\partial x}\right] + Q_{rxn,n} + Q_{rev,n} + Q_{ohm,n}$$

Separator

$$\epsilon_s \frac{\partial c_s}{\partial t} = \frac{\partial}{\partial x}[D_s \frac{\partial c_s}{\partial x}]$$

$$-\kappa_s \frac{\partial \Phi_{2,s}}{\partial x} + \frac{2\kappa_s RT}{F}(1-t_+)\frac{\partial \ln c_s}{\partial x} = I(t)$$

$$\rho_s C_{p,s}\frac{dT_s}{dt} = \frac{\partial}{\partial x}\left[\lambda_s \frac{\partial T_s}{\partial x}\right] + Q_{ohm,s}$$

Positive Electrode

$$\epsilon_p \frac{\partial c_p}{\partial t} = \frac{\partial}{\partial x}[D_p \frac{\partial c_p}{\partial x}] + a_p(1-t_+)j_p$$

$$-\sigma_p \frac{\partial \Phi_{1,p}}{\partial x} - \kappa_p \frac{\partial \Phi_{2,p}}{\partial x} + \frac{2\kappa_p RT}{F}(1-t_+)\frac{\partial \ln c_p}{\partial x} = I(t)$$

$$\frac{\partial}{\partial x}\left[\sigma_p \frac{\partial \Phi_{1,p}}{\partial x}\right] = a_p F j_p$$

$$\frac{\partial c_p^s}{\partial t} = \frac{1}{r^2}\frac{\partial}{\partial r}\left[r^2 D_p^s \frac{\partial c_p^s}{\partial r}\right]$$

$$\rho_p C_{p,p}\frac{dT_p}{dt} = \frac{\partial}{\partial x}\left[\lambda_p \frac{\partial T_p}{\partial x}\right] + Q_{rxn,p} + Q_{rev,p} + Q_{ohm,p}$$

D: diffusion coefficient
t_+: transference coefficient
λ: thermal conductivity
j: pore wall flux
Q: source terms

Length scales:
X-direction: 100 μm
Y-direction: 5 cm
500-1 aspect ratio

Particle radius: 2-5 μm

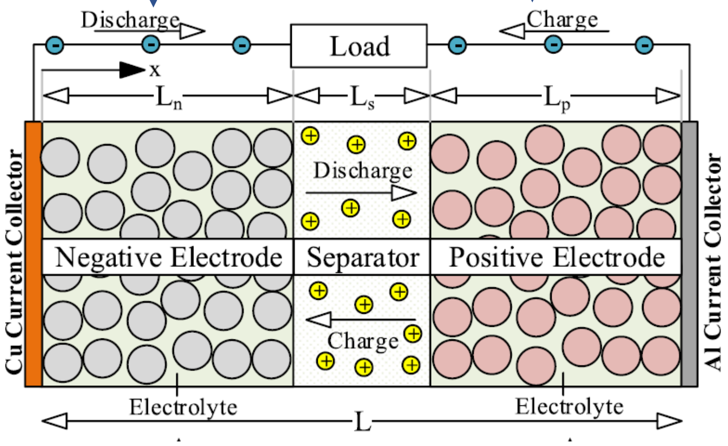

Fick's law of diffusion for spherical particles
conservation laws
Kirchhoff's and Ohm's
Butler-Volmer kinetics
Fourier's heat equation

High dimensional, complex problem domains

- Just to note, there are many such problems across several domains in Science and Engineering.
- One fact that does stand out, and we will talk about this again, is the impact of these complexities on purely data driven re-constructions.
- For systems with high degrees of Nonlinearity, a data driven reconstruction of the system dynamic will typically require dense sampling across the domain of nonlinearity requiring *apriori* knowledge of the nonlinearities.
- Often such data will NOT be available, forcing a sparse representation and the use of techniques such as:
    - Compressed sensing
    - Bandwidth extension
    - and/or incremental learning from failures

Automation → Control → Autonomy

Computer Science construct

Neural Control

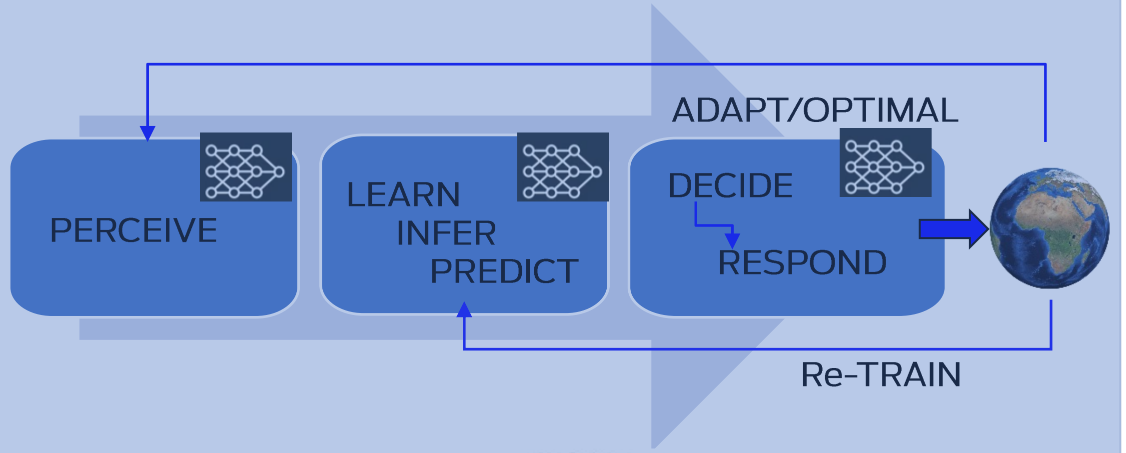

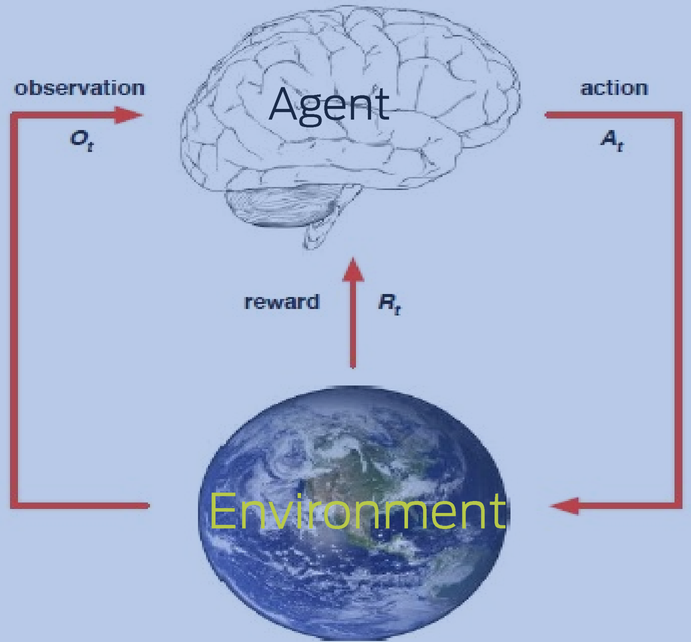

Optimal Control Construct

Modern Control

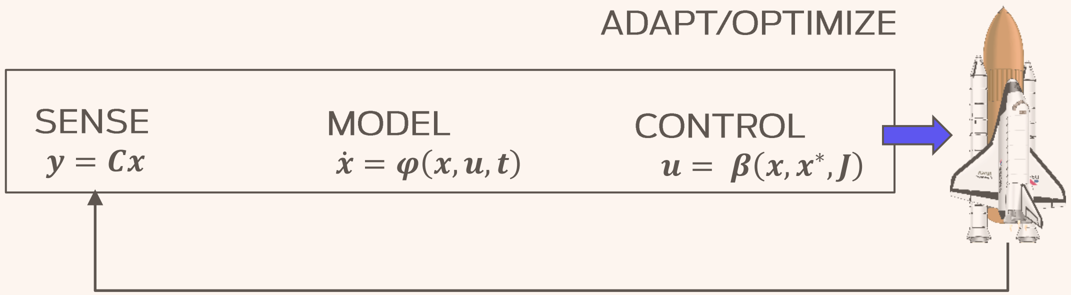

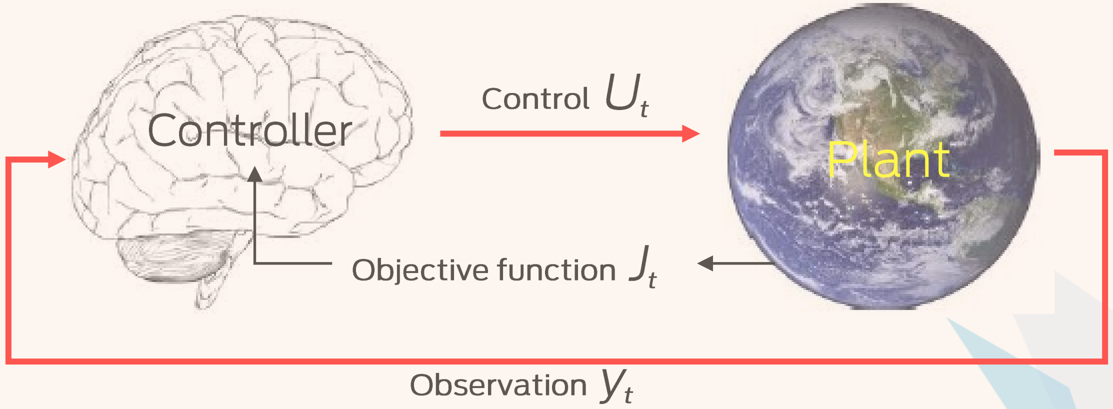

$$\mathcal{R}^{Env} \gg \mathcal{R}^{Plant}$$

$$G^{MC} \gg G^{NC}$$

Formal methods exist for MC, similar methods are being explored (are needed) for NNC

images adapted from D. Silver notes

Autonomy, Mobility, Robotics

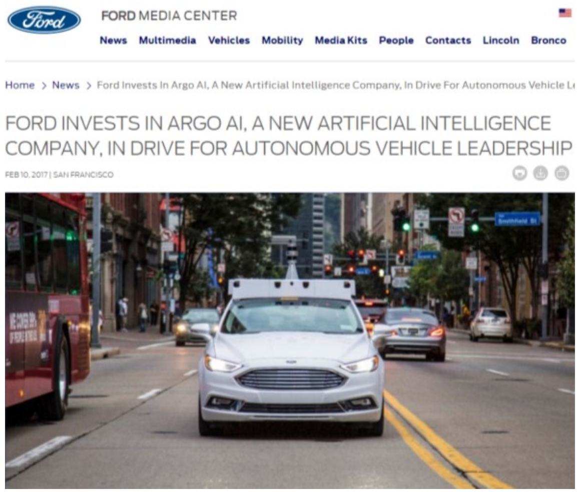

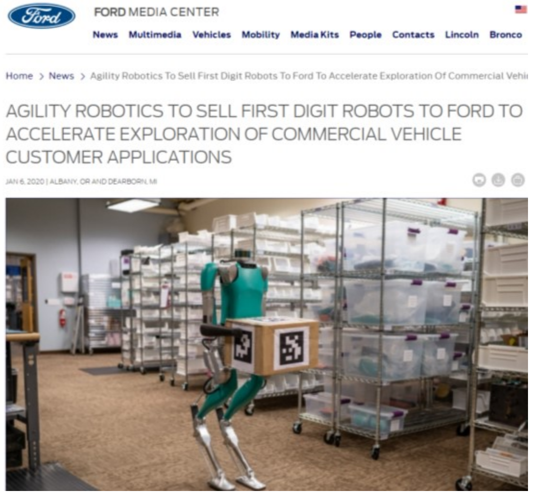

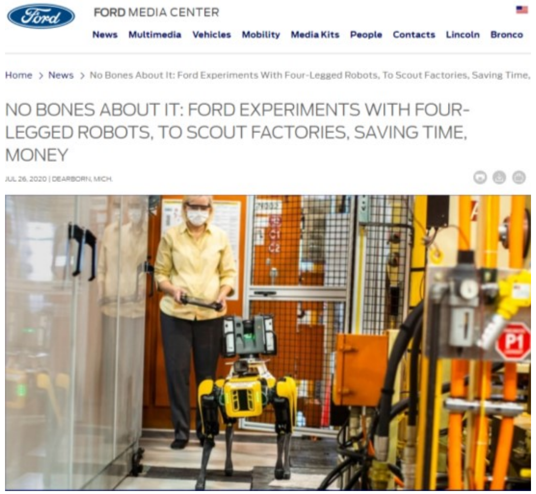

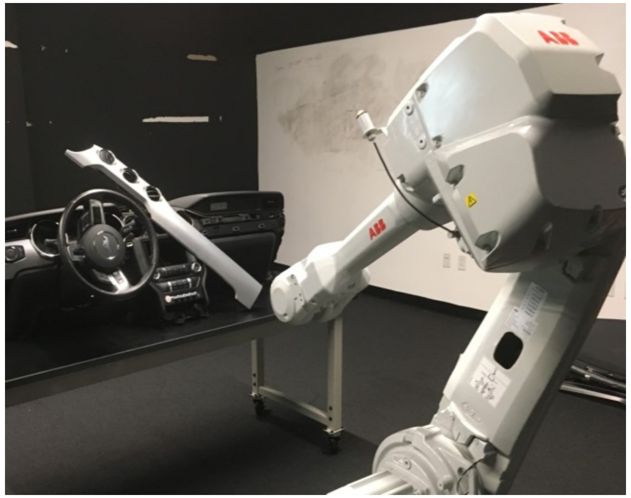

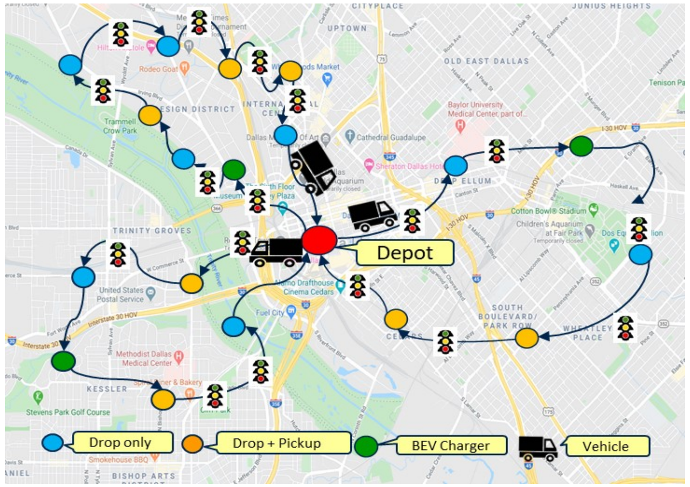

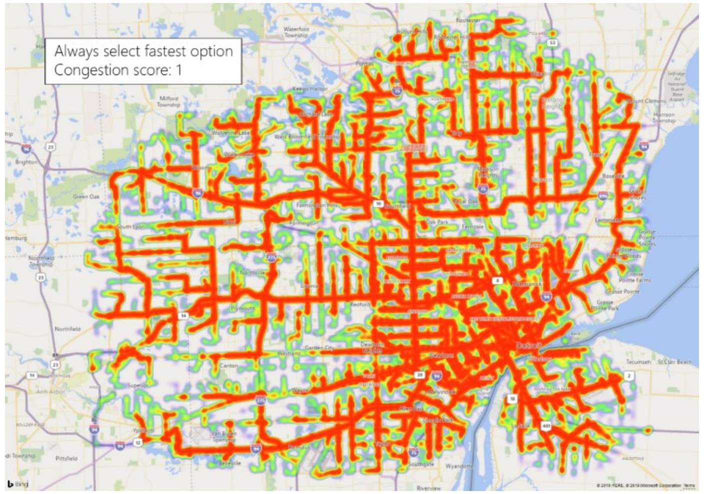

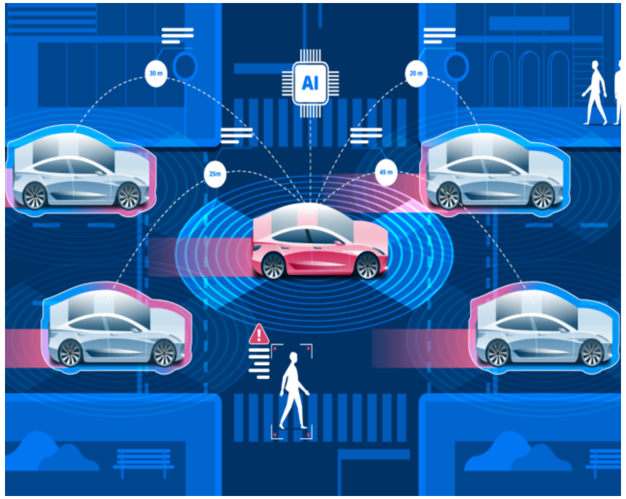

Monitoring: shifting to a distributed paradigm

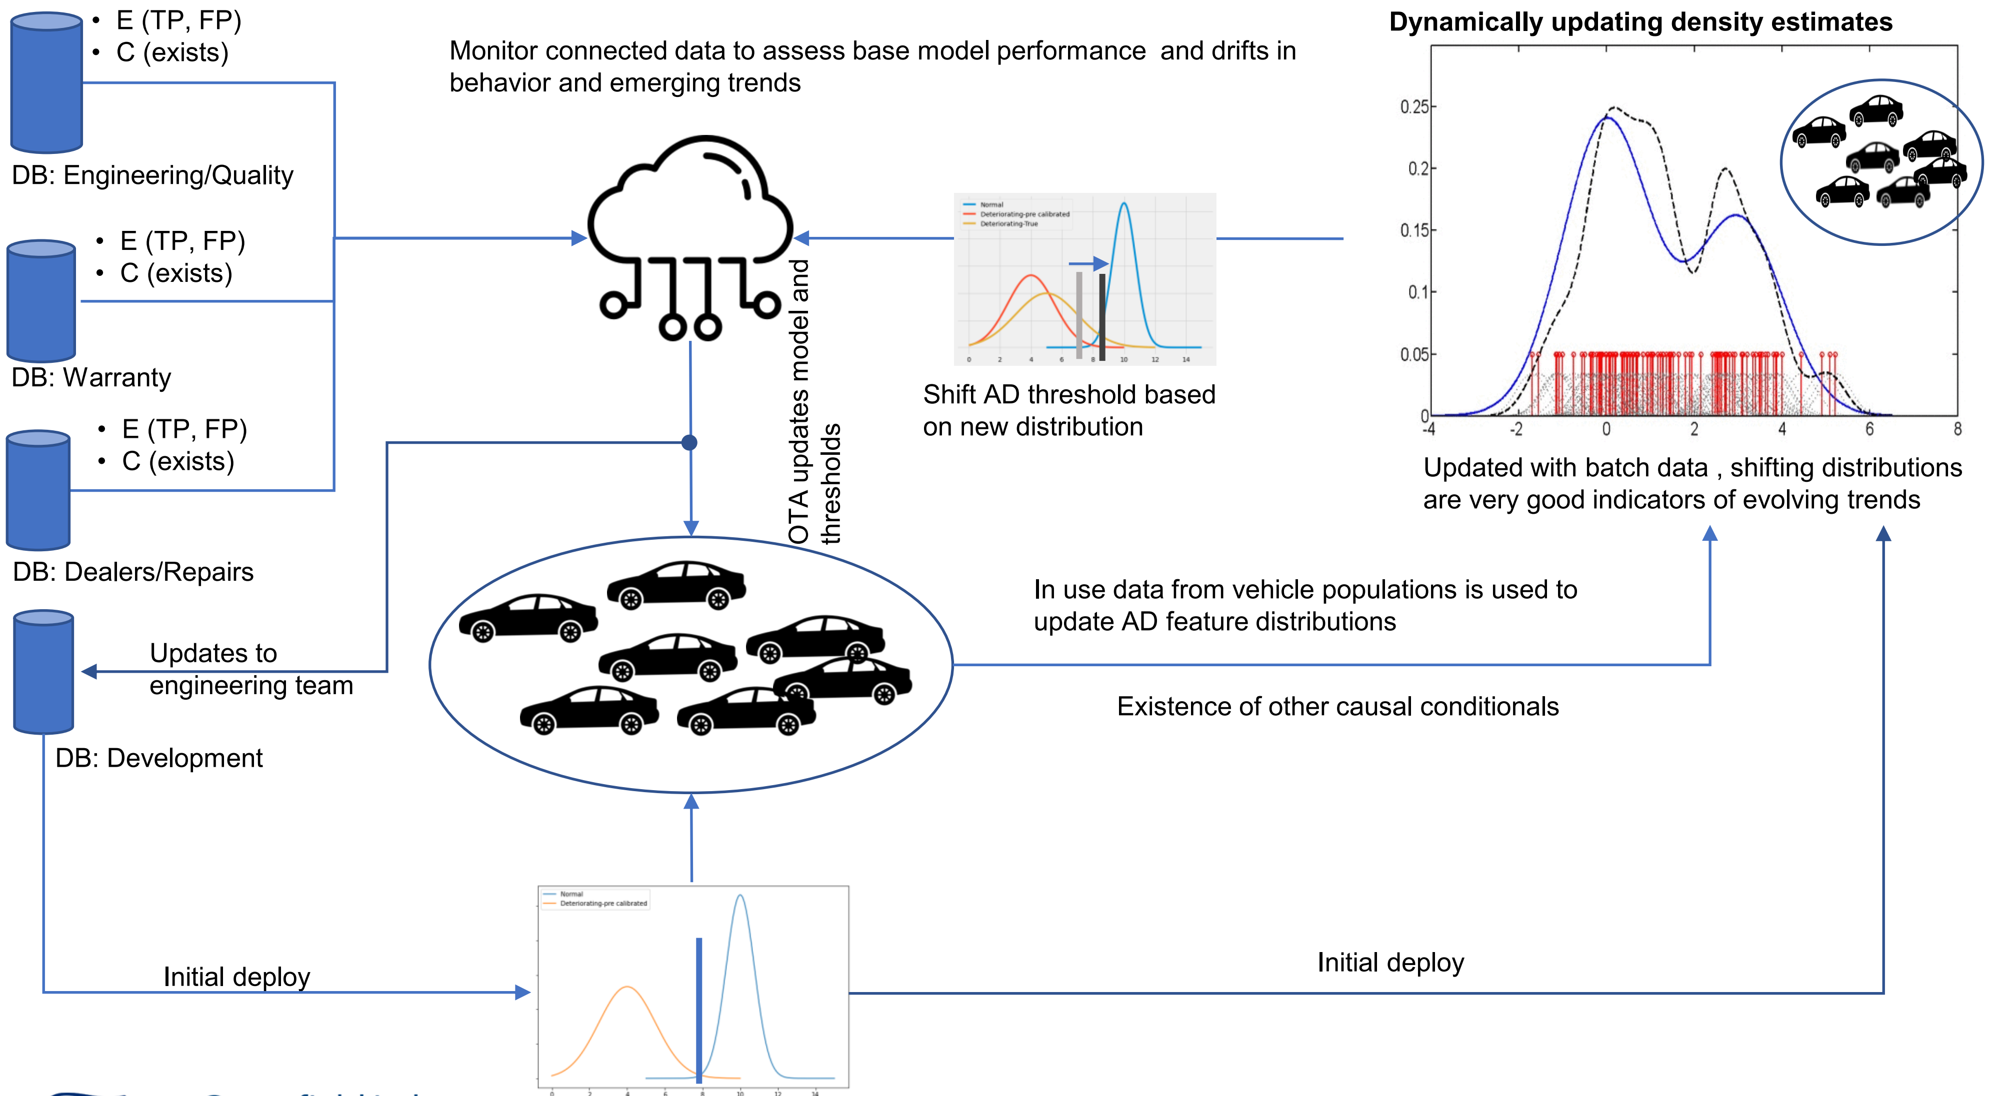

$$C = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix} \longrightarrow \text{cause}$$

$$E \longrightarrow \text{effect}$$

Case studies … some examples

- ❑ Powertrain– Emissions Control
- ❑ Collision -- Vision systems and challenges
- ❑ Manufacturing – process control , Robotics
- ❑ Material Science – using images
- ❑ Charactering Sprays with Images
- ❑ Models (PINNS and Surrogates)
- ❑ Monitoring (OBD, vs AD&RUL)

Diesel Particulate filter – process and constraints

- Diesel vehicles must trap soot in a filter → DPF.
- Filters have finite trapping capacity and must be periodically cleaned or "regenerated".
- Active regeneration involves a targeted burning of the trapped soot by intrusively elevating and maintaining the exhaust gas temperature in the range of 600°C - 700°C.

Formulate as an optimization problem, *SAE-2019-01-0316*

$$J = \delta \times (J_{Heat} + J_{Burn}) + (1-\delta) \times \mathcal{R} \times \frac{(m_{soot} - m^*_{soot})}{m_{soot\_max}} + (1-\delta) \times J_{Load} \times \frac{(m_{soot} - m^*_{soot})}{m_{soot\_max}}$$

Where:

$$\begin{cases} \boldsymbol{\delta = 1, for, m_{soot} \geq m^*_{soot}} \\ \boldsymbol{\delta = 0, \qquad otherwise} \end{cases}$$

$\mathcal{R}(m_{soot}, m^*_{soot})$ is the risk of not regenerating in current drive. $m^*_{soot} = f(\mathcal{D})$is the optimal soot load as a function o the drive $\mathcal{D}$

Soot Management and Penalties

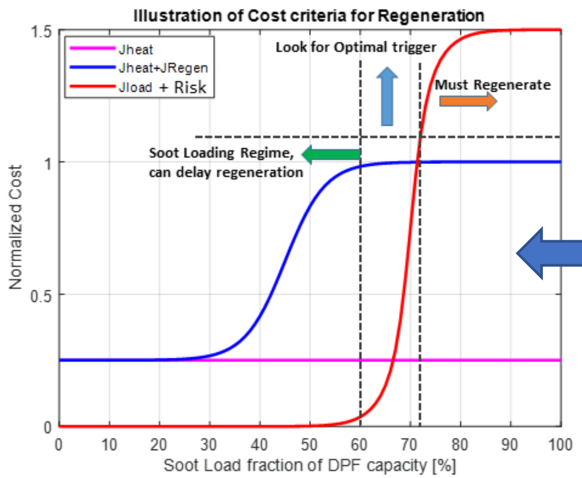

- Regen Fuel Penalty (RFP) → f(Regen Freq, Rgen duration).
- Regenerations also incur NO_x penalty.
- Traffic conditions and drive quality impact Regen Quality (RQ).
- Operating conditions during an active regen may become non-optimal resulting in aborts.
  - **Frequent aborts increase regeneration frequency.**
  - **Frequent regeneration is not optimal**

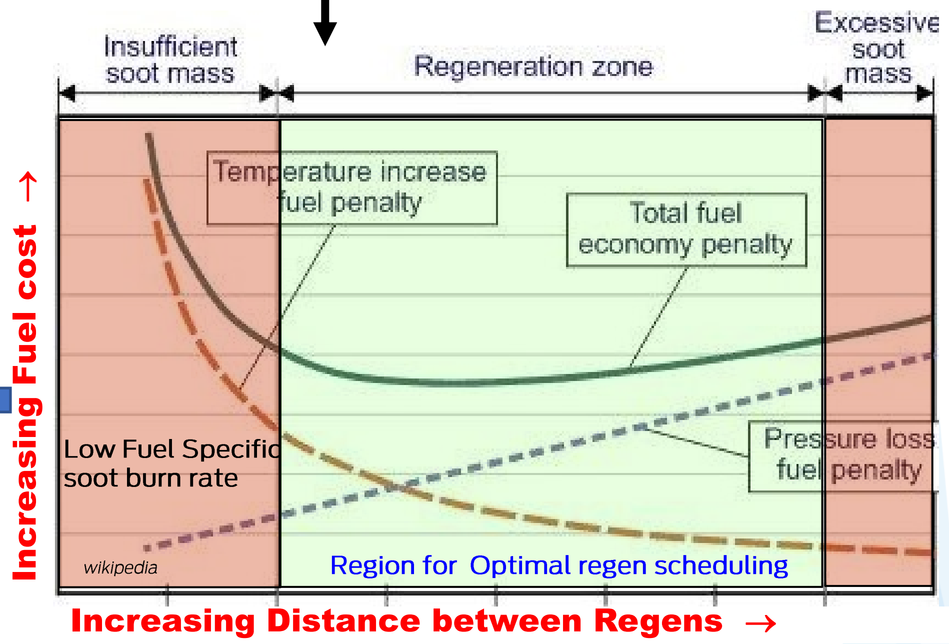

Naïve approach leads to short cycle regens

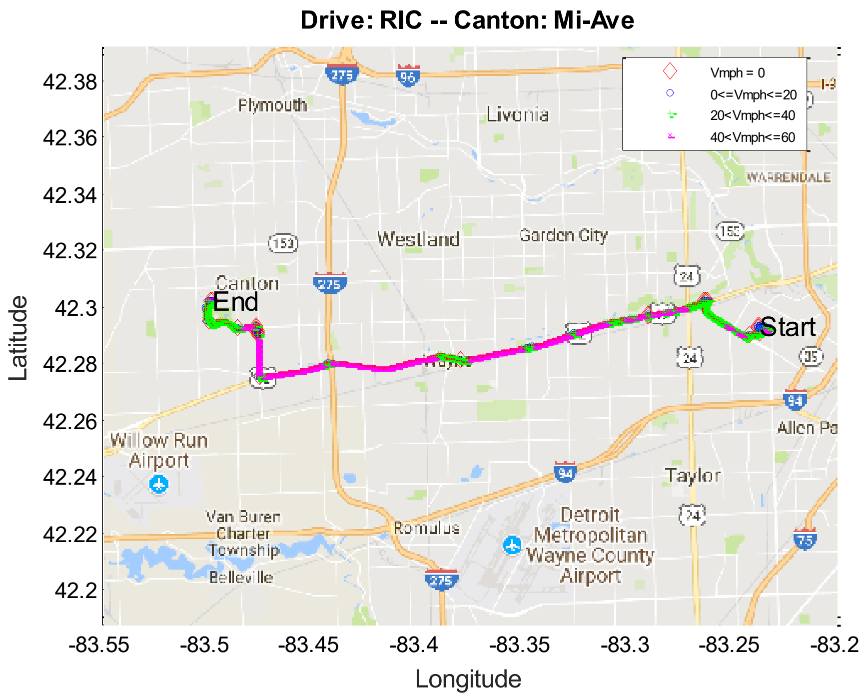

<u>Fixed SL- Threshold</u> based regen decisions can lead to short cycle regenerations resulting in increased regeneration frequency.

The probability of getting a deep regen (≥ 70%) is a strong function of the route and the likelihood of maintaining optimal regen conditions over the driven route.

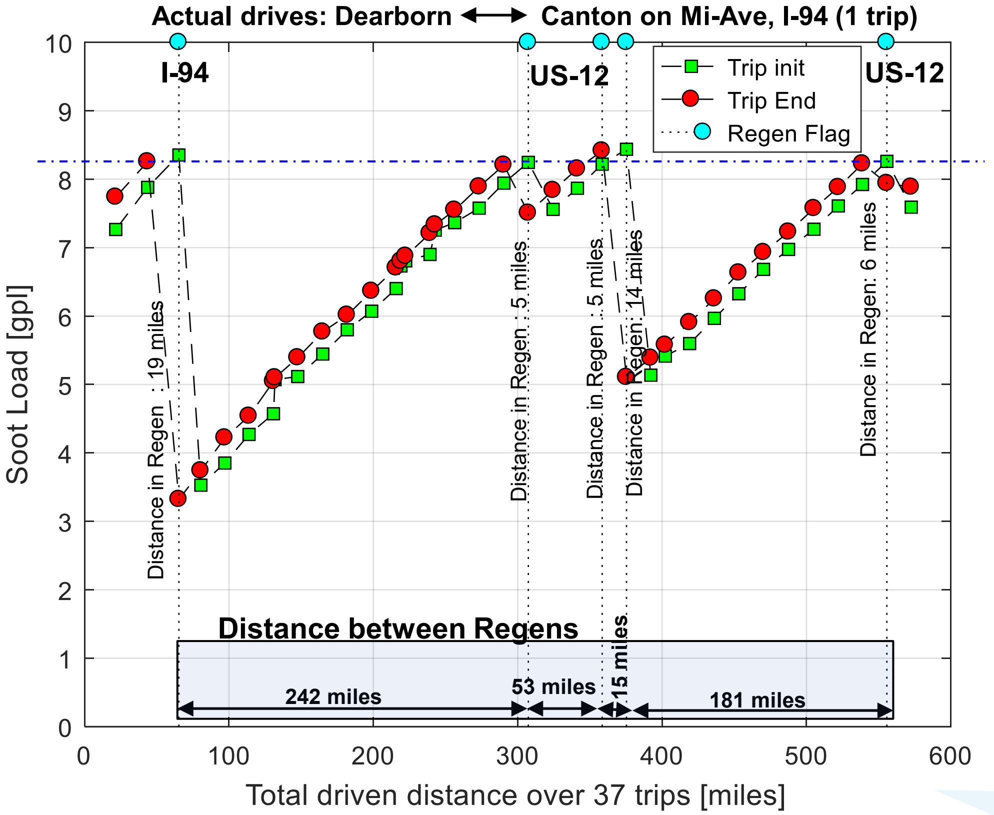

Soot Load (SL) at the start and end of each trip are shown

SL_{end} > SL_{start} → loading cycle

SL_{end} ≤ SL_{start} → Regen cycle. (SL_{start}- SL_{end}) → Soot load depleted

Drive information is an enabler toward more efficient regenerations

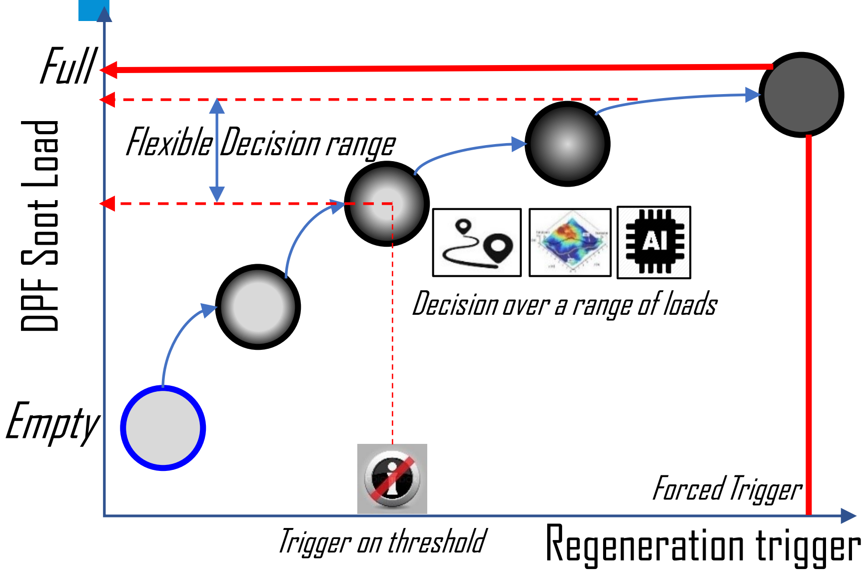

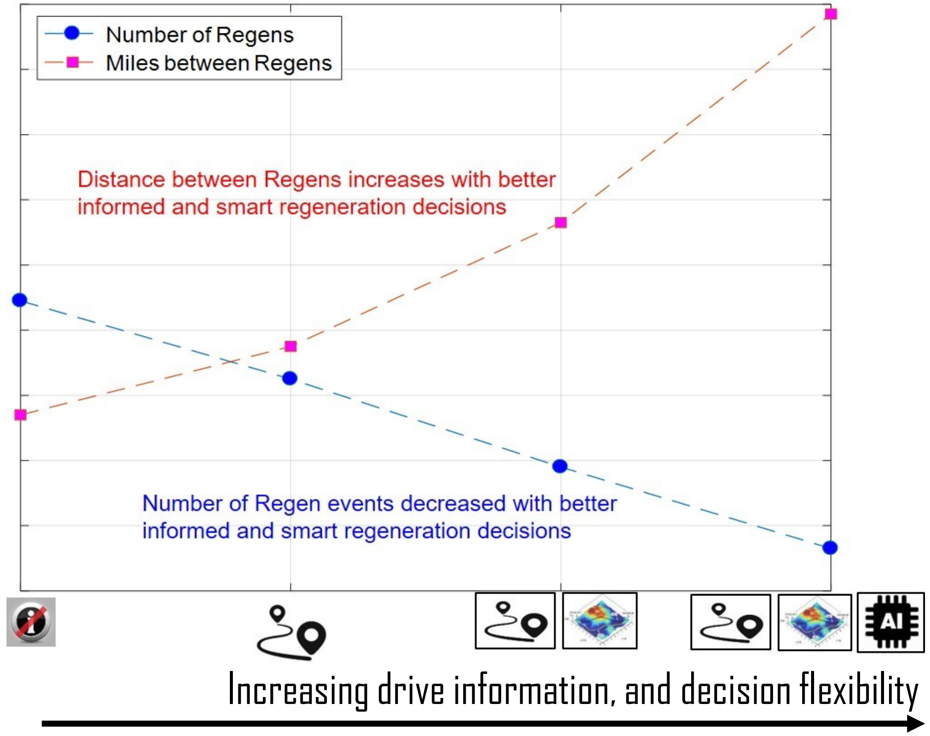

Conventional

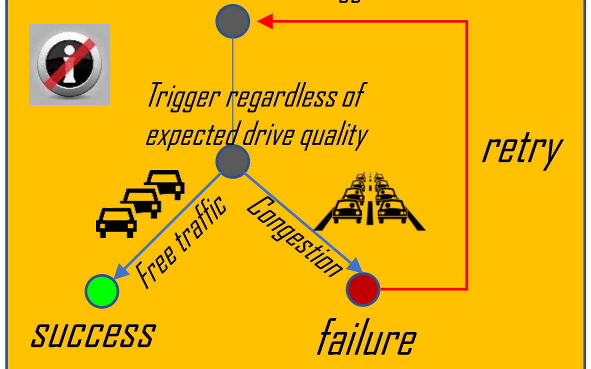

$P(success|traffic, drive, \dots)$
Forces a risk averse policy that still does not guarantee efficiency

Connected and Smart

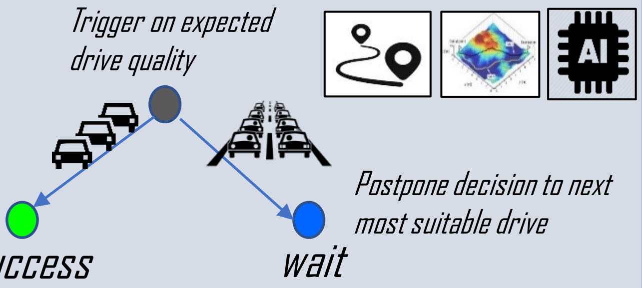

Information allows flexibility in decision from a better understanding of Risk-Reward trade-off

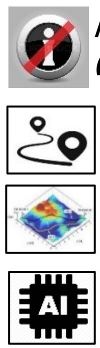

No information forces threshold based (blind) decisions

Destination is known

Route with traffic known

Smart decisions (Gaming, Reinforcement) *with A. Rahimpour*

Vision based Collision avoidance with large animals.

Collisions with large animals such as deer and moose, etc. lead to about 200 human deaths and $1.1 billion in property damage every year.

Backbone	mAP @ IoU=0.50	Train time (s/iter)	Test time (s/iter)	Memory footprint (GB)	# parameters	Model size
ResNet50-FPN	0.99	0.1897	0.0467 (25FPS)	5.2	25 million	98 Mb
Mobilenet v2	0.79	0.1419	0.0269 (50FPS)	3.5	3.5 million	14 Mb

Preprocessing → Segmentation → ROI selection → CNN - Feature Extraction → Classification and Regression

Tracking

Track between frames
Smoothing

Risk/intention prediction

Orientation
Pose
Trajectory
Motion

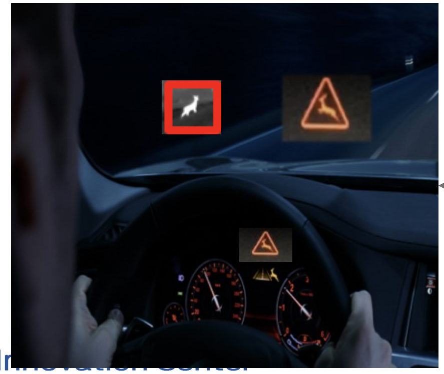

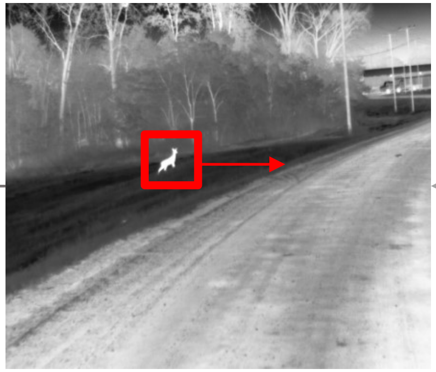

Predicting intent:

- Human intent is difficult
- Animal is even trickier

Manufacturing – Process control , Robotics, Error-proofing

Assembly, Precision tasks, Inspection

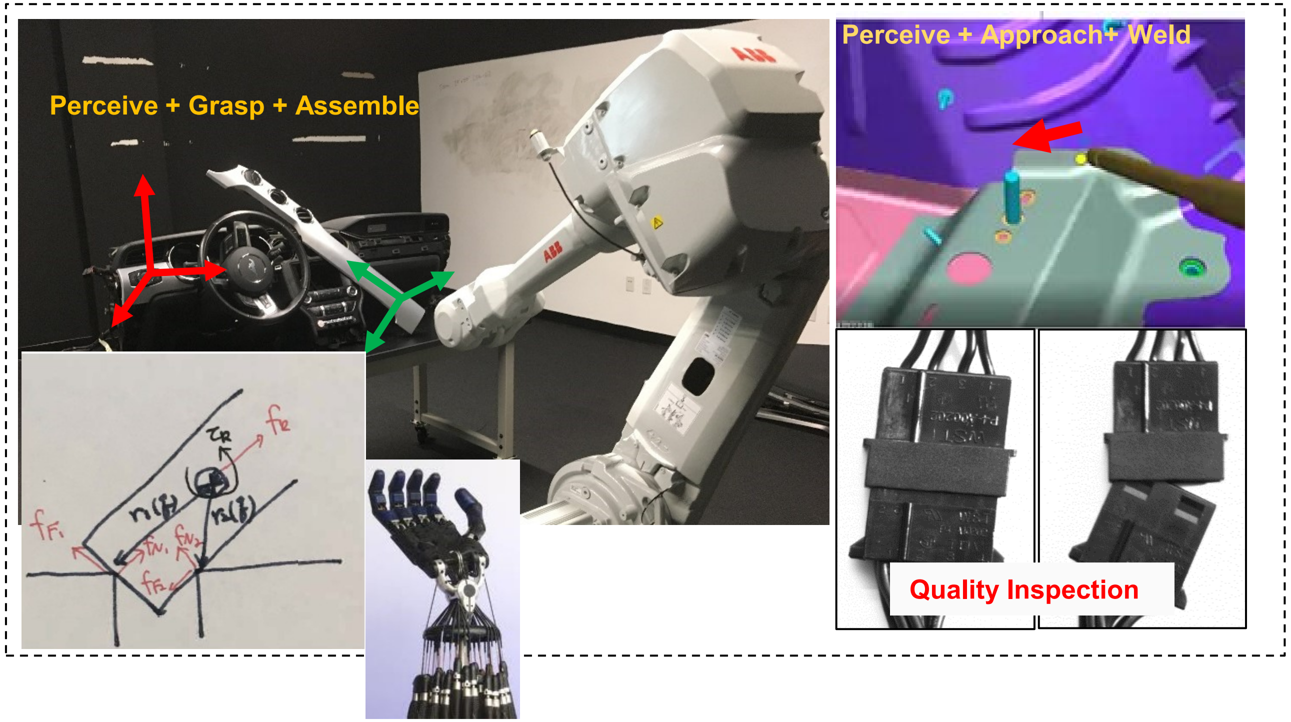

- Robots can assemble with repeated precision
- Grasping and manipulation tasks remain challenging
- Current end effectors are under-actuated systems
- Current robotics problems deal with parts designed for human assembly.

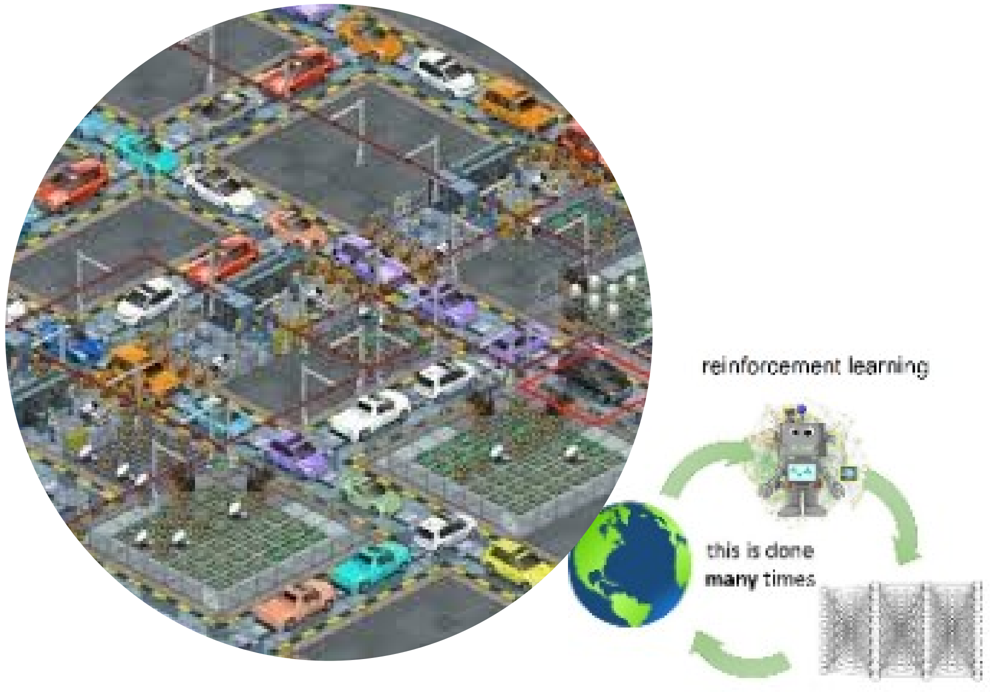

Assembly line process .. Pallett routing

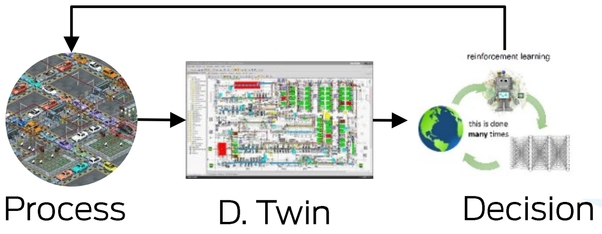

Material Science – using images

- An image editing generative adversarial network, AttGAN, was implemented to generate realistic synthetic SEM images showing subtle effect of the varied synthesis parameter of graphene additive %
- AttGAN avoids the large data requirement in GAN based design.
- Domain scientists can visualize and understand how changes in material attributes, synthesis/ processing conditions impact the sample microstructure features which in turn determine its properties and thus can further be optimized
- Needs fast imaging, SEM/FIB images, 2D vs 3D , defining structure level information is tricky
- Need to include physics based constraints

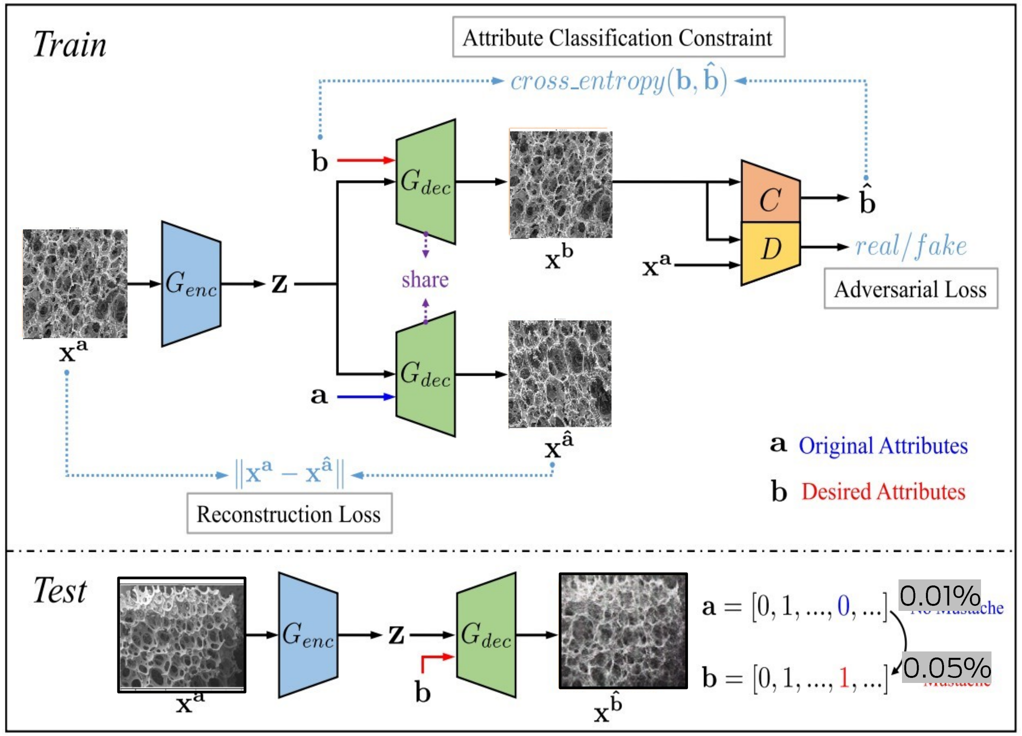

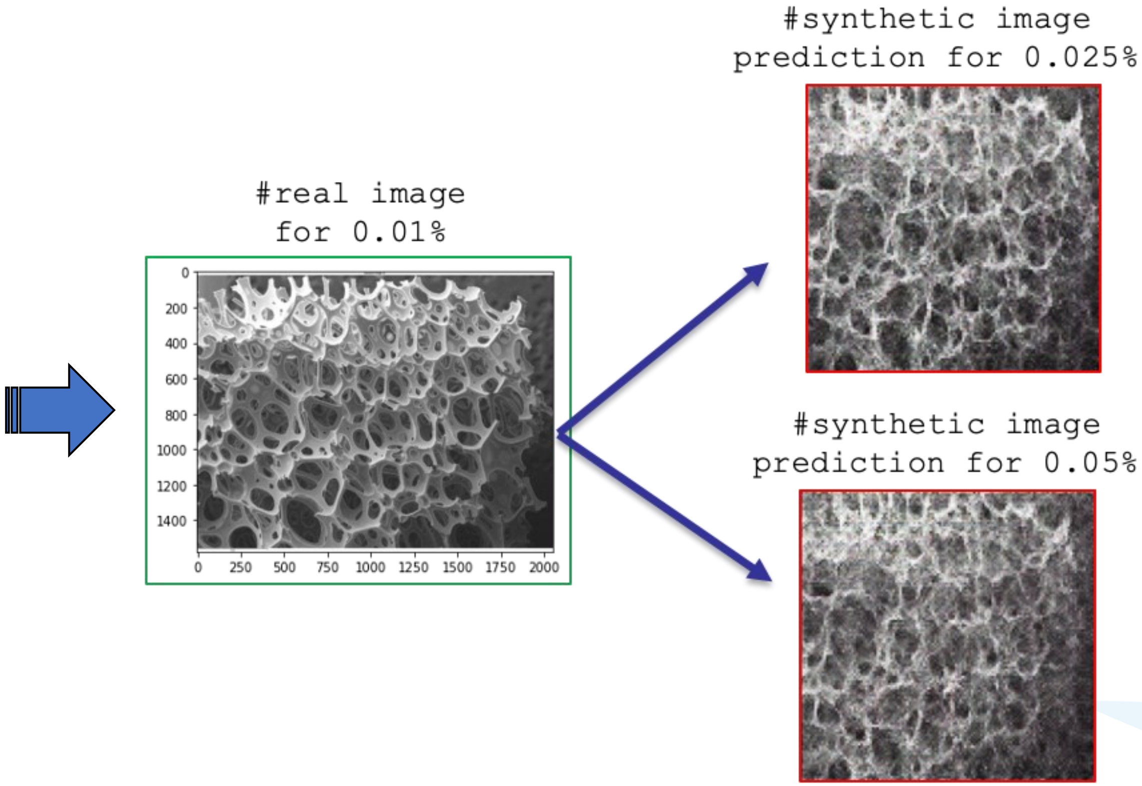

Discussions on going with LLNL on extending these efforts *with Hongjiang Li, et.al*

Charactering Sprays with Images

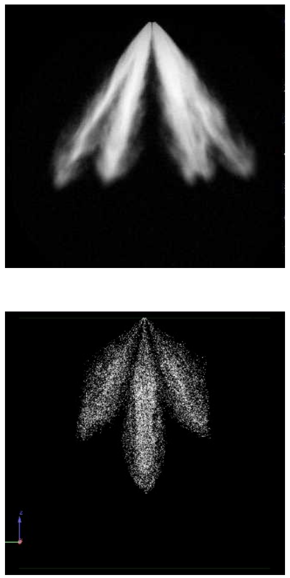

Class A: No collapse

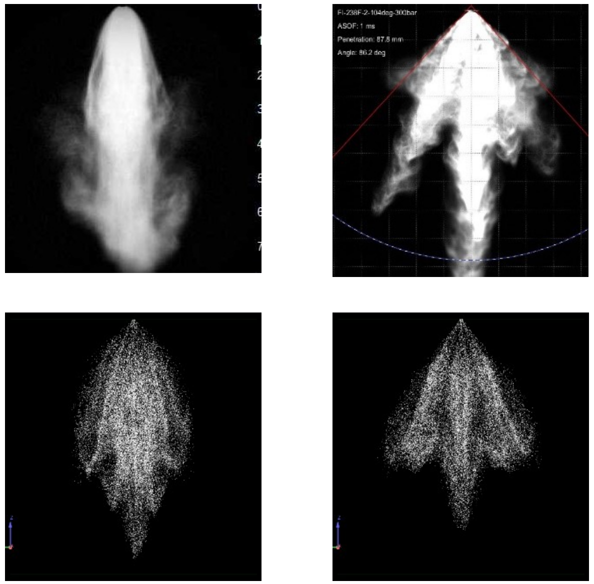

Class B: Spray collapse, and transition

<u>Objective:</u> Injector nozzle design and

<u>Normally → visual + CFD sims</u>

- Visual process too slow to keep up with fast imaging
- Prone to bias →based on SME experience

<u>New → use CNN's</u>

- Major speed up in analysis
- Consistent and highly accurate and quantifiable results
- Transfer to other sprays/injectors
- Can further map to in-cylinder combustion performance measures and integrate into inverse design

with C. Rackauckas, A. Edelman -(MIT), K. Balakrishnan (Ford)

Models (PINNS and Surrogates)

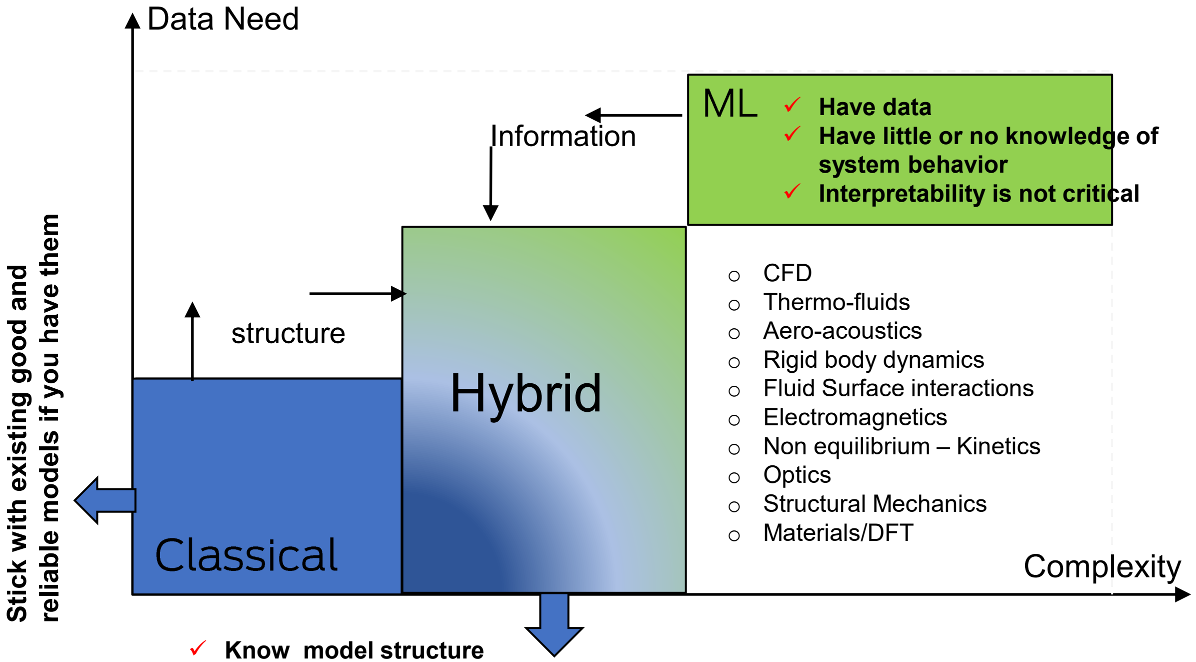

- ✓ **Know model structure**
- ✓ **Models are complex or computationally intensive**
- ✓ **Have data: experimental + Simulations**
- ✓ **Interpretability is essential**

Use fundamental properties:

- Universal approximation property of NNets →given enough layers a NN model can approximate any nonlinear function to within ϵ → good when nonlinearities are not known *apriori*.
  - $y_{ML}(x) = \sigma_3\left(W_3 * \sigma_2\left(W_2 * \sigma_1(W_1 x)\right)\right)$ → eg. a 3 layer network, NN's are function approximations: $\mathcal{R}^n \rightarrow \mathcal{R}^m$
- Differential equations are compact ways of specifying arbitrary nonlinear transforms by mathematically encoding prior structural assumptions → good when nonlinearities are known *apriori* → $\dot{x} = \varphi(x, u, t)$
- Both are differentiable.

Koopman dynamics to learn time evolution

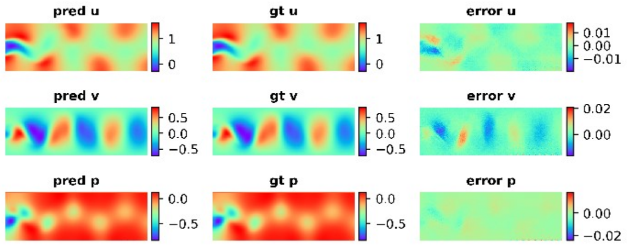

Include domain knowledge in ML : Scientific ML

Neural O/P-DE's

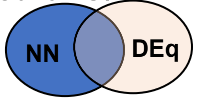

1. NN architectures with activations as DiffEqs
2. NN architectures with some layers as ODE's
3. ODE's are defined with NN embeddings
4. Cost functions on ODE's can be NNets

- *Do not know RHS*
  $\dot{x} = \varphi(x, u, t)? = f_{ML}(x, u, t, \theta)$
- *Know some structure, can decouple*
  $\dot{x}_1 = f_{ML}(x_1)$ ... unknown ... ML solution
  $\dot{x}_2 = \varphi_{2(x_1, x_2)}$ ... known
- *Impute missing term*
  $\dot{x} = \varphi(x) + f_{ML}(x)$
- *Stochastic Diff Eqs*
  $dx = \mu(x, \theta, t)dt + \sigma(x, \theta, t)dW_t$
- *SDE's with jump discontinuity*
  Filippov: $\dot{x} = sign(t)$
- *Partial Diff Eq classes*

Navier-Stokes is HARD.... solutions are attempted by quasi-linearizing the convection term as:
$(u \cdot \nabla)u \approx \psi(u, p, t)\nabla u$
Instead approximate "ψ" by NNet ...

$$\overbrace{\underbrace{\frac{\partial \mathbf{u}}{\partial t}}_{\text{Variation}} + \underbrace{(\mathbf{u}\cdot\nabla)\mathbf{u}}_{\text{Convection}}}^{\text{Inertia (per volume)}} \overbrace{- \underbrace{\nu\nabla^2\mathbf{u}}_{\text{Diffusion}} = \underbrace{-\nabla w}_{\text{Internal source}}}^{\text{Divergence of stress}} + \underbrace{\mathbf{g}}_{\text{External source}}$$

NeuralPDE: Automating Physics-Informed Neural Networks (PINNs) with Error Approximations, to appear in arXiv

Anomaly Detection: Use ML for multivariate problems

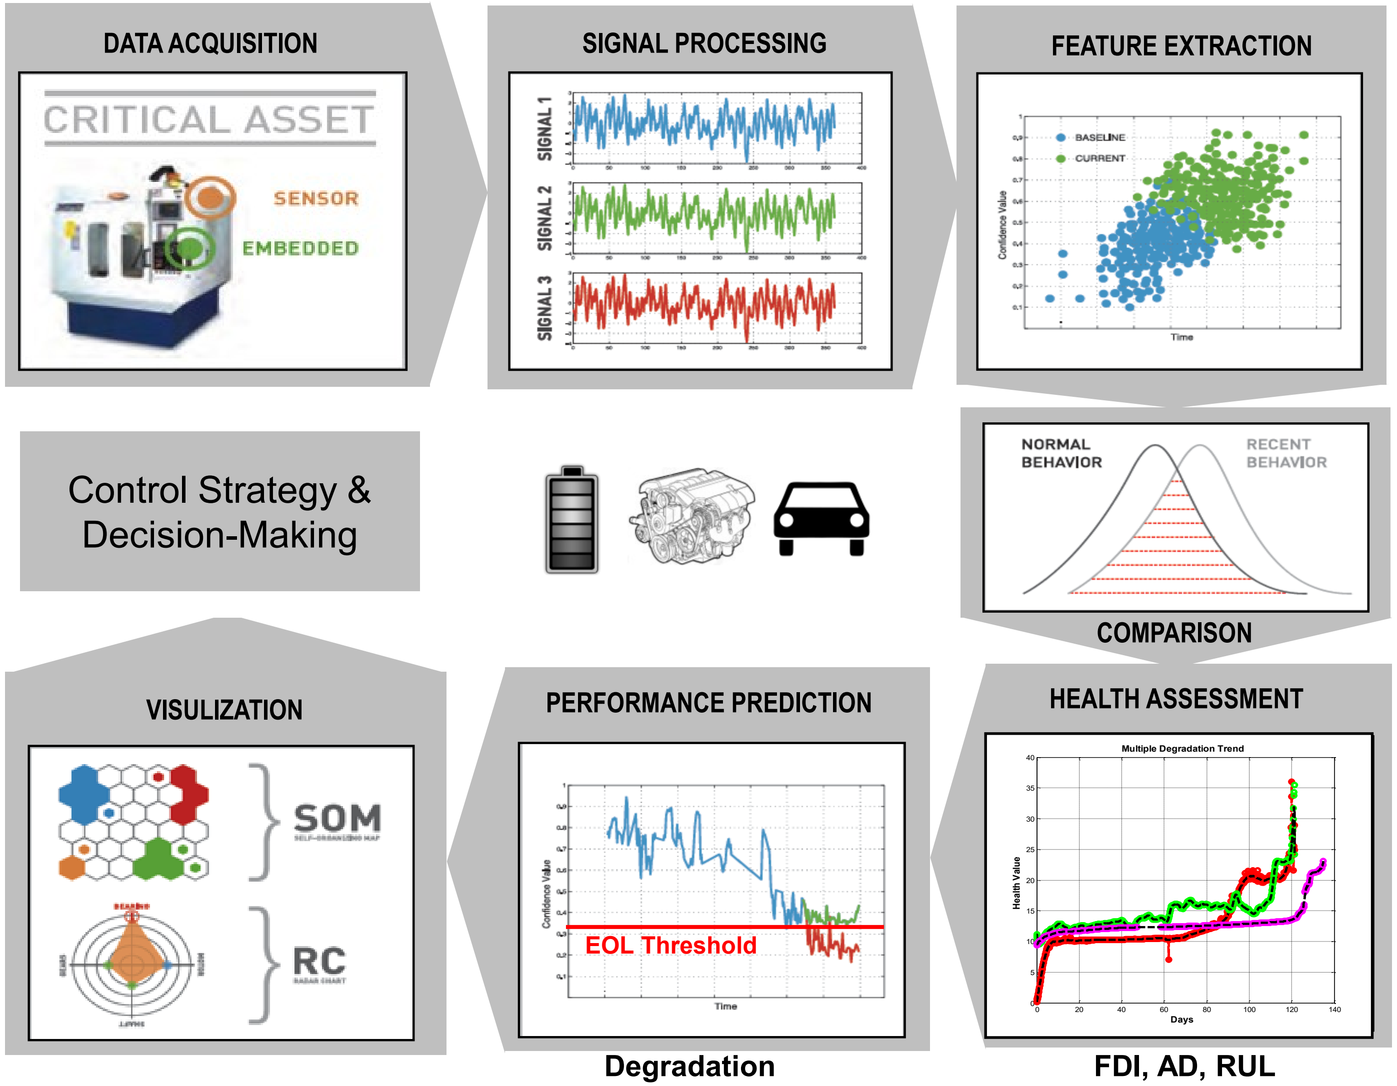

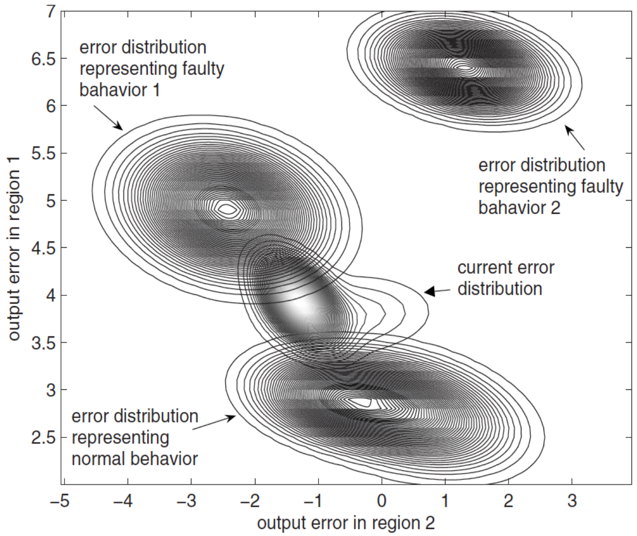

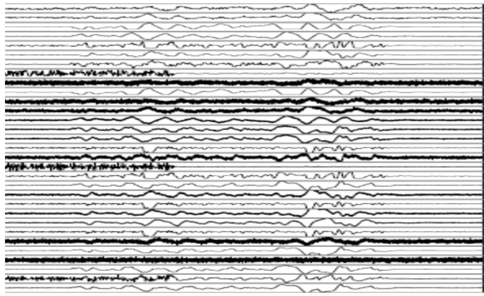

- Parts tested at production may not provide adequate indications of impending in-use failure
- Defective parts in an assembly may go un-noticed
- Can we isolate these “problem child's” at birth ?
- Parts will degrade differently from usage variations. Several populations with varying degradation patterns will evolve.
- Faulty data is sparse.
- What is the scope of domain transfer between problems?
- Do contrastive learning approaches work beyond simple examples.

On-Board Diagnostics (OBD)with future rule making

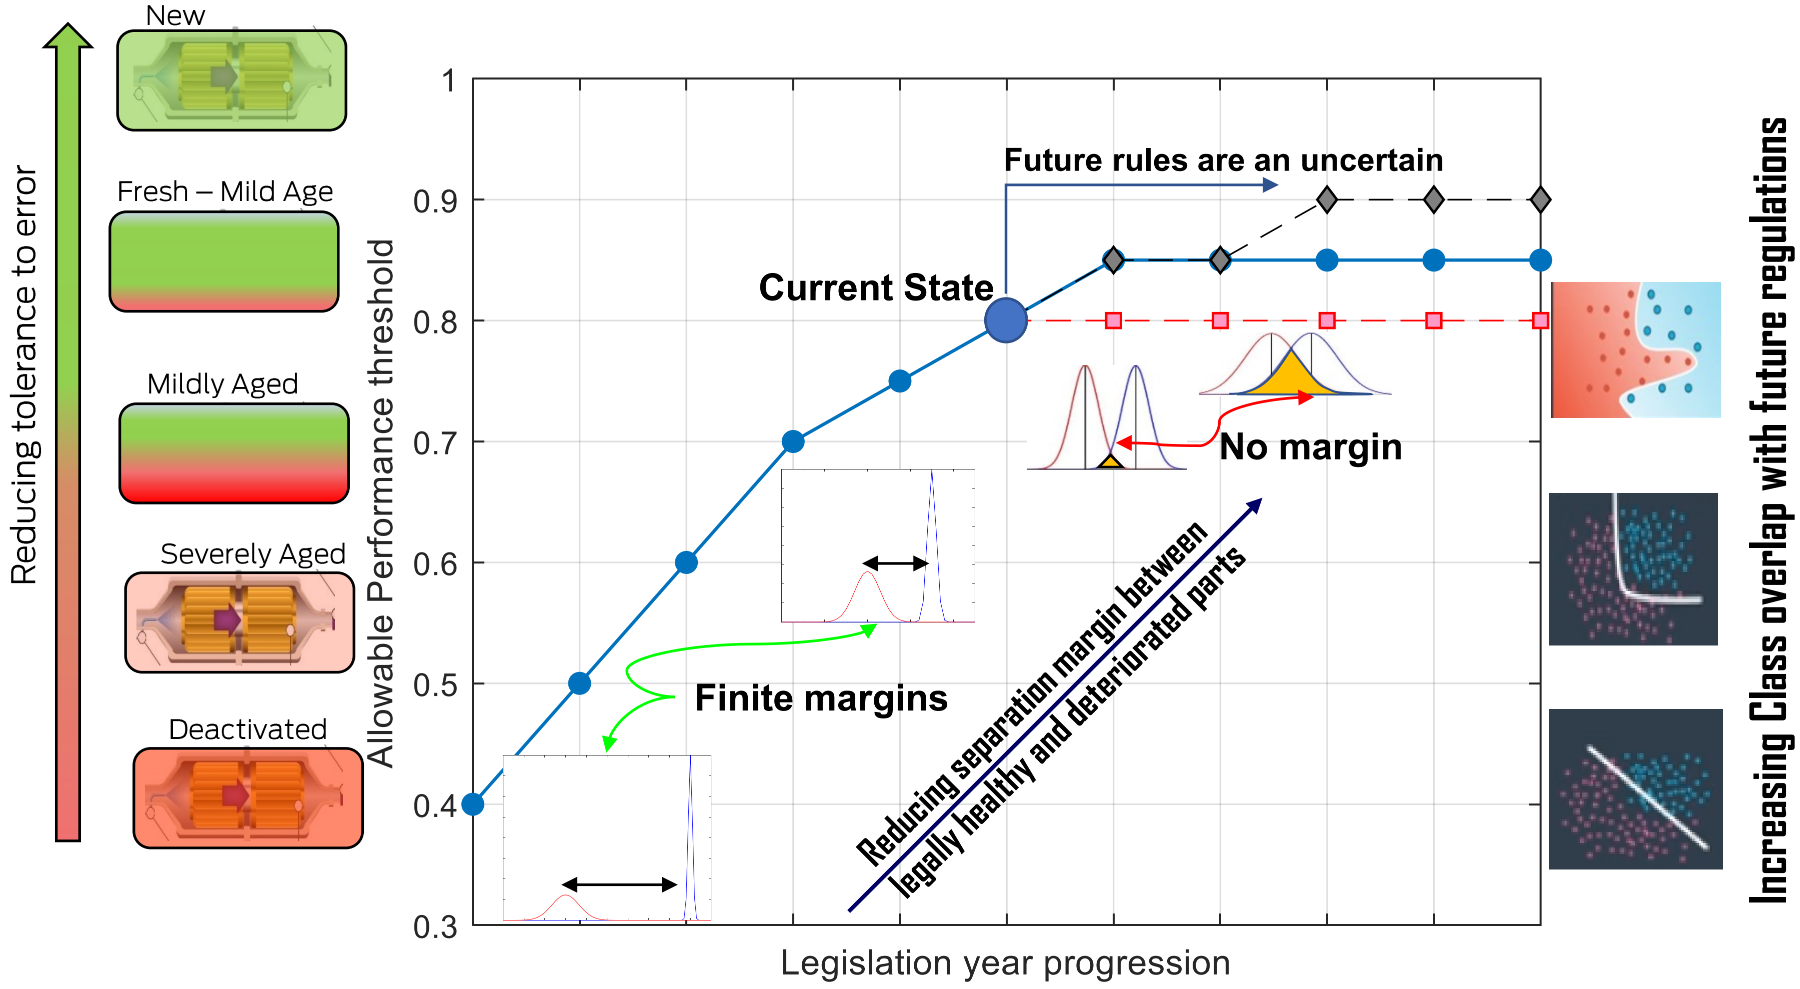

Isolating failed parts is easy (most of ML based AD today), isolating functional parts with marginal deterioration and multiple dependencies … not so easy

Outlook, Opportunities and Challenges

Data

- How much data and how to get it optimally
- The data trap: “too much data with very little information”

Robustness

- Data drift/bias/coverage
- Models drift/bias
- Safety– in RL (exploration), neural controllers, controllers with perception

Causality

- Generalization /Explainability
- UQ (of Data and Model)
- Embedding physics → PINNs, Hybrid methods

Cautions

- Regulations , Safety and Guidance
- Deployment, CI-CD is now CI-CD-CM
- Fact Vs Fiction
- The “AI-Hammer” effect

Data: How much ? Costs? Manage, Label, Drift

OpenAI disbands its robotics research team
VB July 2021

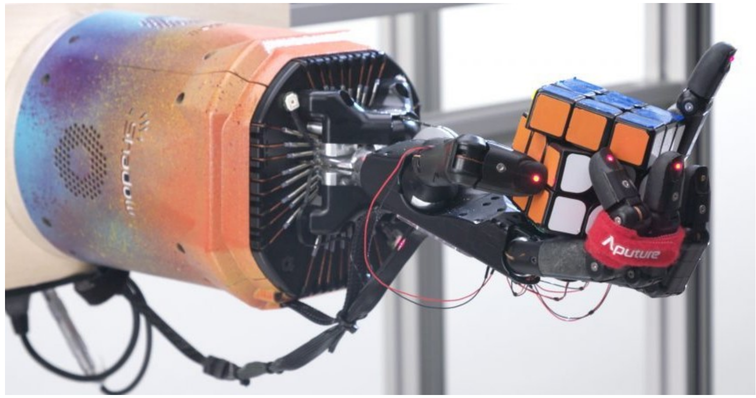

OpenAI has disbanded its robotics team after years of research into machines that can learn to perform tasks like solving a Rubik's Cube. Company cofounder Wojciech Zaremba quietly revealed on a podcast hosted by startup Weights & Biases that OpenAI has shifted its focus to other domains, where data is more readily available.

"So it turns out that we can make a gigantic progress whenever we have access to data".

Slightly misleading its not just data but LARGE DATA

- Is data a good substitute for Physics ? Ignoring physics implies data must re-learn the system → Increased data burden and all associated risks/errors.
- Extracting full domain representation from data alone requires a data set with very dense representation → which usually translates to aggregation of a lot of data with sparse representations
- Data needing labeling (indirect sensing) are prone to noisy labels and impact learning !
- Purely data driven models need to learn continuously (hence need new data) until full coverage of domain, but by then system/domain may have drifted !
    - Tesla trains FSD with 1.5PB → building their own cluster 5760 A100 GPUs
- Need an "intelligent" Data and Learning policy.

Robustness (of ML models)

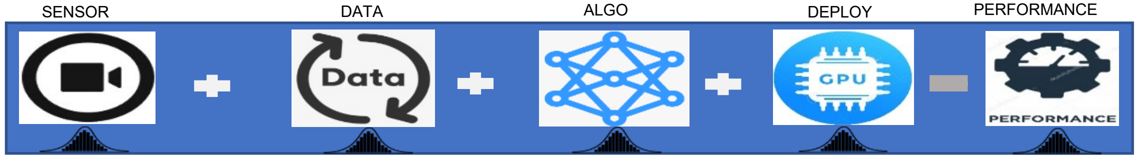

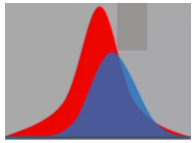

- Industrial and Engineering systems need guarantees and must be explainable!
  - Design, Safety, Performance, …
- ML methods, in general, lack formal approaches to understanding robustness, reliability, stability etc, but this may be changing !
- ML models rely on data, both data and models will have biases and can also drift → requires robustness checks at build as well as constant monitoring post deployment. →*Past performance is no guarantee of future results !!*
- Corner cases (more frequent than edge) vs Edge cases (did not think of these but Pr () >0 = unknown unknowns).
  - One may inadvertently transition into regions where Pr(Edge) increases!
- ASIL-B → < 1 Failure in 10^7 hours of driving. At ~ 60mph that is 6×10^8 miles of driving = 0.6B miles of driving. At 150K Miles vehicle useful life this is equivalent to data from 4000 cars driven to full useful life !!
- Currently Robustness implies massive amounts of testing, Simulation + Real life → different impacts depending on data class.

Corner or edge ?

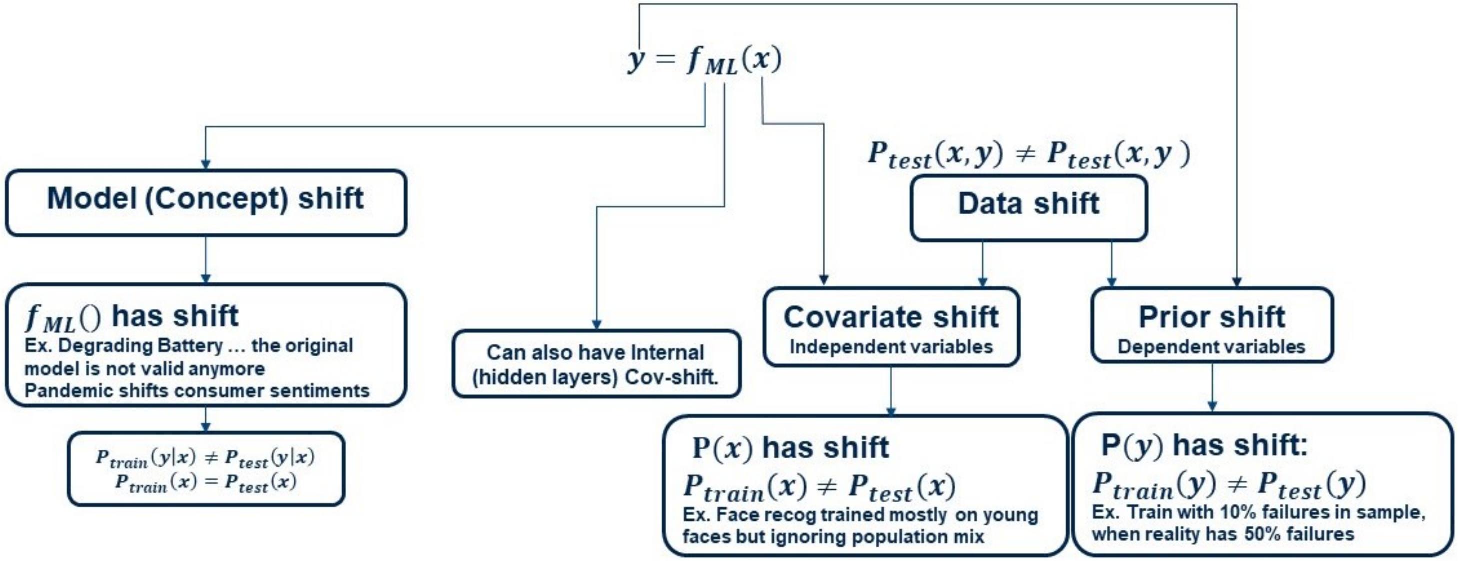

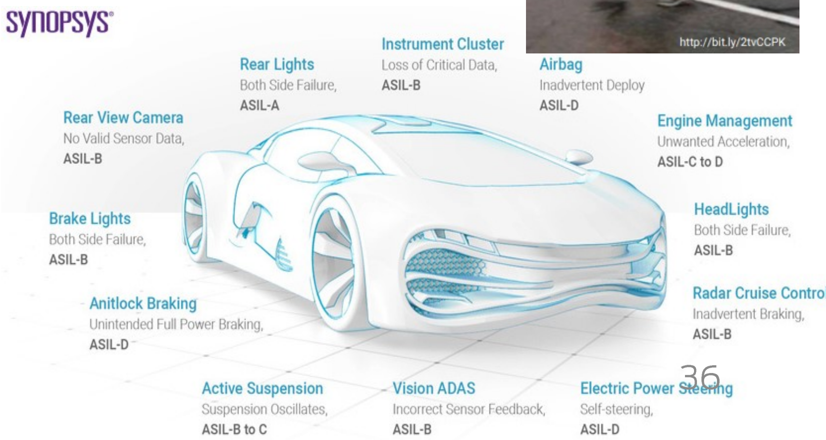

Causality

- Recently there has been growing interest in the CS-ML community to address causality (Yoshua Bengio, Hinton, LeCunn), Others, Pearl Marcus, etc have been asking for this.
- Understanding causality is critical in engineering/manufacturing systems
  - What is the root cause ? Formally (statistically) ... determine if a change in a given “treatment” leads to a change in some outcome → sensitivity analysis (engineering speak)
  - Humans can only manage short chain credit assignments, for multivariate influences … need compute
  - Extracting causality from pixel space is more difficult
- Discover causal relations by analyzing statistical properties of purely observational data
  - Expert knowledge: collection of facts and heuristics about the system
  - Granger causality: causality in time series data
  - Structural Equation Modeling (SEM): what factors determined the variable value
  - Casual Bayesian Network (CBN): what the probability of the variable changes when changing factors
  - Causal calculus (‘Pearl’)
- Causality in the action space – recovering from changes
  - Distributions (system representations) will change due to environmental pressures and/or direct intervention
  - Good causal models allow Causal induction from interventions .. Can we estimate the intervention

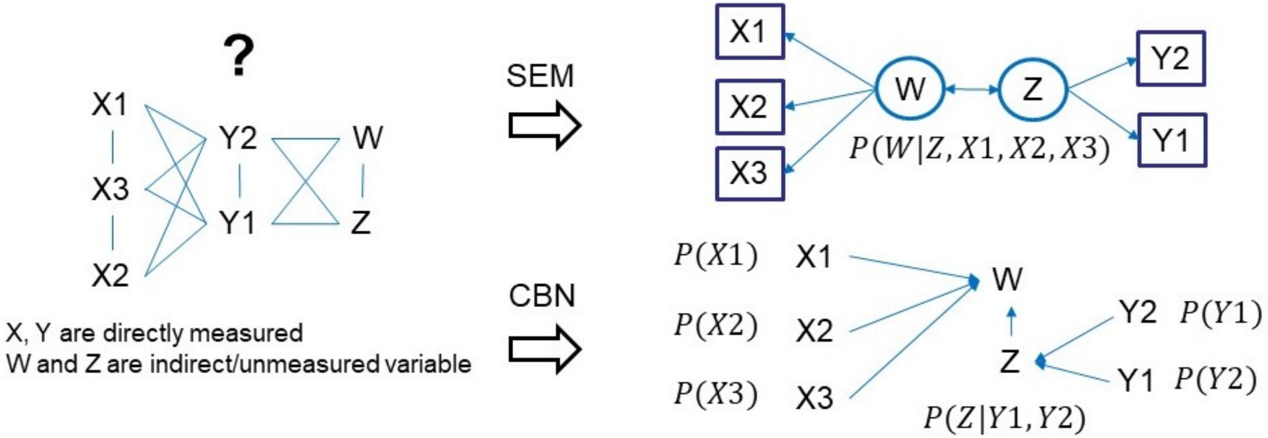

Extract the largest subgraph with target attribute

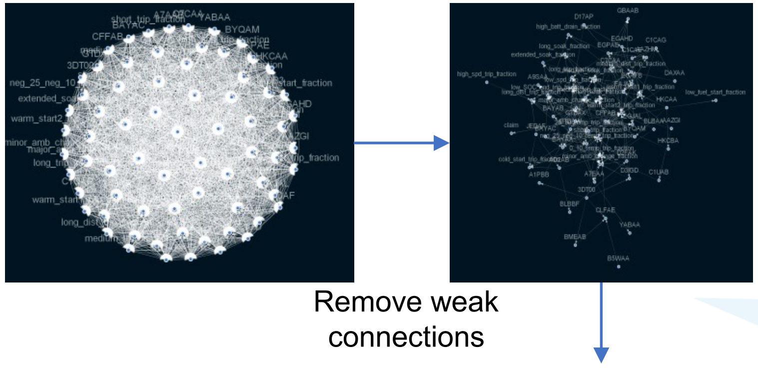

Classify

Cautions

- Regulations , Safety and Guidance
  - Most ML models will not satisfy current regulatory obligations, typically requiring > 99.99 accuracies.
  - Currently there is no mechanism to re-certify models that re-train on the fly.
  - Blame assignment remains an open question
  - Are simulators a good substitute, for real world testing ? This drives the need for “certifiably correct simulators”
  - Synthetic data and Photo real vs PBR (Light Structure interactions)
- Deployment, CI/CD is now CI/CD-CM
  - ML SW when deployed requires continuous monitoring (CM) given that models and data will drift as we discussed
- Fact and Fiction
  - Buyers beware ! This drives the need for in-house expertise.
  - There is an explosion of jargon that can be unnerving and quite tricky to navigate.
  - Beware of “Dashboards” , front ends can be dangerous.
- The “AI-Hammer” effect
  - Everything looks like a nail, and it sells (surprisingly)
  - OTS, pretrained networks are becoming commoditized, and there is a tendency to just train and re-train.
  - Strict quality control of ML models, in an industrial setting, is still evolving

Thank you for your time

- This presentation had contributions from several people on the Core AI-ML-QC team in Ford Research
  - Alireza Rahimpour, Alemayehu Admasu, Huanyi Shui, Harshal Maske, Hongjiang Li, Kaushik Balakrishnan, Mike Hopka, Patrick Blanchard (and his team) and the rest of the Core AI-ML-QC team.
- Guidance and support from: John Schneider (Director ESCAIT), Dimitar Filev (Tech. Fellow)
- We also included some results on PINNs from an ongoing project with MIT (Prof. Alan Edelman, Chris Rackuackas, et.al