

AI for Sustainability

Deploying Intelligent Diagnostics
and Optimal Control to Reduce
the Climate Burden of Google
Office Buildings

Herrick Conferences 2024
Compressors | Refrigeration | Buildings

John Sipple
sipple@google.com

Google Research



Effects of Climate Change

Extreme climate events comprising conditions beyond which many species are adapted are occurring on all continents, with severe impacts

Climate change is affecting ecosystem services connected to human health, livelihoods and well-being

Climate impacts on urban population health, livelihoods and well-being are felt disproportionately, with the most economically and socially marginalised being most affected



135 EJ



Global operational energy
demand in buildings

+4%



Increase from 2020, and 3%
previous peak in 2019

10 GtCO₂



Emissions exceed the
pre-pandemic all-time high in
2019 by 2%

22.3



¢/kWh average Commercial
Price for year ending April 2023
in California

86%



Commercial Electricity Price in
California above average for the
rest of the US

1



California's commercial
electricity prices were the
highest among the contiguous
states and DC

“ We are the first major company to make a commitment to operate on 24/7 carbon-free energy in all our data centers and campuses worldwide...by 2030. ”

Sundar Pichai

Potential Applications of AI for Smart Buildings

- Monitoring and Fault Detection
 - Anomaly Detection
- Fault Explanation and Root-Cause Analysis
 - Explainable AI and Causal Reasoning
- Optimal Control to minimize Energy Use and Carbon Emission
 - Reinforcement Learning and Model Predictive Control
- Demand and Emission Forecasting
 - Sequence models and time-series regression
- Scenario Analysis and Simulation
 - World models and physics-enhanced ML

Challenges for AI in Smart Buildings

- Non-deterministic behavior
 - Infinitely many solutions lead to many possible behaviors
 - Not suitable for some safety-critical tasks
- Brittleness under seasonality
 - Assumption of stationarity leads to overconfidence
 - Requires periodic retraining and evaluation
- Local Minimum
 - Potentially learn suboptimal behavior
- Data Gremlins
 - Bias/Drift
 - Confounders/Uncontrolled Factors
 - Undersampling/Aliasing



Tomorrow



Renewable (Replenishable)



Solar



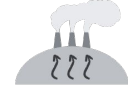
Hydro



Biomass



Wind



Geothermal

Non-renewable (Gone once used up)



Oil



Coal



Natural gas

We believe our near-term goals of reducing energy consumption and greenhouse gas emission at global scale can be achieved by applying AI for **intelligent diagnostics** and **optimal control**.

AI-Based Solutions for Smart Buildings

Optimal Control

Reduce energy use and CO2 emission

Explained Anomaly Detection


Explainable AI and Anomaly Detection to detect HVAC faults

Project Freon

Detect leaky Air Conditioners

Intelligent Diagnostics

Provide root-cause analysis and recommend fixes



Optimal Control for Reducing Carbon
Emission and Energy Consumption

Research Hypothesis

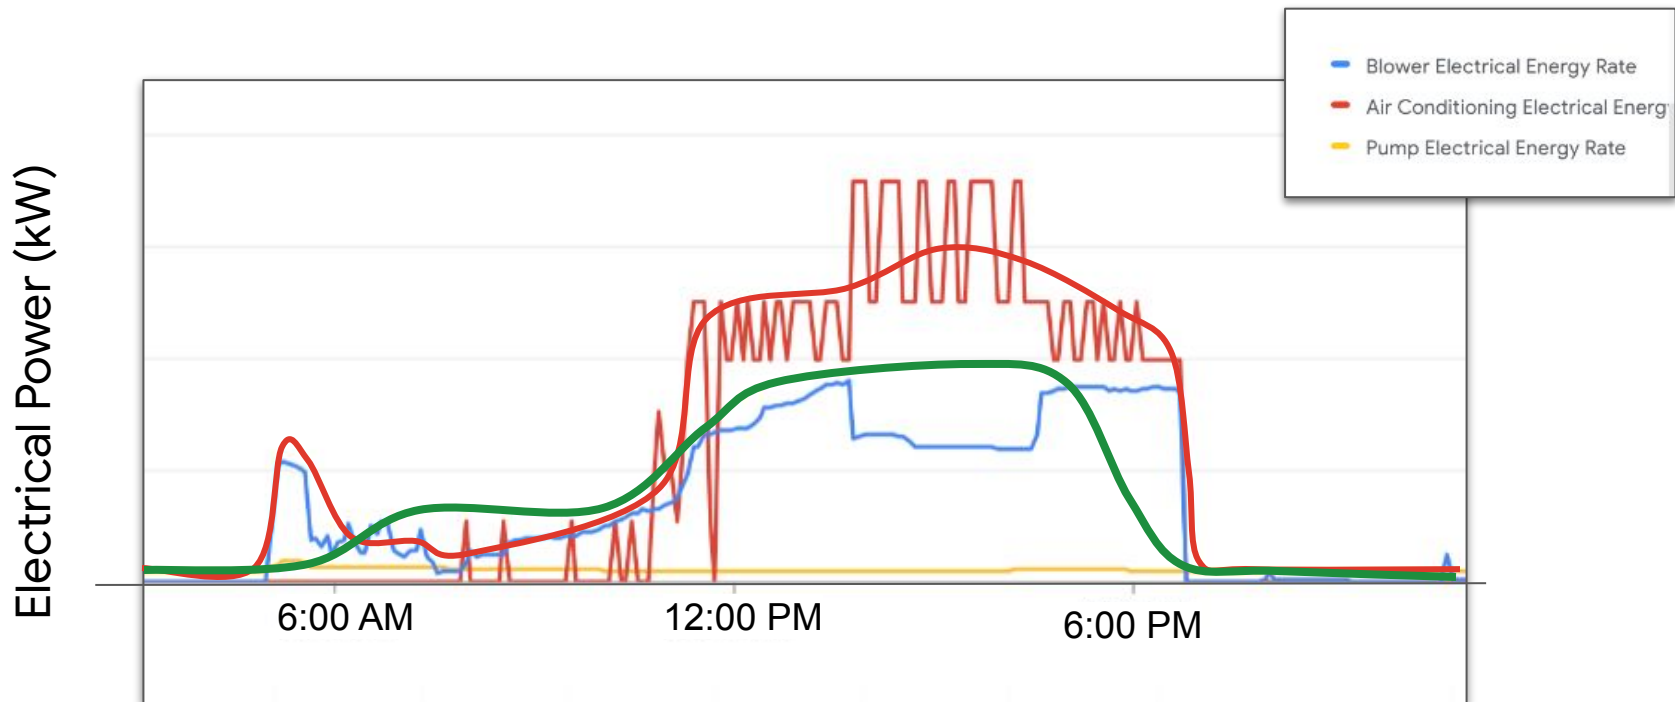
Objective:

Adaptive control can reduce energy consumption and greenhouse gas emission from HVAC devices in commercial office buildings without compromising occupant comfort.

Minimal Success Criteria:

- 5%+ reduction in carbon emission,
- 5%+ reduction in energy consumption,
- maintain occupant comfort conditions

A typical workday



Exterior Environment



Chiller



Air Handler



Meters



Boilers

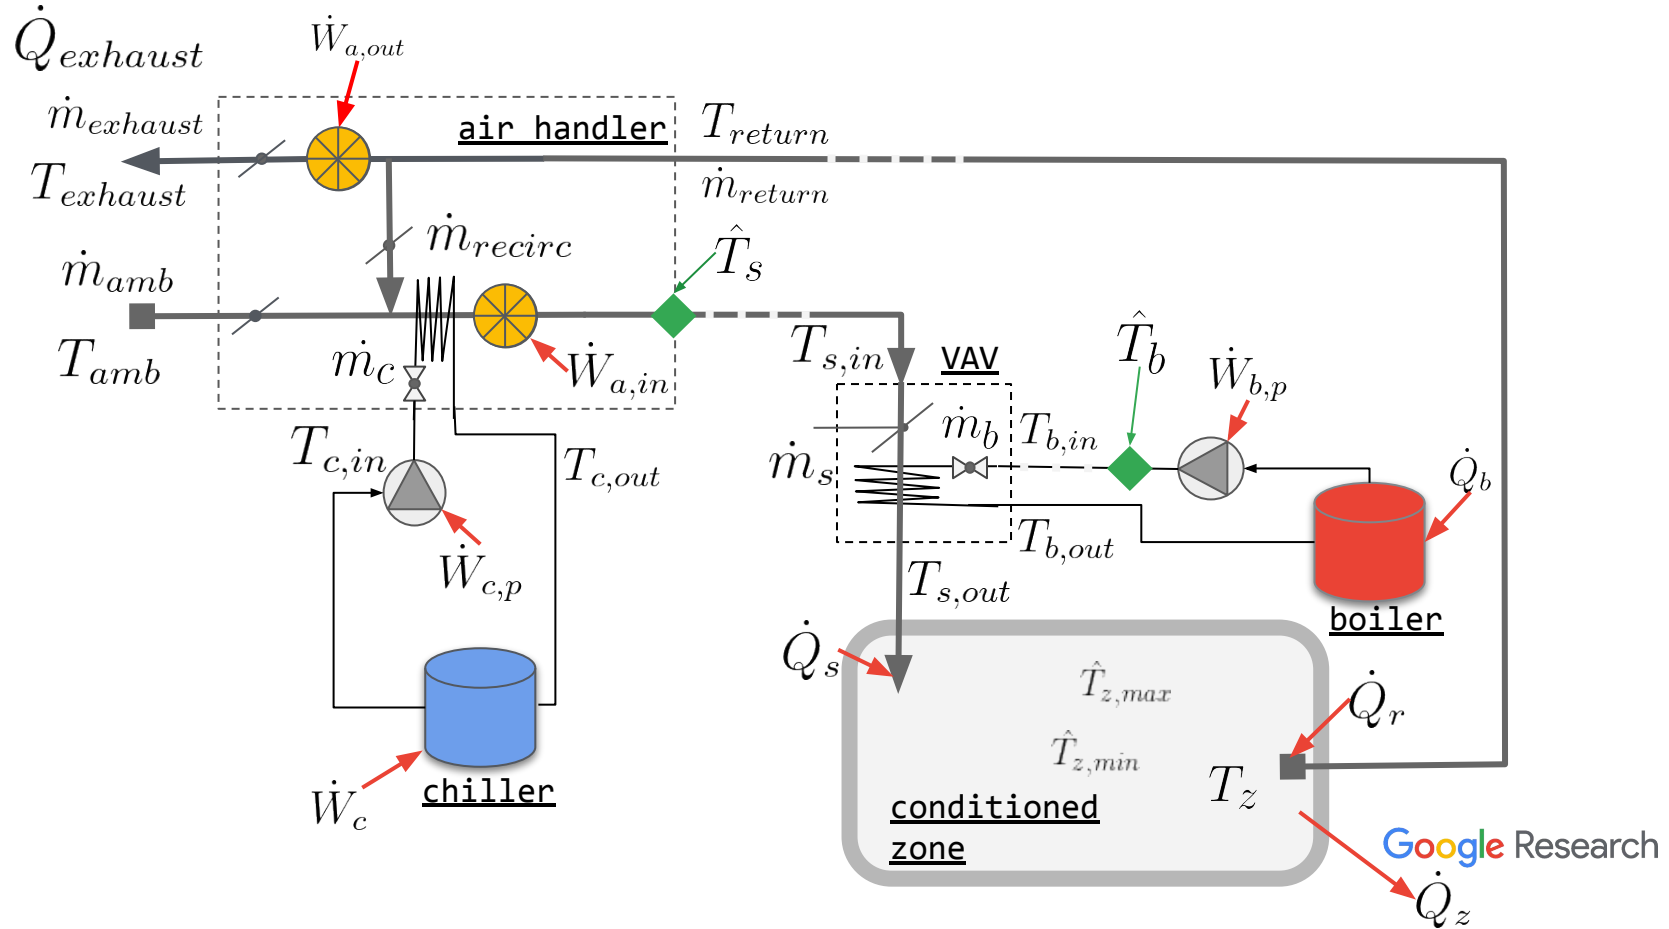


Interior Environment

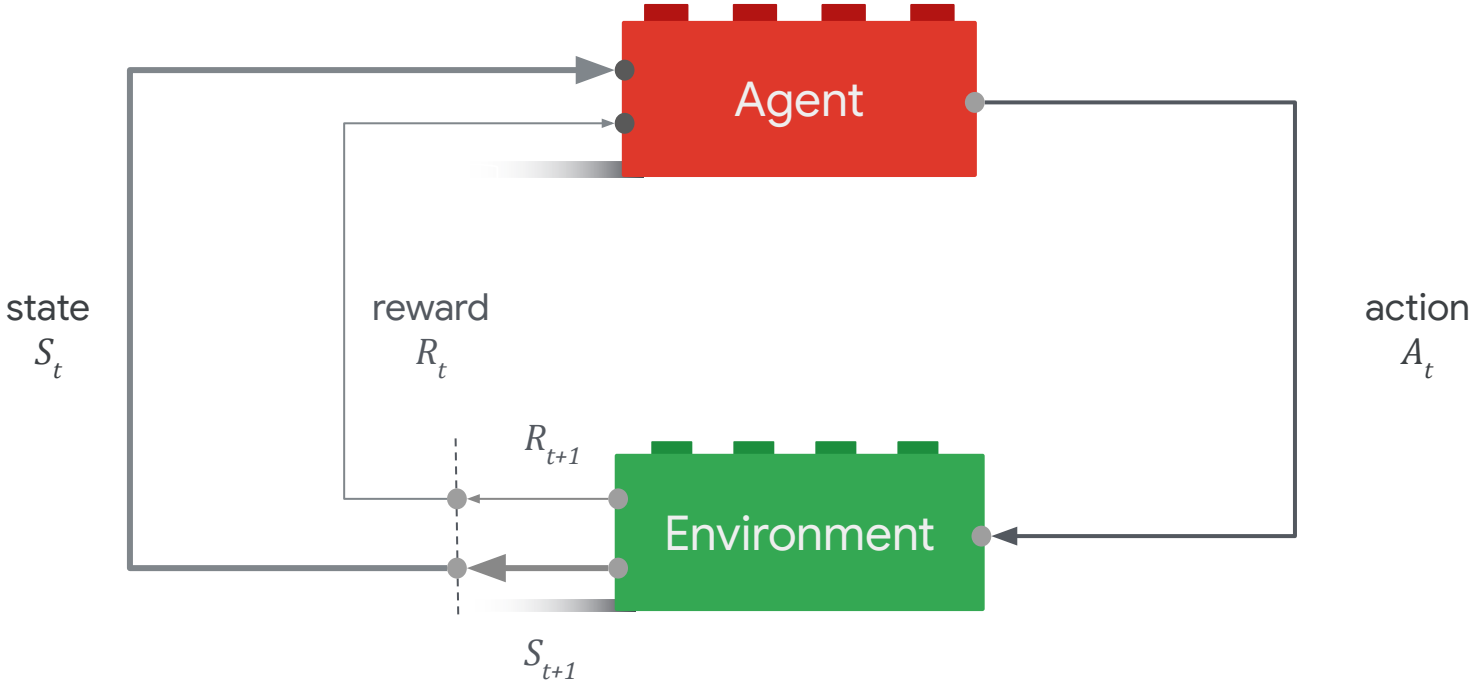


Variable Air Volume (VAV)

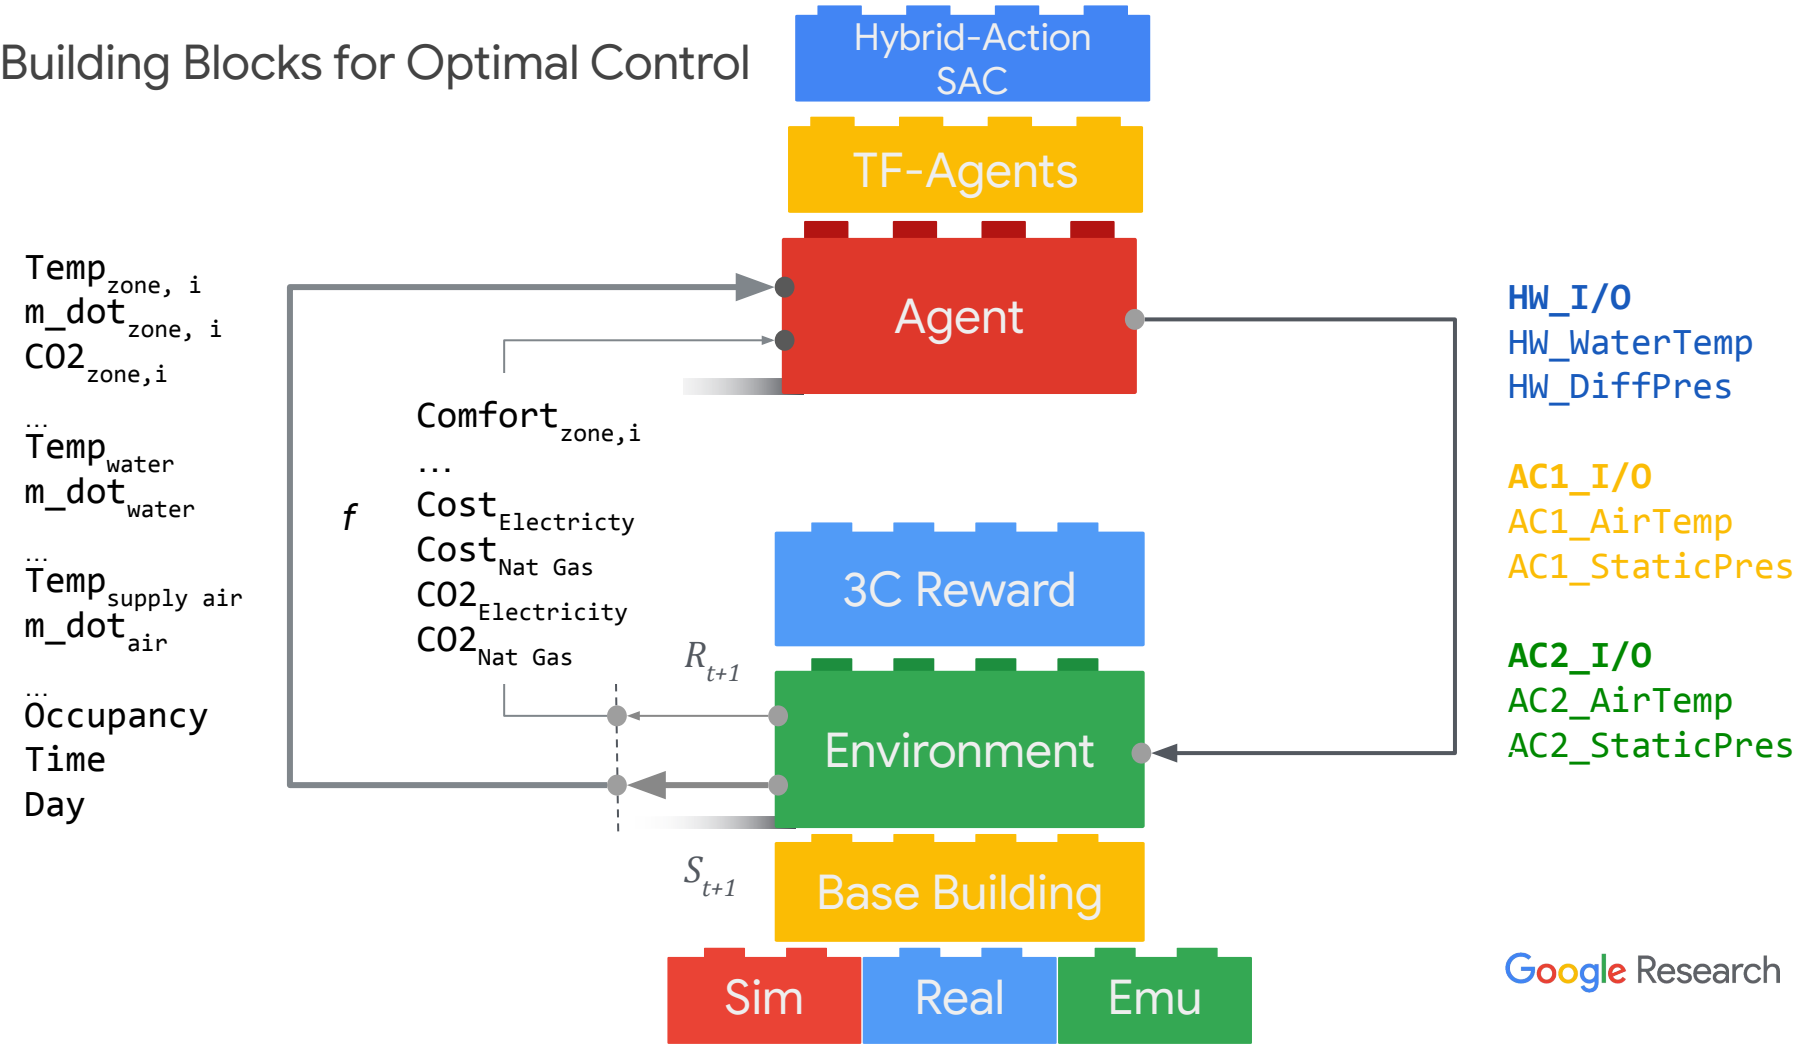
Environment Thermal Model



Optimal Control with Reinforcement Learning



Building Blocks for Optimal Control



Reward Function

3C Reward Function

Comfort

Google exceeds standard comfort conditions for office buildings. Comfort conditions include room temperature and ventilation/CO2 levels. The agent must maintain these standards while minimizing cost.

Cost

Electrical energy cost varies by time of day and day of week. Natural gas costs fluctuate less, but also vary monthly. Typically, buildings are biased to comfort and tend to consume more energy than required to meet comfort conditions.

Carbon

Both natural gas and electricity emit carbon dioxide into the atmosphere. The agent should minimize carbon emissions.

$$r = u \times r_{comfort} + v \times r_{cost} + w \times r_{carbon}$$

Offline Training

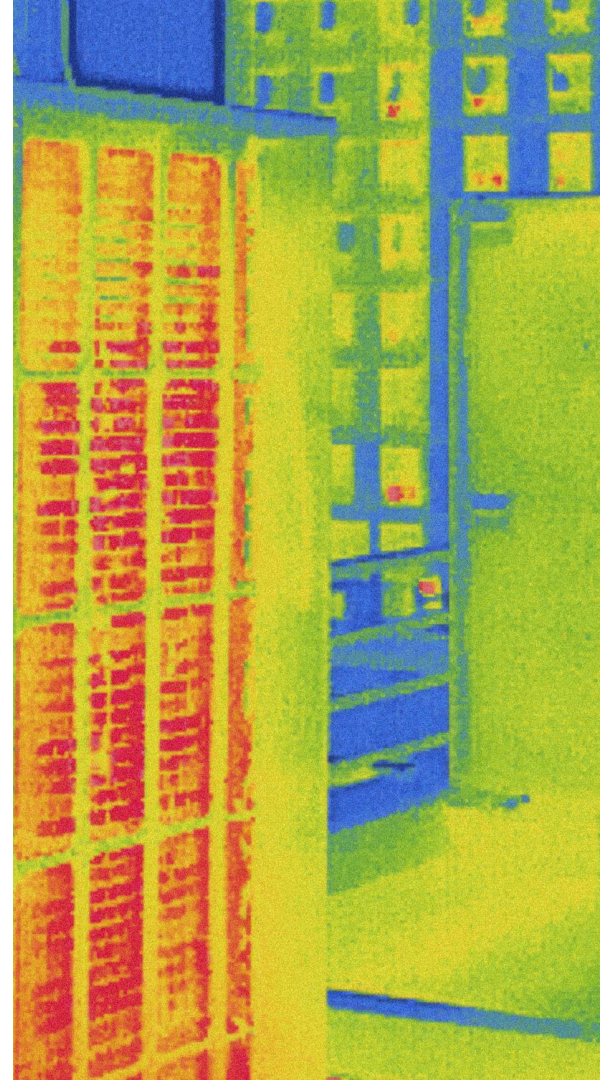
Offline Training

Need to train a viable agent offline

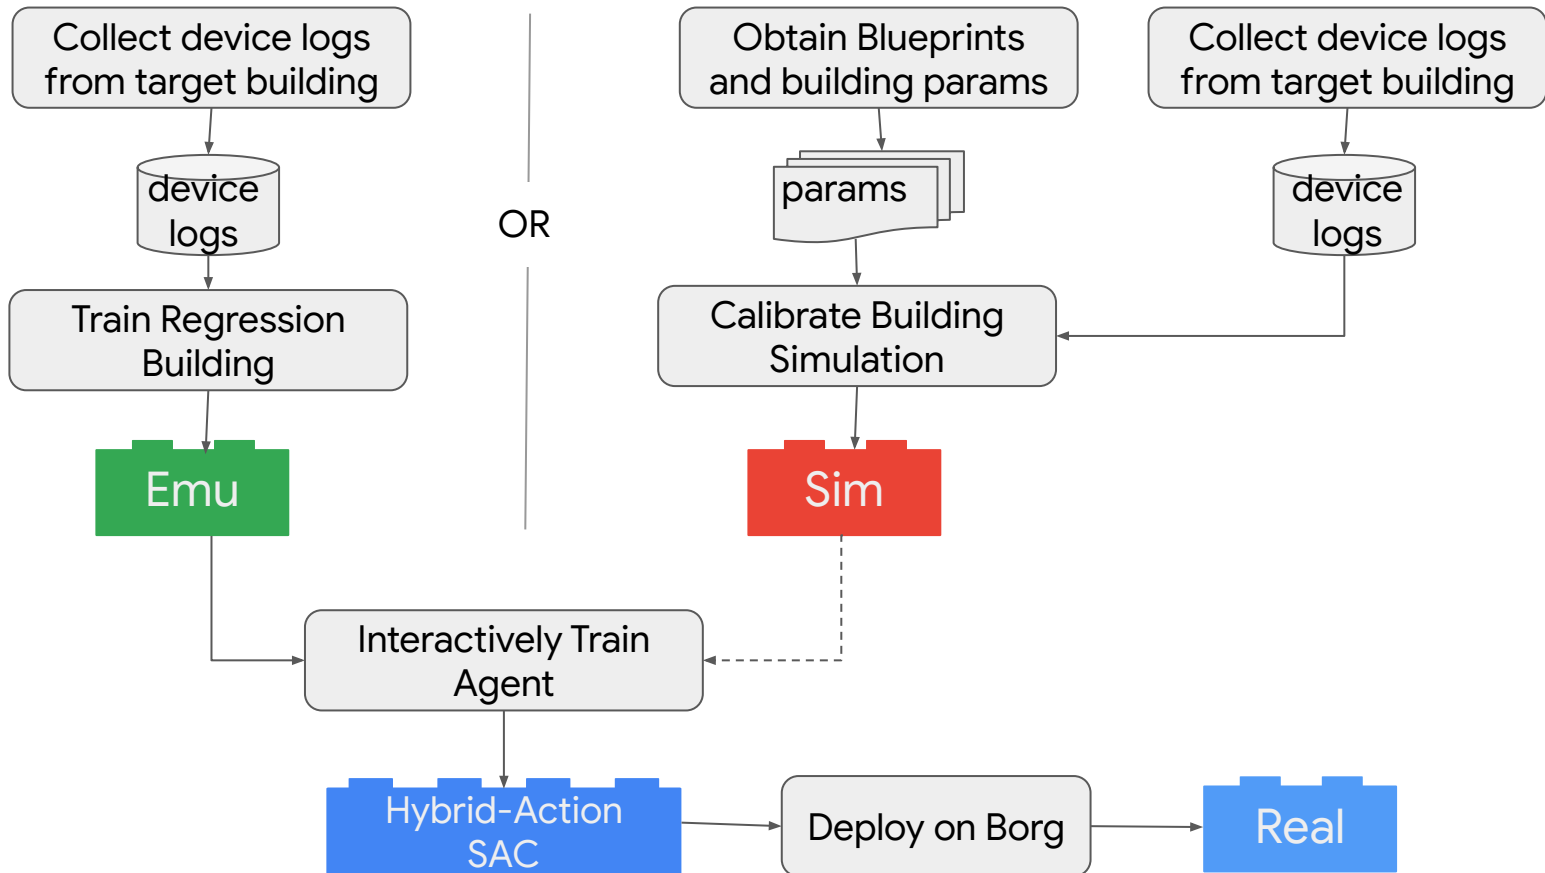
- Training requires many iterations
- Thermodynamics is slow

“*Digital twins*” for offline training

- Simulation: Model the thermodynamics
- Emulation: Multivariate regression



Offline Training Approaches: **Emulation** vs. **Simulation**



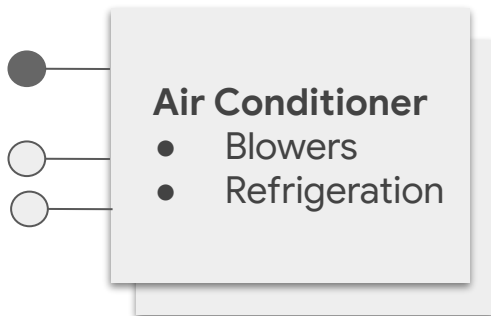
Agent Design

Hybrid Action Spaces

AC_I/O = ON

AC_StaticPres = 0.69iwc

AC_AirTemp = 71.2F



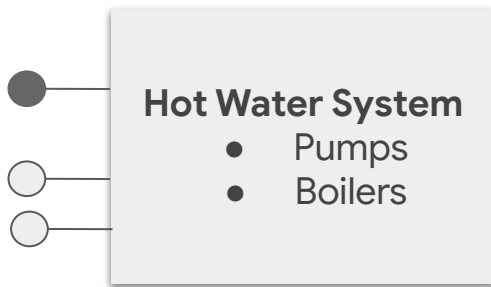
Ventilation

Cooling

HW_I/O = OFF

DiffPres = 12.1psi

WaterTemp = 101.3F



Circulation

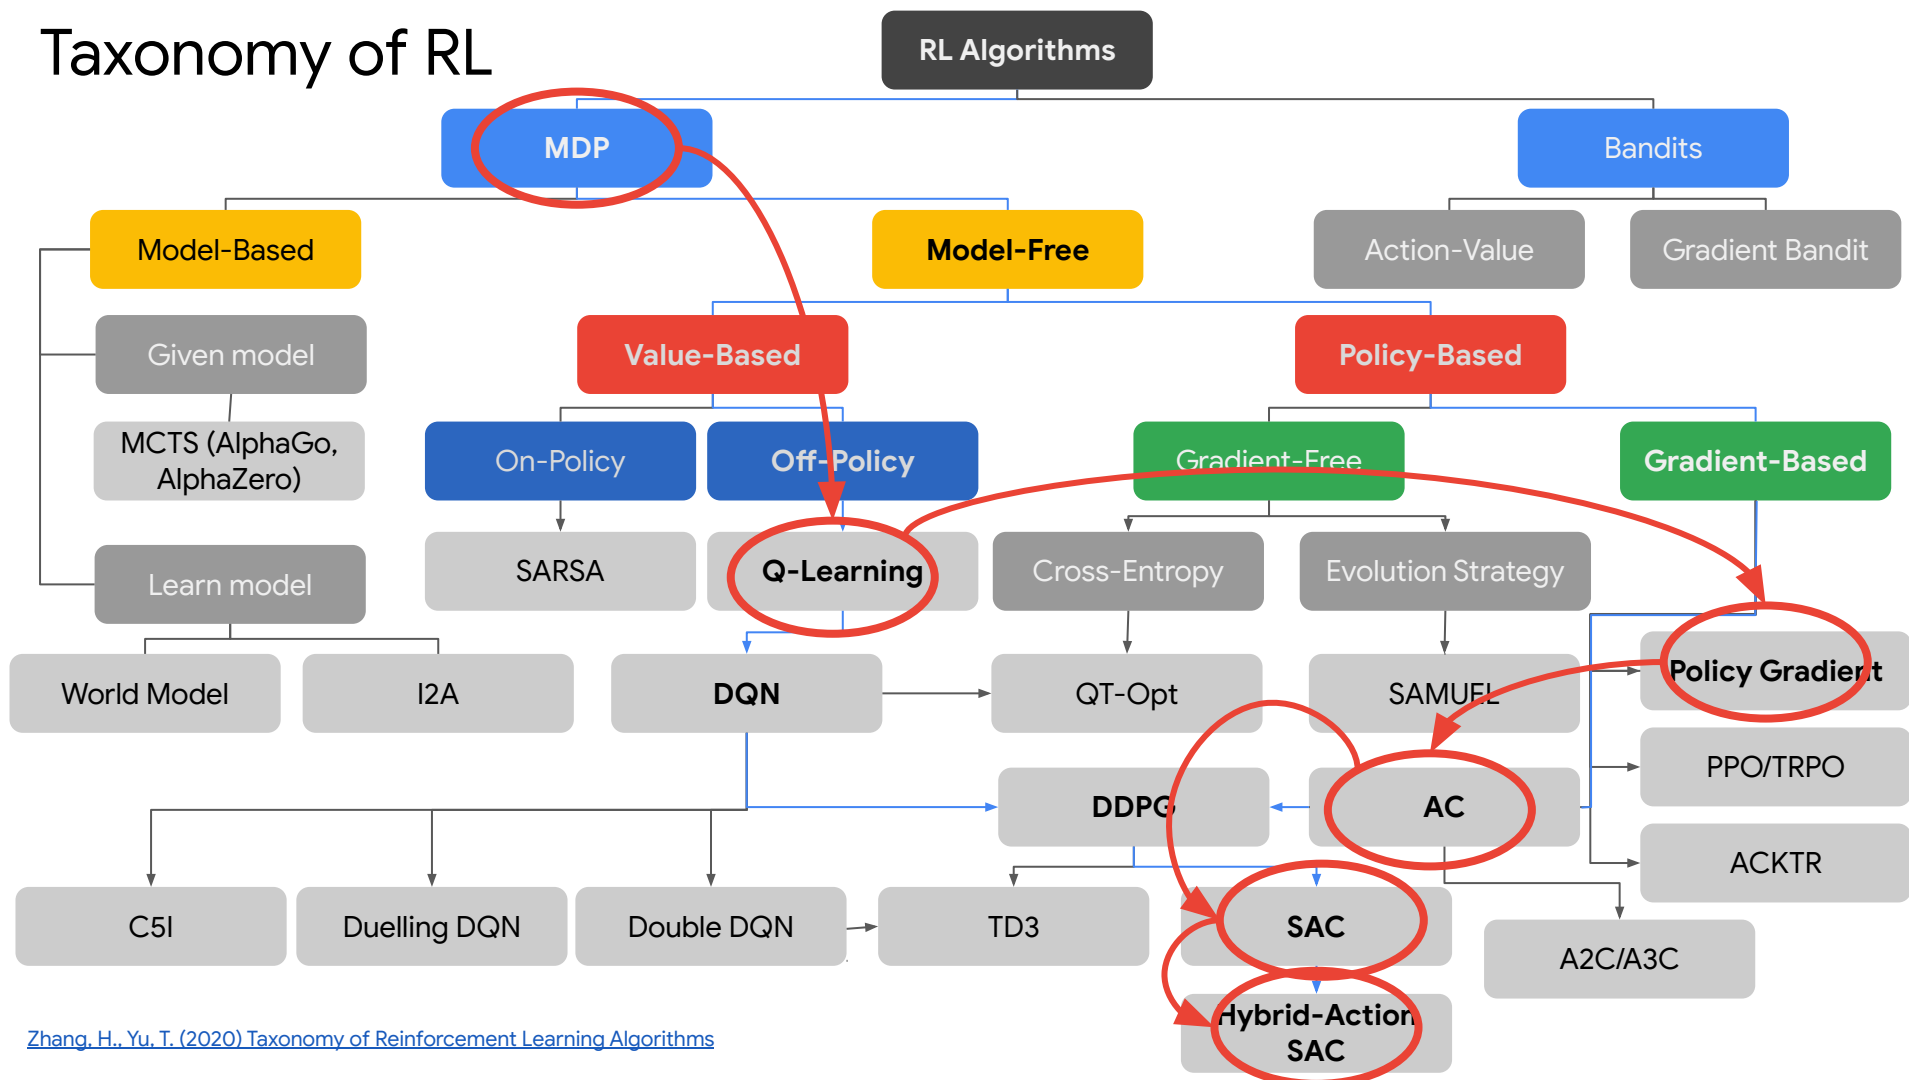
Heating

Actions are both continuous and discrete.

Standard RL algorithms are generally discrete or continuous, but few are both

Very few benchmark RL environments that expose hybrid actions

Taxonomy of RL



Hybrid-Action Soft Actor-Critic

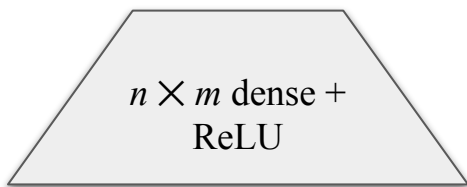
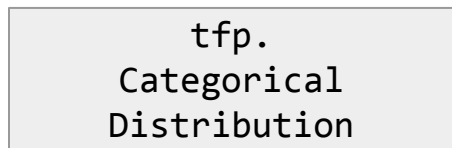
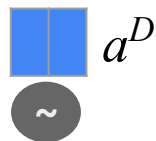
Applies existing “best practices”

- Policy Gradient
- Actor-Critic
- Experience Replay Buffers
- Target actor and dual critic networks
- Maximum Entropy RL
- Continuous Action Spaces

Introduces two new concepts:

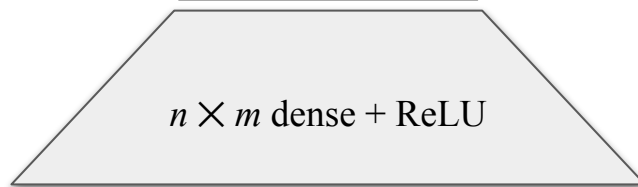
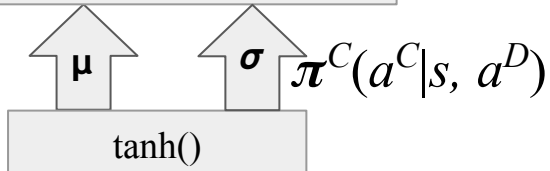
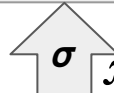
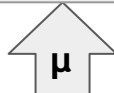
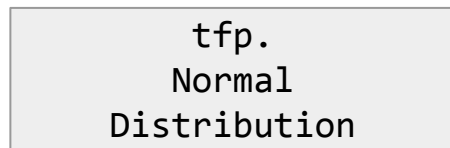
- Action Pretraining
- Hybrid Action Spaces

Hybrid-Action SAC Networks



s

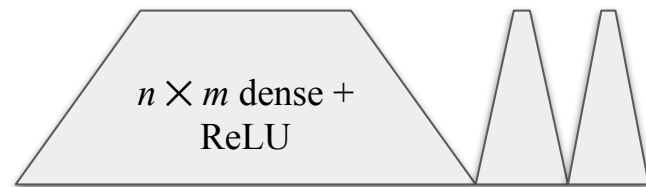
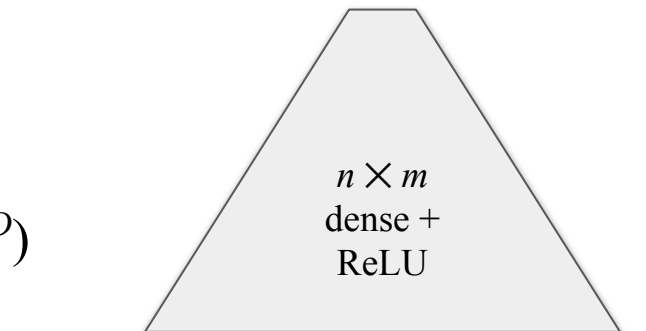
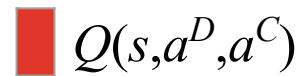
Discrete Actor Network



$\pi^D(a^D|s)$

s

Continuous Actor Network



s

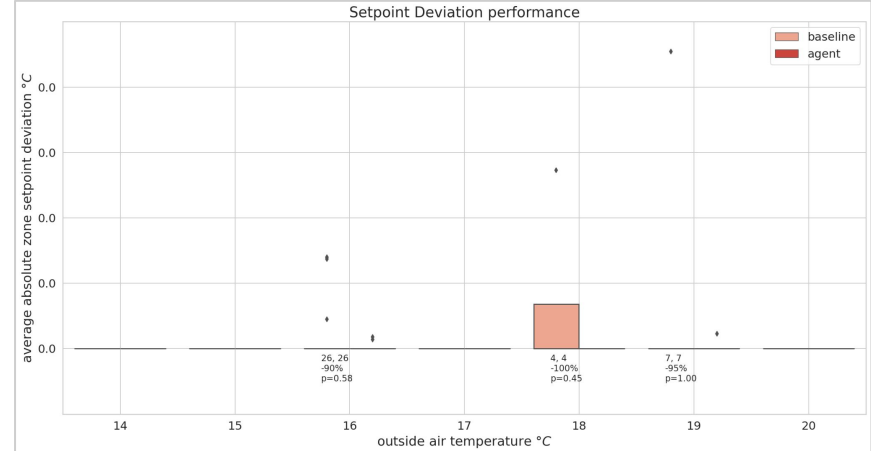
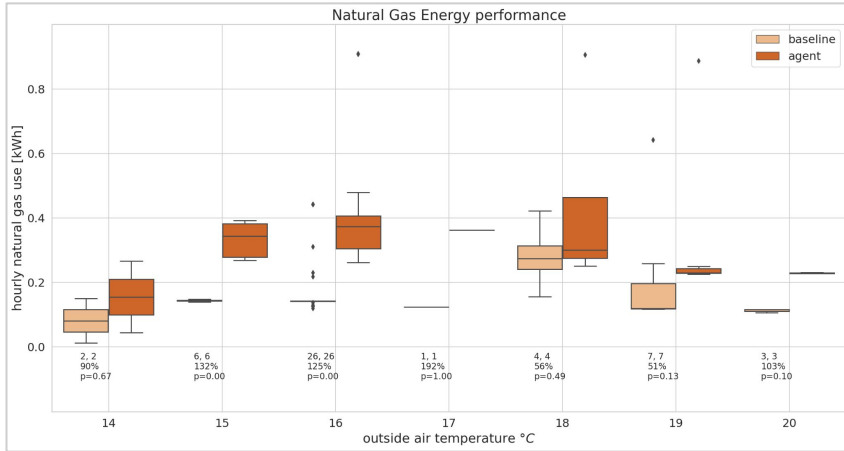
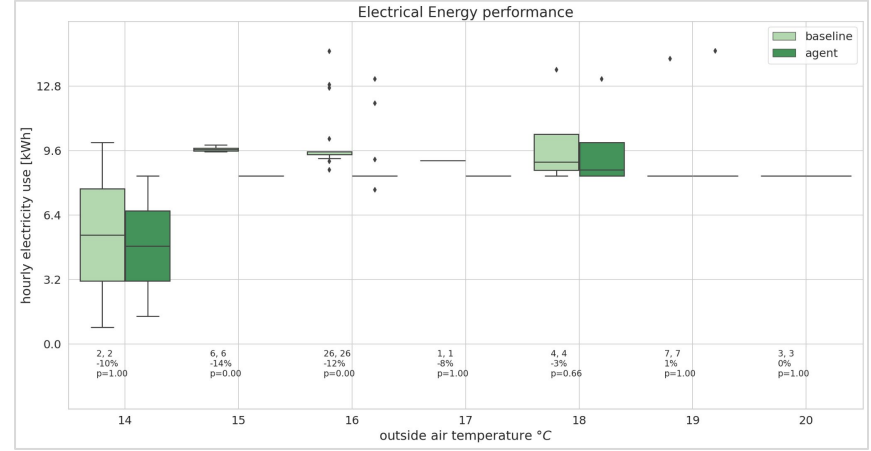
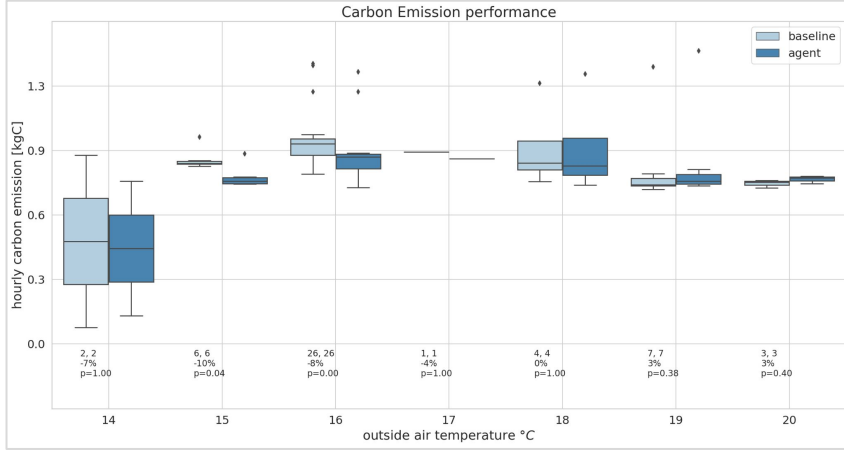
a^D

a^C

Critic Network

Validation

Validation Metrics





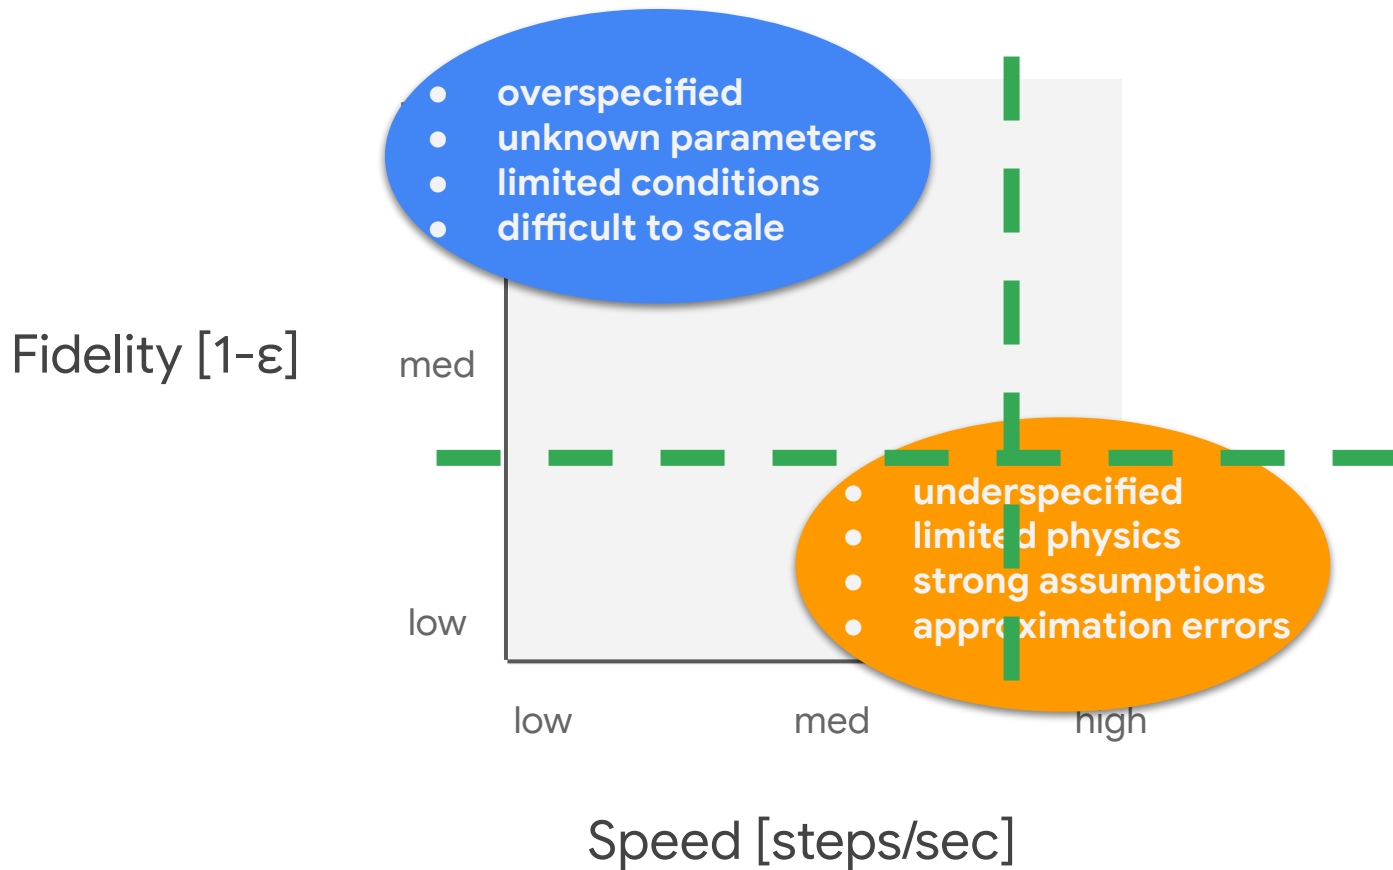
Google

1055 JOAQUIN RD



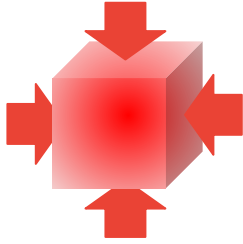
Calibrated Building Simulation to estimate building energy consumption and carbon emission enabling offline agent training

Sim Fidelity vs. Efficiency



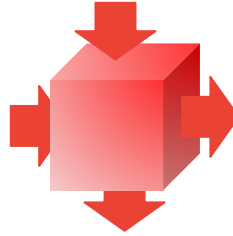
Energy Balance for unsteady thermal diffusion

Thermal
Absorption



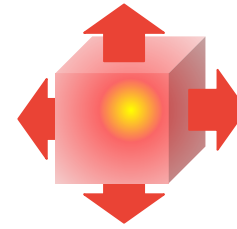
+

Thermal
Flux



=

Thermal
Source



$$\frac{d}{dt} \int_{V(t)} Q dV + \oint_{S(t)} \mathbf{n} \cdot \mathbf{F} dS = \int_{V(t)} P dV$$

t - time

V - volume

Q - energy absorption by volume

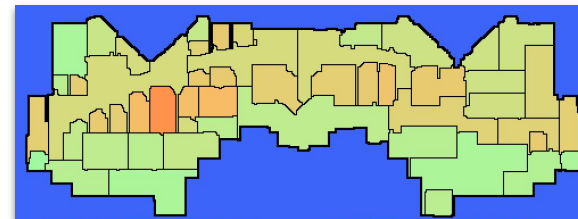
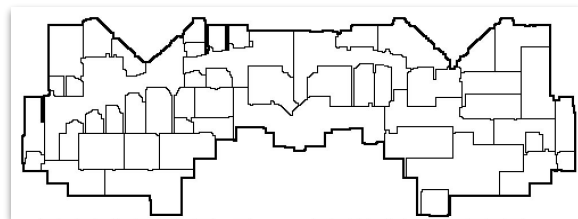
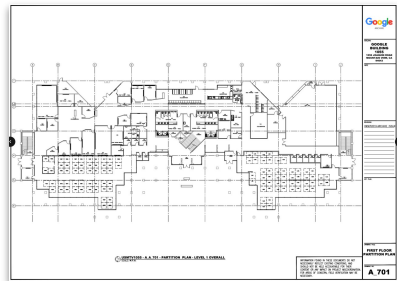
S - surface area

F - Flux, energy per surface area

n - Surface normal vector

P - Energy source per volume

Customizing the Simulation for Other Buildings



Extract
Floorplan

Ingest an image of a floorplan, and using OpenCV extract a floorplan grid.

Pre-
process

Preprocess floorplan grid to make sure it is uniform. Add device parameters.

Label
Devices

Interface that allows a human to quickly assign devices to zones

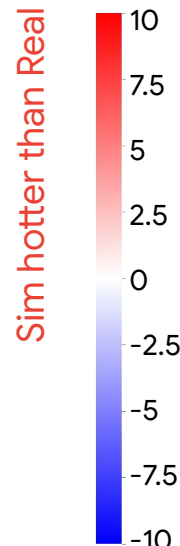
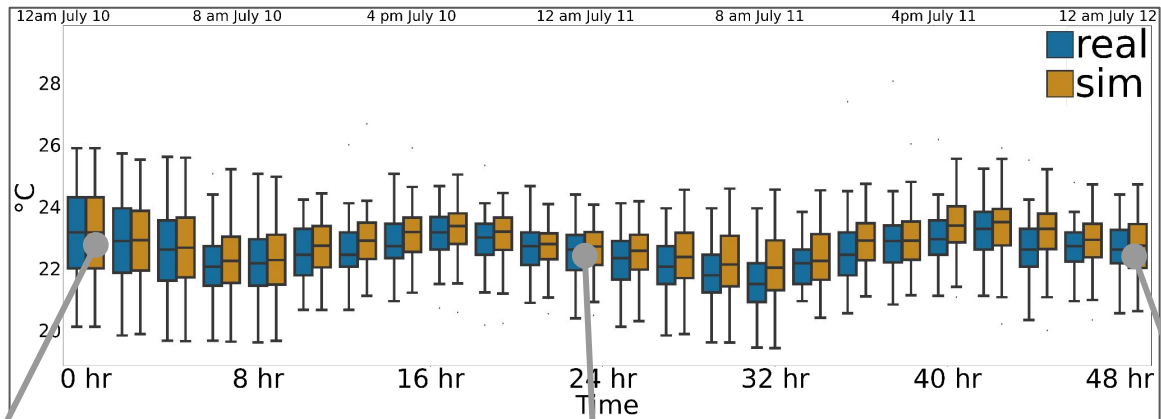
Generate
Simulator

Generate a simulated RL environment that matches the real building in physical layout, device layout, and HVAC structure.

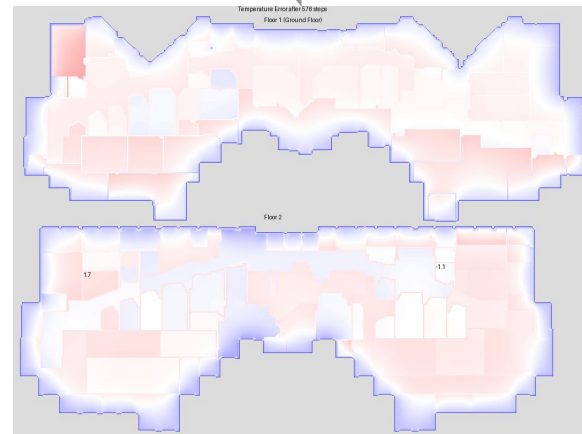
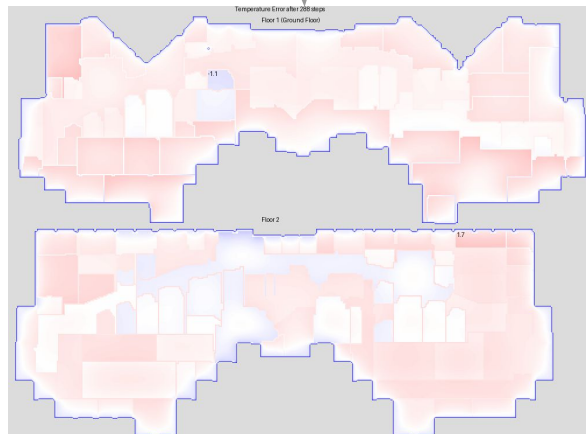
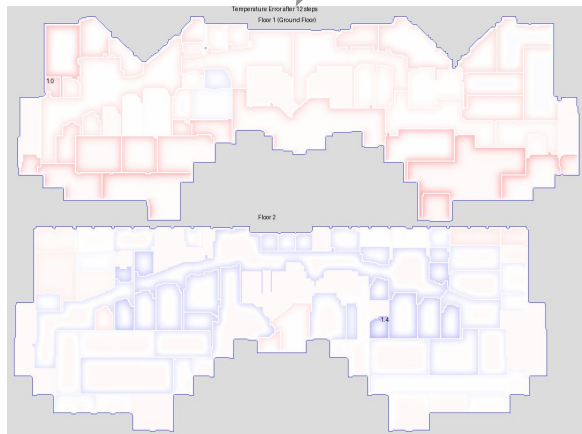
Calibrate
w/ Real
World

Use real world data to hyperparameter tune the simulator, to aid in sim-to-real transfer.

Calibrated Simulator Performance on Zone Air Temperatures



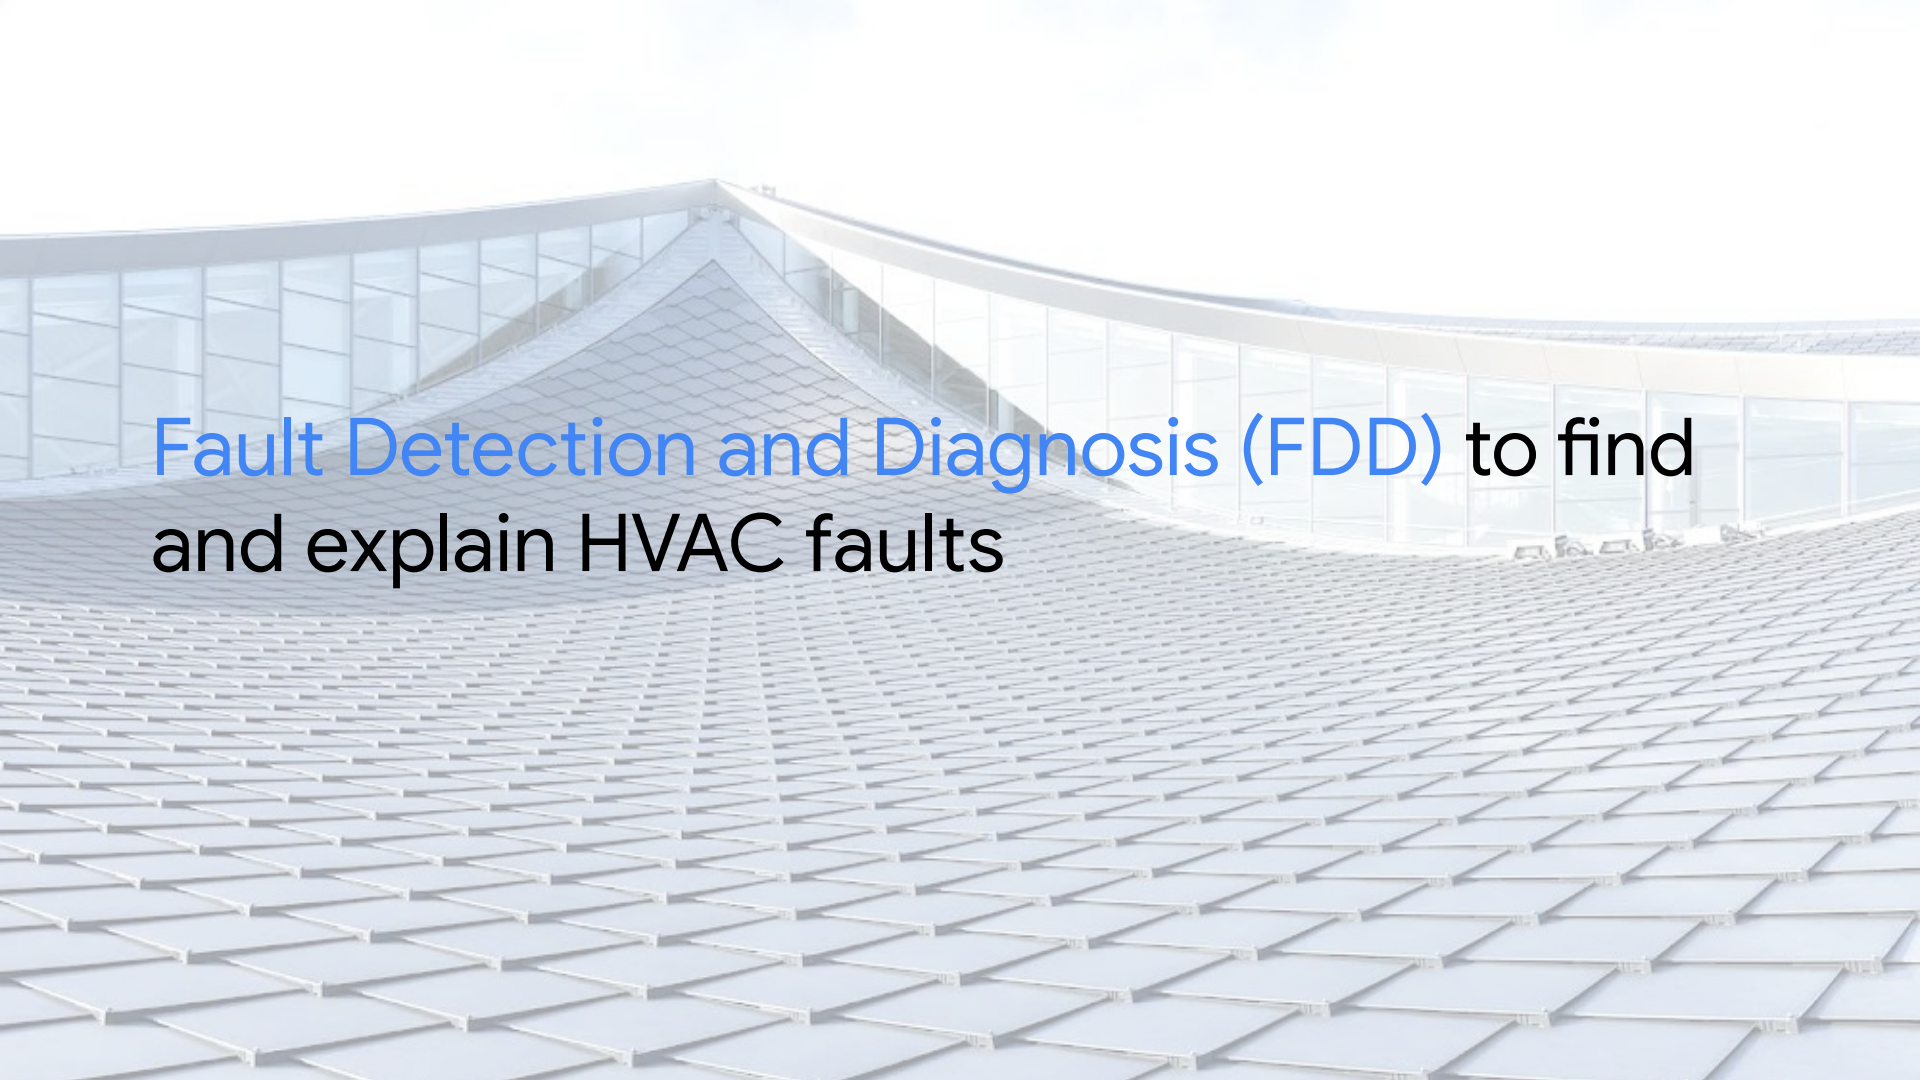
Sim cooler than Real



Mean and Median Absolute Error on $N = 72$ prediction window

Metric	Tuning Data (a)	Validation (b)	Validation (c)
MAE	0.64 °C	0.63 °C	1.18 °C
Median	0.01 °C	0.18 °C	0.98 °C

Goldfeder, J., Sipple, J. (2023) [A Lightweight Calibrated Simulation Enabling Efficient Offline Learning for Optimal Control of Real Buildings](#)

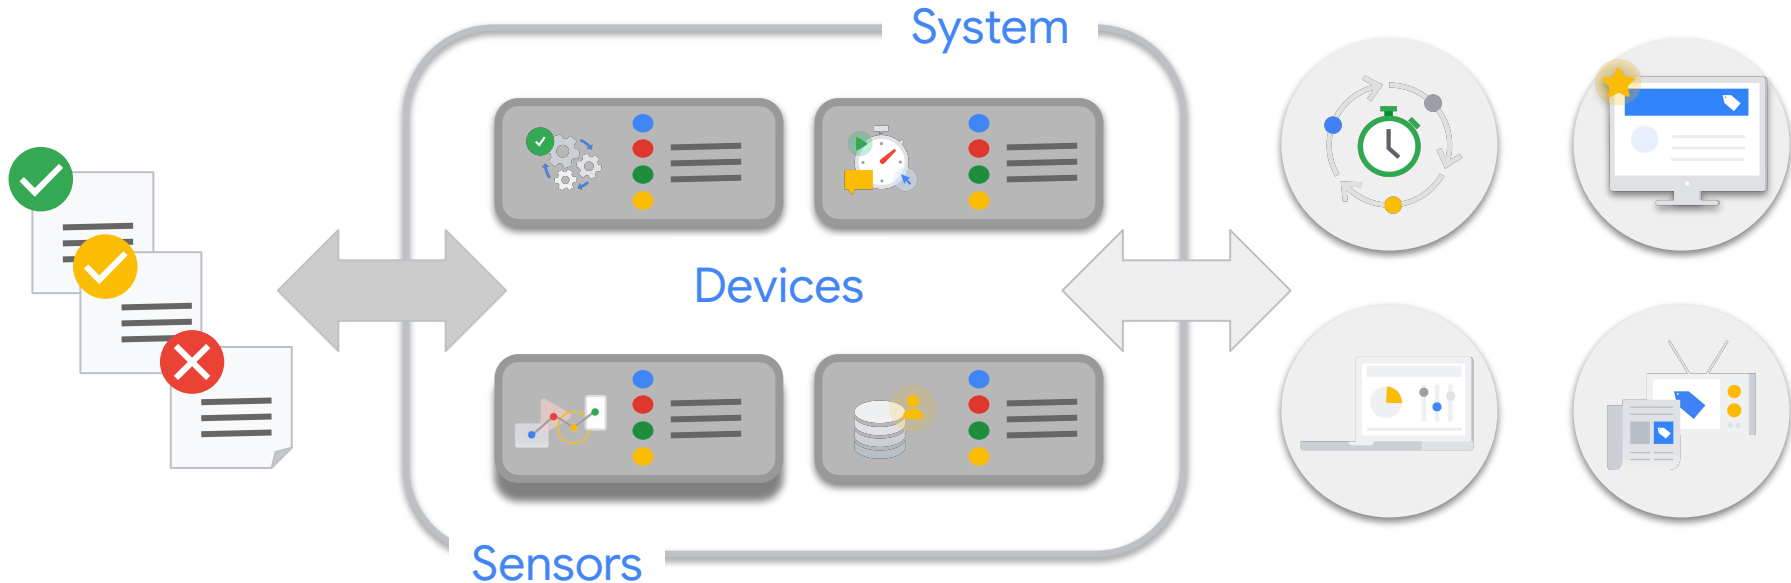


Fault Detection and Diagnosis (FDD) to find
and explain HVAC faults

Telemetry Environment

Alerting/Diagnostics

Services/Utilities



Heating, Ventilation, Air Conditioning (HVAC)

Problem: ACs, Hot Water Systems, VAVs fail, resulting in uncomfortable conditions and wasted energy and CO2 emission

Stakeholder: REWS, FORT

Fleet: 35k+ devices in 200 buildings (200+ types)

Datastream: Unlabeled numeric multivariate, temps, pressure, etc., 5 min interval

Anomaly Detector: MADI



Anomaly Detection

Device generates a sequence of D -dimensional **observations** x

$x = [\text{zone air temp} = 18.0 \text{ }^\circ\text{C}, \text{static pressure} = 128 \text{ pa}, \dots]$

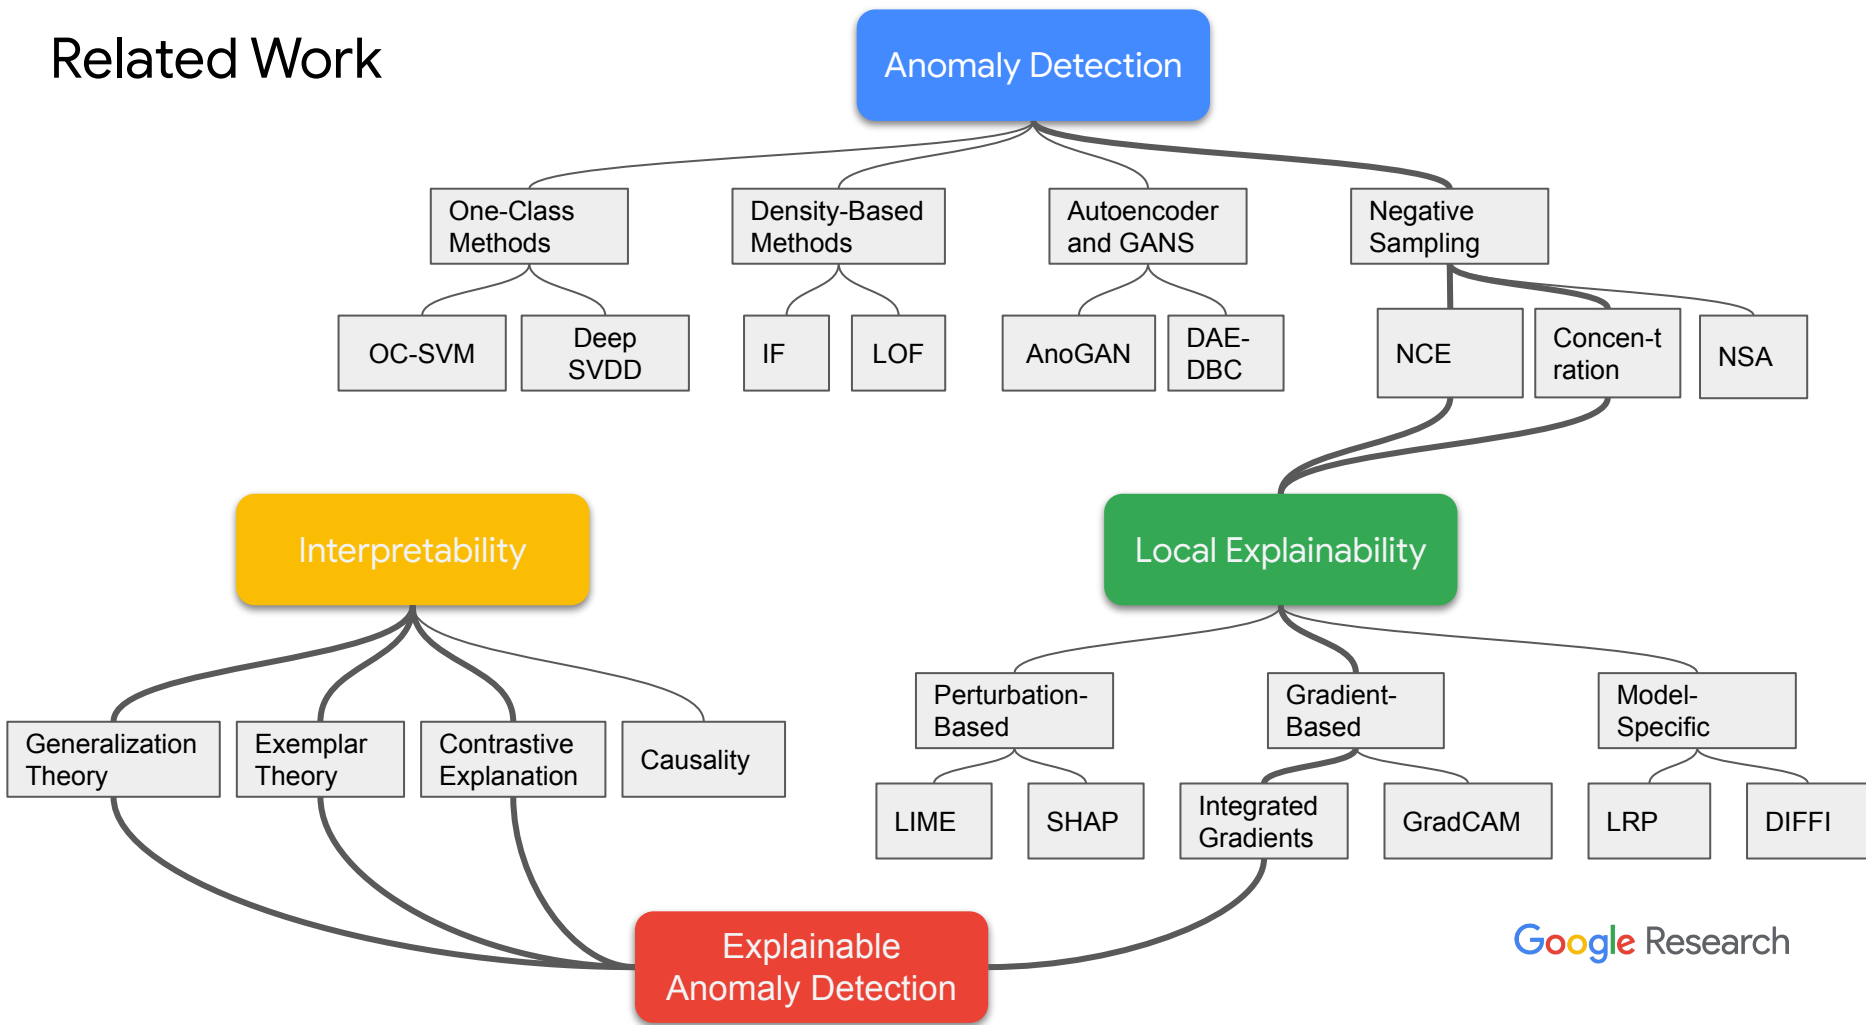
Anomaly Detection is the task of determining whether x is Normal

Anomaly Detector is a differentiable function
 $F: \mathbb{R}^D \rightarrow [\textit{Anomalous} \approx 0 \text{ to } \textit{Normal} \approx 1]$

What is "normal"?

How do we test?

Related Work



Anomaly Detection with MADI

Multivariate Anomaly Detection with Interpretability (MADI)

- Sipple J., (2020) Interpretable, Multidimensional, Multimodal Anomaly Detection with Negative Sampling for Detection of Device Failure, ICML 2020
- Sipple J., Youssef, A., (2022) A General-Purpose Method for Applying Explainable AI for Anomaly Detection, ISMIS 2022

Anomaly Detection based on Noise-Contrastive Estimation and Neural Nets

Anomaly Explanations using Integrated Gradients

Anomaly Detection Results

ROC-AUC %	OC-SVM	Deep SVDD	Iso Forest	Extended Iso Forest	NegSampleRnd Forest	NegSample Neural Net
Forest Cover*	53 ±20	69 ±7	85 ±4	93 ±1	80 ±2	86 ±4
Shuttle*	93 ±0	88 ±9	96 ±1	91 ±1	93 ±7	96 ±5
Mammography*	71 ±7	78 ±6	77 ±2	86 ±2	85 ±4	84 ±2
Mulcross*	90 ±0	54 ±4	88 ±0	66 ±4	94 ±1	99 ±1
Satellite*	51 ±1	62 ±3	67 ±2	71 ±3	65 ±4	73 ±3
Smart Buildings	76 ±1	60 ±7	71 ±7	80 ±4	95 ±1	93 ±1

* Courtesy of ODDS Library [<http://odds.cs.stonybrook.edu>].
Stony Brook, NY: Stony Brook University, Department of Computer Science

Anomaly Explanation

Given $F(x) \approx 0$, what make x anomalous?

- Describe the Symptoms
- Map x to an **attribution**
- May also use a **contrastive normal point** x'

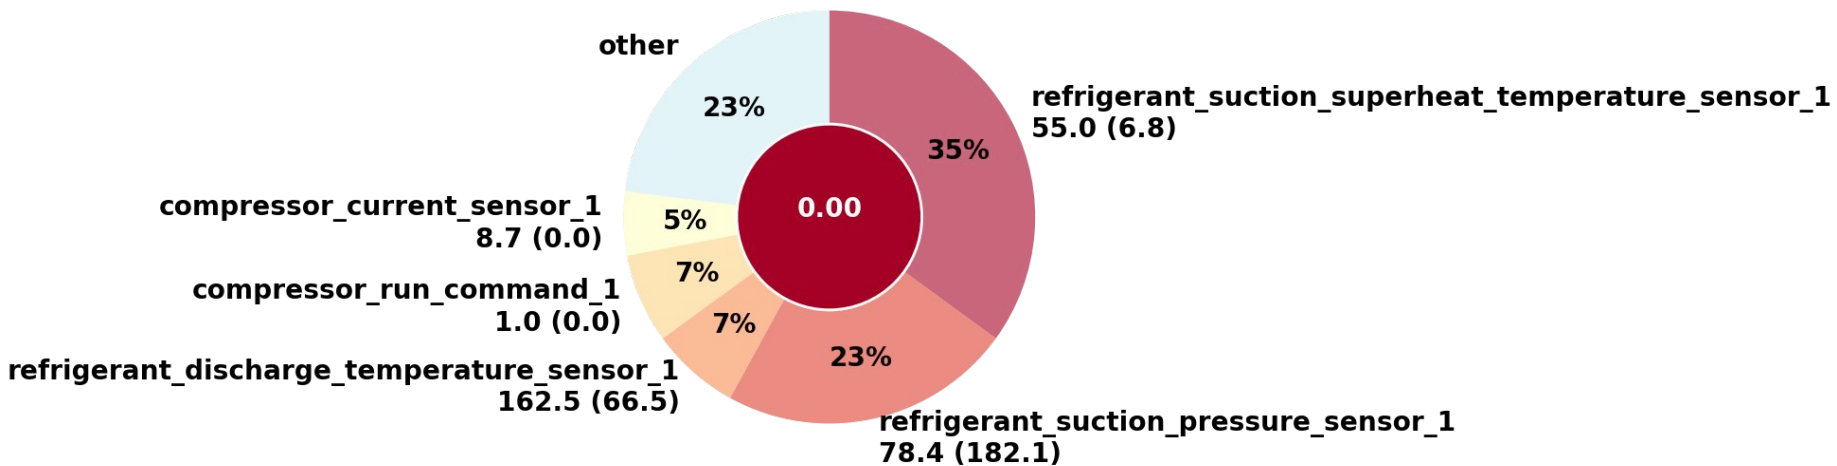
Why is it “anomalous”?

Enables an expert's **interpretation** that maps
Symptoms to Diagnosis

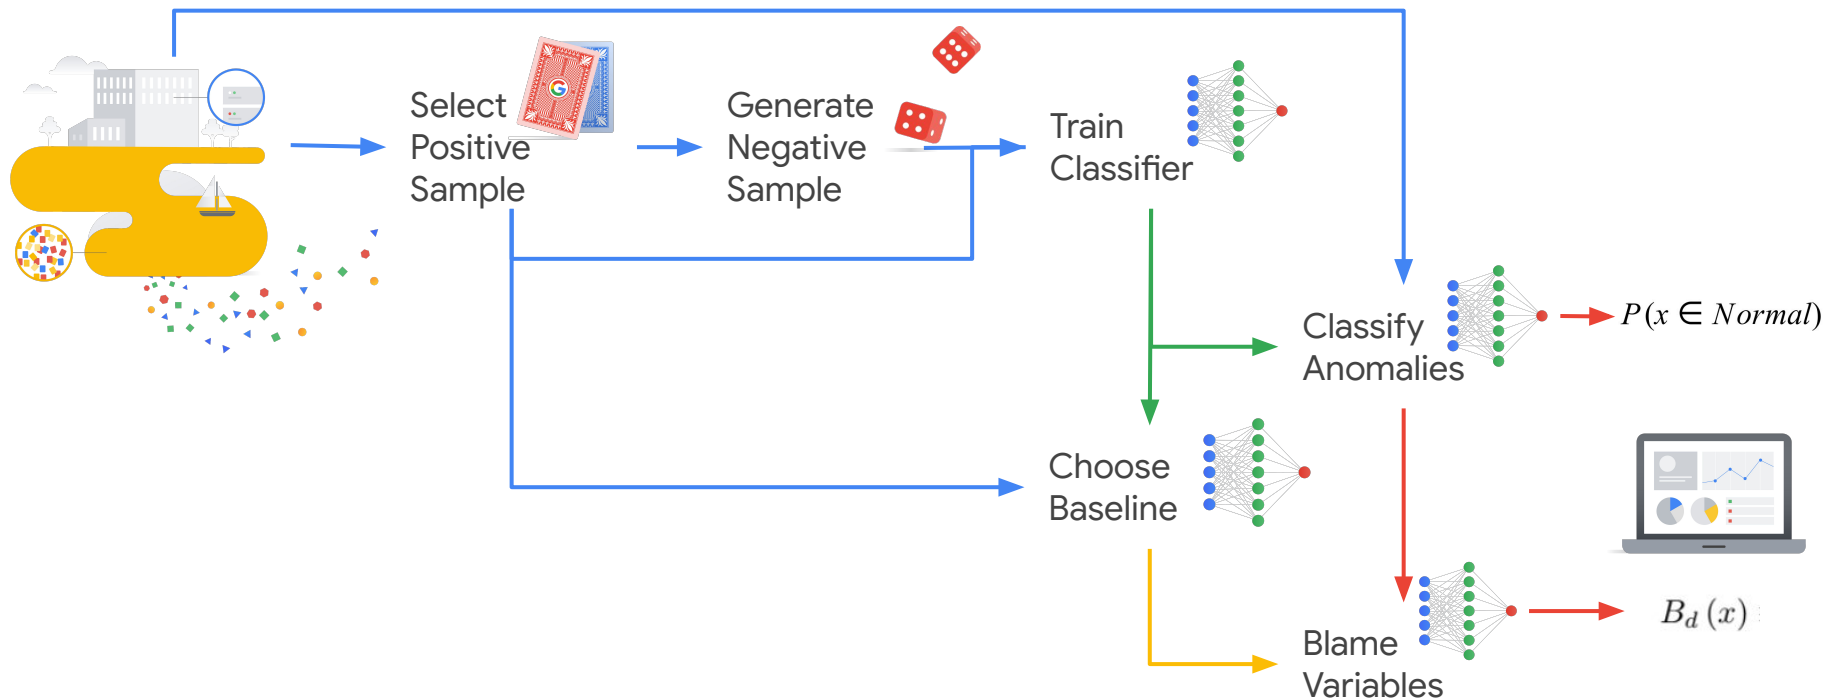
Comparing AI Explainability Methods

	Contrastive	Complete	Sensitive	Proportional
LIME Ribiero2016	-	+	-	-
SHAP Lundberg2017	-	+	+	+
OC-DTD Kaufmann2020	-	+	-	-
LRP Bach2015	-	+	-	-
IG Sundararajan2017	+	+	+	+

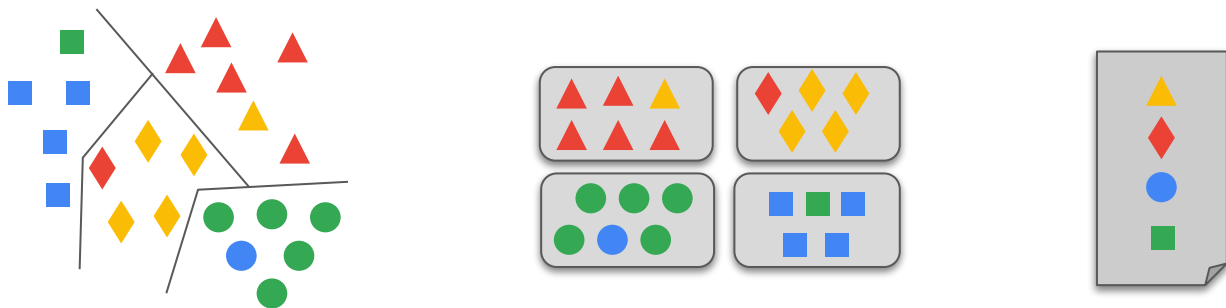
Anomaly Detection with Explanation



Anomaly Detection Pipeline with Interpretability



Distributed Anomaly Detection Framework



Segment the population into homogeneous cohorts.

Resegment periodically.



Launch independent Anomaly Detection instances on borg for each cohort.

Update the cohort membership when population changes.



Aggregate the results from all the instances.

Rank order the anomalous members by severity.



Publish top N most anomalous devices via PubSub topic to consumers.

Devices in the Cohort

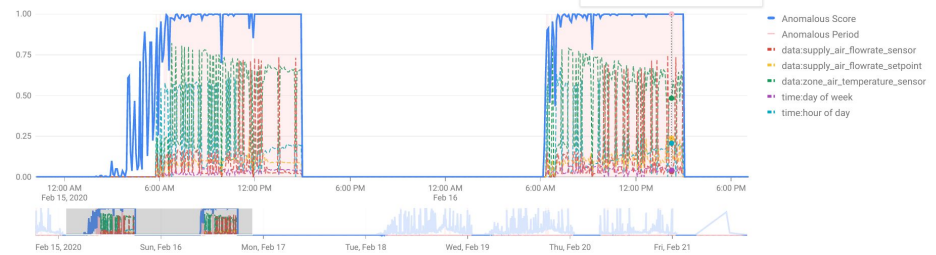
Device	Building	External Id
6 VAVCO 2-4-28	US-MTV-44	2643297080993457.CDM
7 VAVCO 1-1-9	US-MTV-44	2601703867660998.CDM
8 VAVCO 2-4-41	US-MTV-44	2634001211856729.CDM
9 CO 4-4-20	US-MTV-CL4	2547155771674383.CDM

2020-02-16 14:15:00

- Anomalous Score: 1.00
- data:zone_air_temperature_sensor: 0.48 | 65.50 (Expect: 72.40)
- data:supply_air_flowrate_setpoint: 0.24 | 2160.00 (Expect: 2700.00)
- time:hour of day: 0.21 | 14.00 (Expect: 14.00)
- time:day of week: 0.04 | 6.00 (Expect: 2.00)
- data:supply_air_flowrate_sensor: 0.03 | 273.86 (Expect: 188.65)

Device Name
428
1-9
441

Anomalous Scores



Flowrate fields



Project Freon to detect leaky Air Conditioners

Detecting Leaky Air Conditioners



Refrigerants

- 1,000s times more potent contributor to global warming than CO₂
- Fastest growing GHG emission, at 4.6% per year in the last decade
- Reducing emission starts with early leak detection

Google Research

Sensors

Temperature (F)

- 1 Suction Temperature
- 2 Discharge Temperature
- 3 Liquid Temperature
- 4 Liquid Saturation Temperature
- 5 Vapor Saturation Temperature
- 6 Subcooling Temperature
- 7 Superheat Temperature
- 8 Discharge Air Temperature
- 9 Return Air Temperature
- 10 Mixed Air Temperature
- 11 Zone Air Temperature
- 12 Outside Air Temperature

Pressure (psig)

- 1 Suction Pressure
- 2 Discharge Pressure
- 3 Liquid Pressure

Current (A)

- 1 Discharge Fan Current
- 2 Compressor Current
- 3 Condensing Fan Current

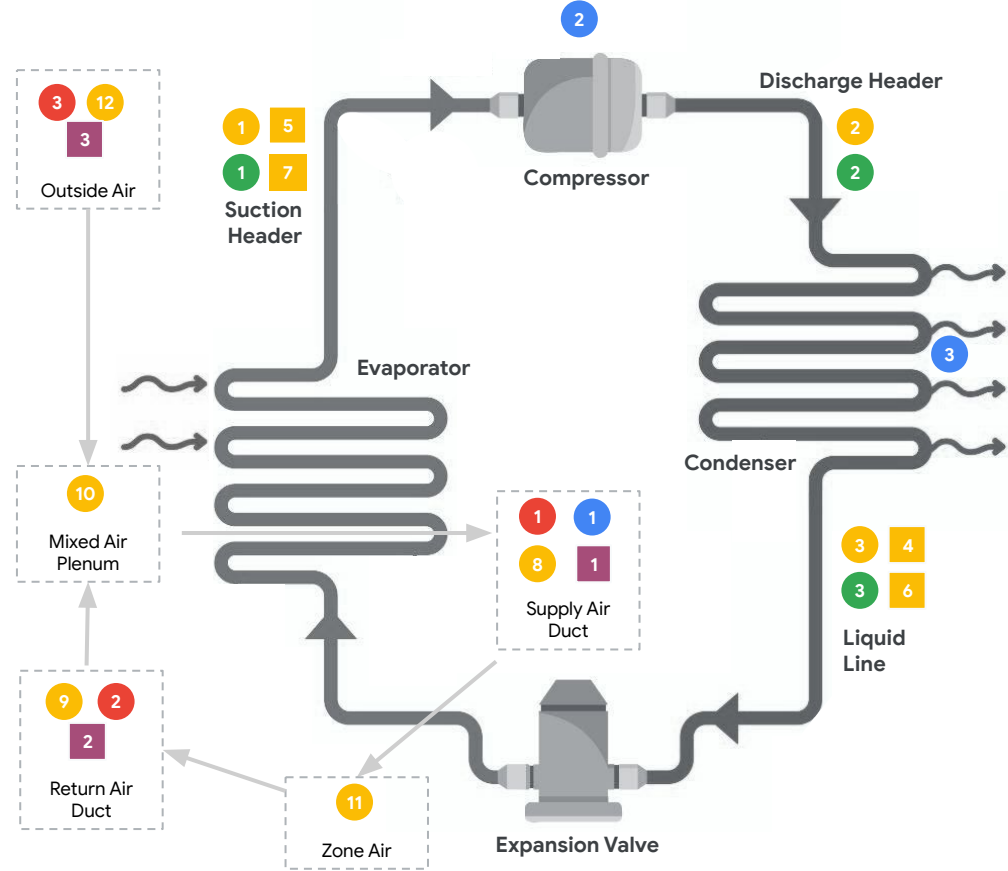
Relative Humidity (%)

- 1 Discharge Air Relative Humidity
- 2 Return Air Relative Humidity
- 3 Outside Air Relative Humidity

Enthalpy (BTU/lbm)

- 1 Discharge Air Enthalpy
- 2 Return Air Enthalpy
- 3 Outside Air Enthalpy

- Sensor (captures data not typically measured)
- Calculated data



Project Freon Test Sites





Technical diagrams and wiring schematics pinned to the left side of the cabinet.

⚠ DANGER
High Voltage
Electrical Shock Hazard

20
100
150

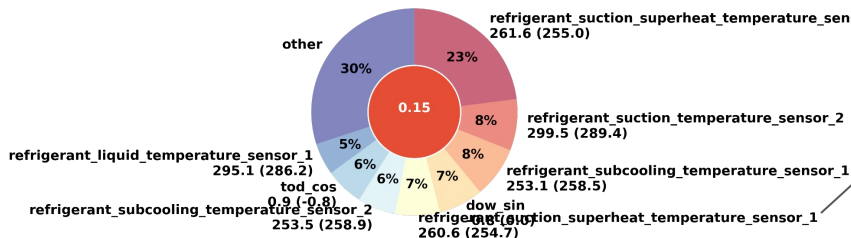
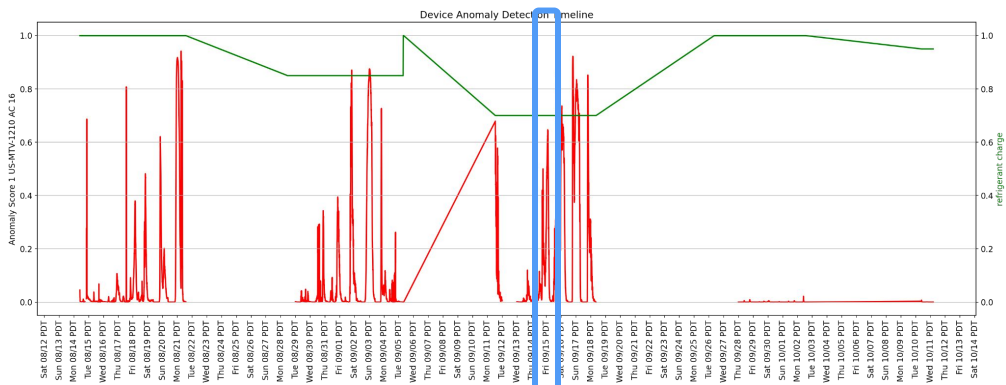


AG 6A

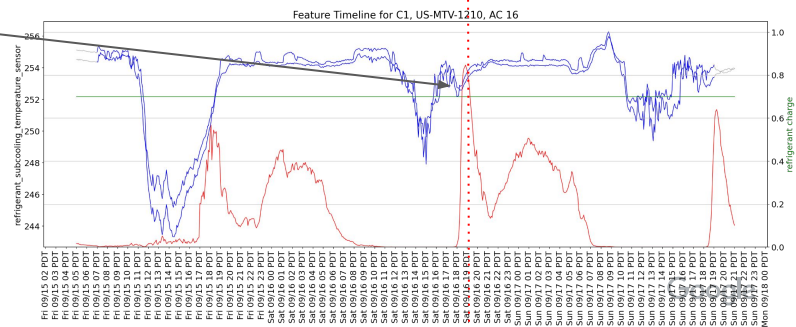
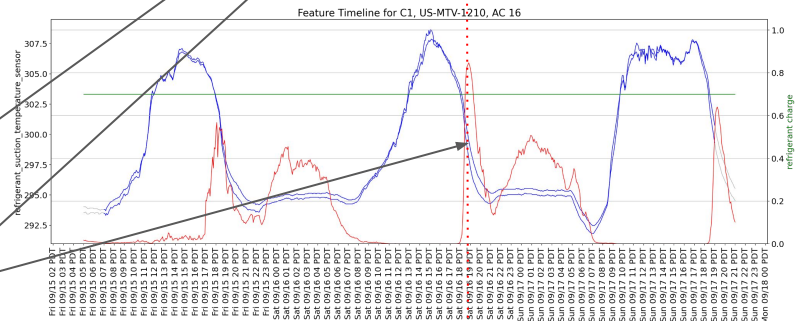
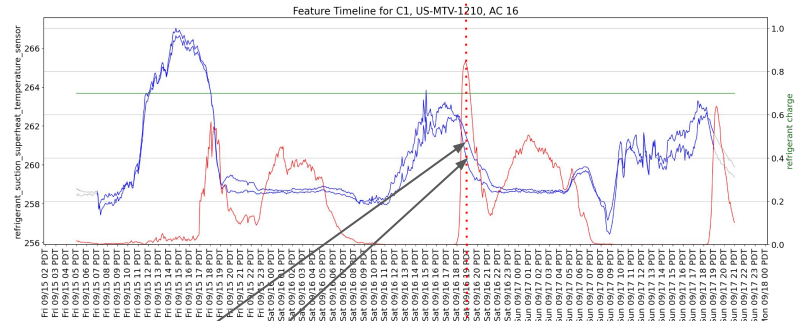
1228
14500
CL550

此處有線路，請
注意，勿動。

Explainability Analysis Example (1)



High refrigerant suction superheat temperature in circuits 1, 2; high refrigerant suction temperature circuit 2; low refrigerant subcooling temperature circuit 1, (on the descending edge of the duty cycle).



Qualitative Assessment of Detection Performance

Weather Conditions

Cooling Degree Days

Hot >8	Detectable	Detectable	Detectable
Mild 4 - 8	Detectable	Marginal	Marginal
Cool 0-3	Marginal	Marginal	Undetectable
	Severe < 75%	Significant 75% - 85%	Initial > 85%

Leakage
Charge Level



Intelligent Diagnostics provide root-cause analysis and recommend fixes

Diagnosis and Causal Reasoning

Diagnosis: An investigation or analysis of the **cause** or nature of a condition, situation, or problem

Merriam-Webster Dictionary

Diagnose: identify the nature of (an illness or other problem) by examination of the symptoms.

Oxford Languages

Actual Causal Reasoning is a fundamental step in forming a diagnosis



LLMs and Actual Causal Reasoning

LLMs outperform Structural Causal Models on the CRASS Counterfactual Reasoning benchmark and **6% lower than human annotators**

CRASS Counterfactual reasoning benchmark example

Premise	Counterfactual Question	Multiple-choice Answers
A bird lands in a forest.	What would have happened if a plane had landed in a forest?	(a) The plane would have crashed. (b) Everything would have been fine. (c) The plane would have landed safe and sound. (d) In a forest, you will find lots of planes.

LLMs and Reasoning about Normality

BIG-Bench Hard Causal Judgements Task (example)		
Passage and Question	Correct Answer	LLM rationale
<p>Lauren and Jane work for the same company. They each need to use a computer for work sometimes. Unfortunately, the computer isn't very powerful. If two people are logged on at the same time, it usually crashes. So the company decided to institute an official policy.</p> <p>It declared that Lauren would be the only one permitted to use the computer in the mornings and that Jane would be the only one permitted to use the computer in the afternoons.</p> <p>As expected, Lauren logged on the computer the next day at 9:00 am. But Jane decided to disobey the official policy. She also logged on at 9:00 am. The computer crashed immediately.</p> <p>Did Jane cause the computer to crash?</p>	<p>Jane logging in at 9:00 am is an abnormal event because she followed policy</p>	<p>Jane's decision to log onto the computer at 9:00 am despite the official policy is abnormal, as it intentionally violates the company's policy put in place to prevent computer crashes. This action is unexpected and goes against the established social and workplace rules.</p> <p>(correct)</p>

Large Language Models achieve nearly 70% accuracy scores on generative causal reasoning tasks.

Intelligent Diagnostics

Detection



Continuous Monitoring of device telemetry, a stream of multidimensional time-stamped observations from a fleet of devices.

Explanation



Anomaly Detection trains anomaly detection models on unlabeled telemetry, and scores each observation with an anomaly score.

Interpretation



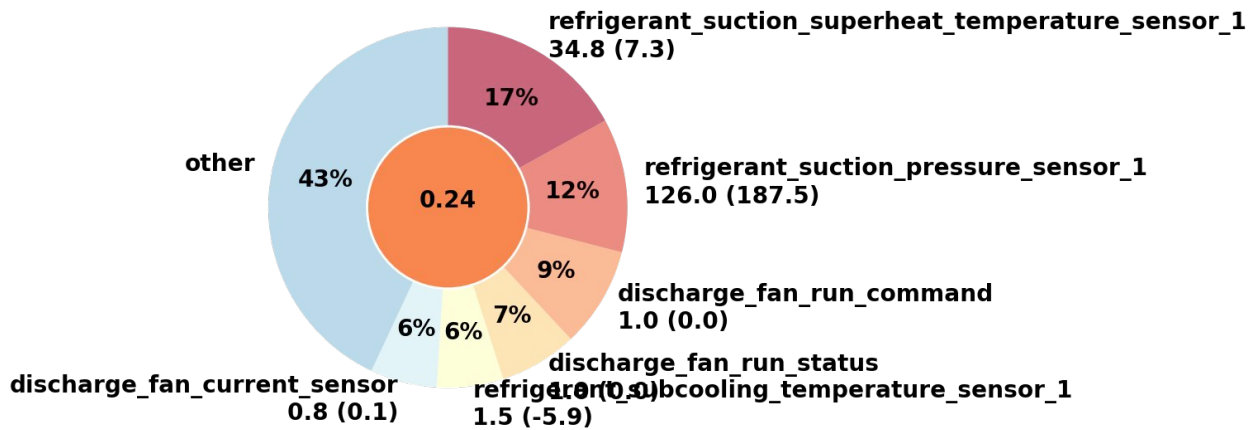
Explainable AI provides rich explanations of each anomaly, including

- Attributions
- Contrastive Normal



Large Language Models combine the explained anomaly with knowledge base to provide a diagnosis and suggested treatments.

From Explanation to Prompt



Consider an anomalous air conditioner. The refrigerant superheat temperature has 17% importance, observed is 34.8 degrees celsius, but expected to be 7.1 degrees celsius. The refrigerant suction pressure has 12% importance and is observed 126.0 psi, but expected 187.5 psi. Provide a diagnosis of the problem with justification.

Gemini-based Diagnosis (Problem Understanding)

The problem

The air conditioner is low on refrigerant.

Justification

The superheat temperature is a measure of how much heat is left in the refrigerant after it has passed through the evaporator coil. The higher the superheat temperature, the less heat is being removed from the air. In this case, the superheat temperature is 34.8 degrees Fahrenheit which is much higher than the expected value of 7.1 degrees Fahrenheit. This indicates that there is not enough refrigerant in the system to remove the heat from the air.

The suction pressure is also a measure of the amount of refrigerant in the system. The lower the suction pressure, the less refrigerant is in the system. In this case, the suction pressure is 126.0 psi, which is much lower than the expected value of 187.5 psi. This further confirms that there is not enough refrigerant in the system.

Gemini-based Diagnosis (Treatment Suggestion)

The solution

The solution to this problem is to recharge the air conditioner with refrigerant. This can be done by a qualified HVAC technician.

Other possible causes

There are a few other possible causes of the problem, but they are less likely. These include:

- A dirty evaporator coil.
- A malfunctioning expansion valve.
- A restriction in the refrigerant line.

If the superheat temperature and suction pressure are still abnormal after the air conditioner has been recharged, then one of these other causes may be the problem. A qualified HVAC technician will be able to diagnose the problem and make the necessary repairs.

Parting Thoughts

AI has an enormous potential to offset the climate burden from buildings

We must work collaboratively to leverage the potential of AI

Google is willing to partner as you unlock the opportunity of AI to increase efficiency and reduce emission



