

Realtime Operation Monitor

Wind direction S'E

Wind speed 2.1 m/s

Hor. Radiation 531.0 W/m²

Slop Radiation 521.0 W/m²

Gen. Efficiency

117.1 kW

KEPCO

W/T

1.2 kW



PV
20 kWp



Battery
50 kWh



6.5 kW

Fare Information

Today's Peak

2234

kW

tCo2/Exp

Won/Exp

kWh

kWh

Wh

MWh

Situational Awareness

State Estimation, Predictive Maintenance, & Forecasting

6.S893: AI for Climate Action (Power & Energy Systems)

Spring 2026

Supply Power

139.5 kW

Power Consumption

139.5 kW

KEPCO 117.1 kW

BATTERY

6.5 kW

Light 20.0 kW

Heating

1.8 kW

PV 17.2 kW

WIND TURBIN

0.0 W

Outlet 14.0 kW

EV Charge

0.0 kW

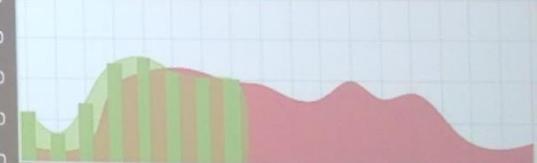
Cooling 4.0 kW

Etc (LG U+)

23.0 kW

250

Current Demand Exp. Demand



(h)

Administrivia

Paper presentations

- Format: See <https://ai4climateaction.github.io/assignments/>
 - Likely you will want to spend most of the time on (a) summary of the application and its broader context, (b) summary of the method, and (c) critique. However, different papers may warrant different time distributions.
 - We will not be pedantic about the exact format – this is less about “quizzing” you on the paper, and more about creating a good open space for discussion.
- Signups: Some papers still need presenters – if you are currently pair-presenting a paper, please consider moving one of your signups to one of these unclaimed papers.
- Slide submission: Submit your slides on Canvas ahead of class.

Reflections: Reminder that these are due every Tuesday during the fork (submit via Canvas).

Final projects: Rubric out today. Project “groups” of 1-2 people; please sign up by next week.

Outline

Wrap-up: Overview of ML for power & energy, and important considerations

Background: AC grids and power flow

Overview of situational awareness applications, challenges, and methods

- State estimation
- Predictive maintenance
- Near-term forecasting

Challenge highlight: Distribution shift

Outline

Wrap-up: Overview of ML for power & energy, and important considerations

Background: AC grids and power flow

Overview of situational awareness applications, challenges, and methods

- State estimation
- Predictive maintenance
- Near-term forecasting

Challenge highlight: Distribution shift

Last time: Introduction to power & energy systems

🔧 **Energy systems transformation** is needed due to climate change mitigation, adaptation, and sustainable development

- Key axes: Operations and planning

⚡ **Power grids are “the world’s largest machine”**

- Requires *constantly* maintaining energy balance, including decision-making across many timescales (sub-seconds to decades)
- Power systems are diverse, rapidly changing, and increasing in complexity
- Many different stakeholders: Regulators, operators, suppliers, consumers, etc.

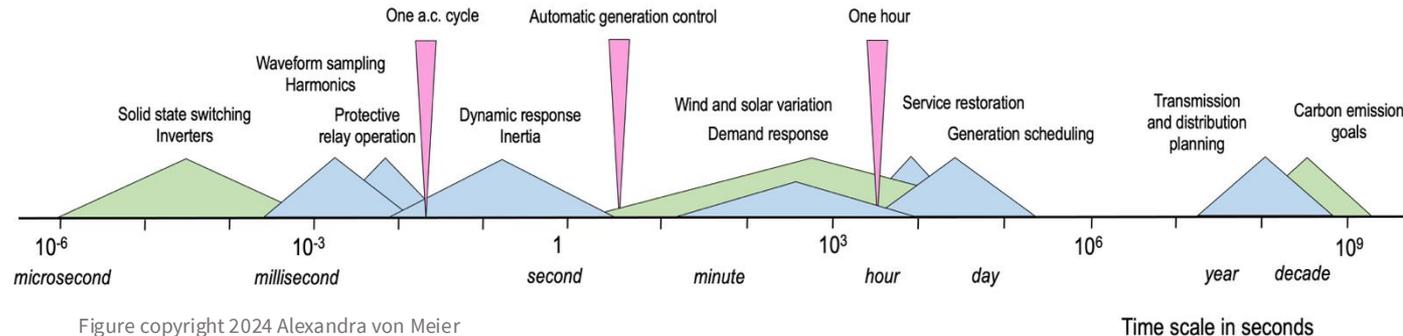


Figure copyright 2024 Alexandra von Meier

Time scale in seconds

Overview of ML applications in energy systems

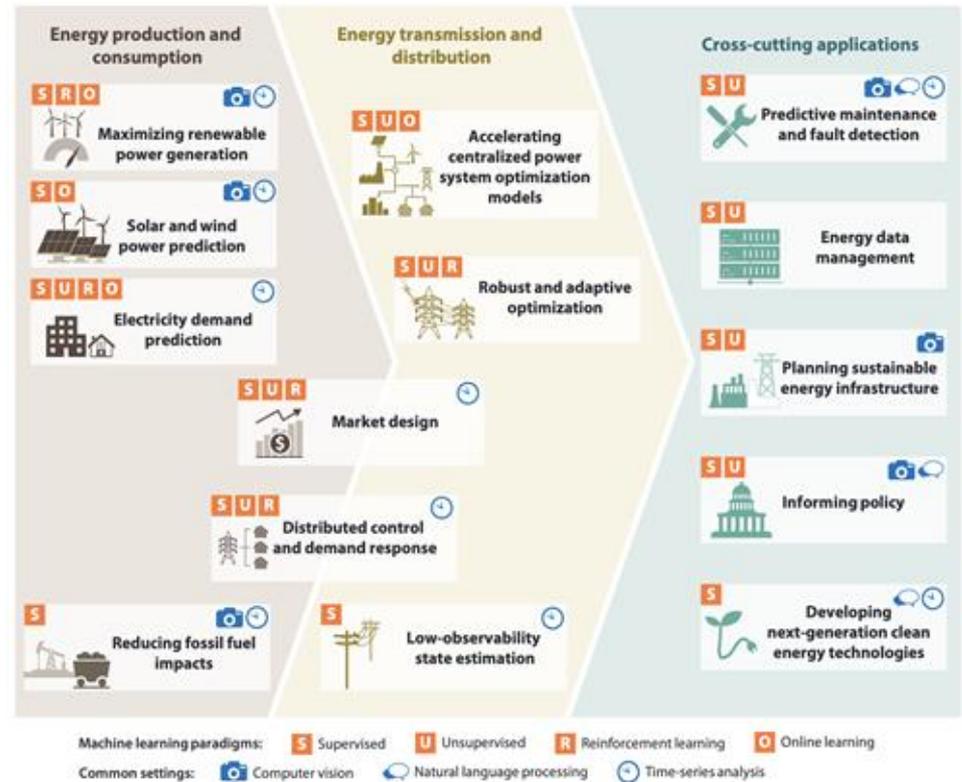
Operations

- Situational awareness
- Prediction
- Optimization & control

Planning

- Infrastructure mapping
- Speeding up simulations
- Scenario generation

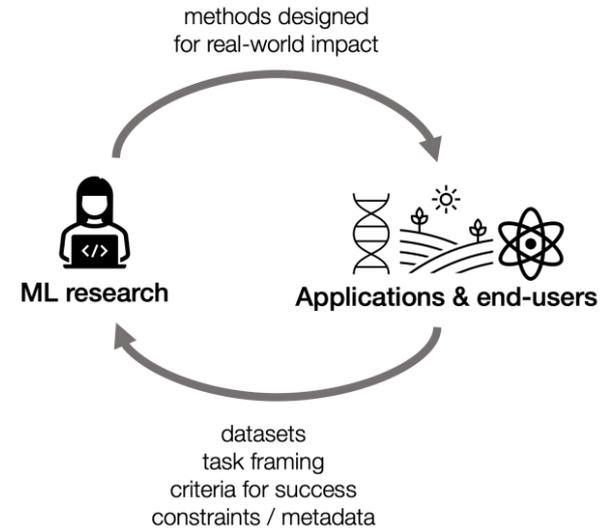
Enablers – facilitating innovation, policy & markets, and management of data



Considerations for ML in power & energy systems

Different requirements for ML models and their outputs, depending on the context

- Accuracy/solution quality (better than SOTA)
- Safety & physical feasibility
- Robustness
- Interpretability, explainability, & auditability
- Uncertainty quantification
- Fast running time
- Hardware integration
- Data efficiency
- Generalizability
- Multi-agent and human-in-the-loop
- Privacy preservation
- Usability and accessibility
- Meeting regulatory standards



David Rolnick, Alan Aspuru-Guzik, Sara Beery, Bistra Dilkina, Priya L. Donti, Marzyeh Ghassemi, Hannah Kemer, Claire Monteleoni, Esther Rolf, Milind Tambe, Adam White.
"Position: Application-driven innovation in machine learning." *ICML 2024*.

Responsible AI in power & energy systems

Mitigating biases in data and models

- E.g., Power infrastructure data: Geographic disparities in availability
- E.g., Weather models: Calibration may be optimized for particular regions

Improving trustworthiness and accountability

- Safety and robustness: Critical in, e.g., power system operations
- Interpretability, auditability, and human-in-the-loop approaches: Critical in, e.g., policymaking contexts

Centering equity and climate justice

- Centering diverse stakeholders: E.g., in industrialized vs. emerging economies
- Avoiding centralization: Democratized capacity and compute, digital divide
- Avoiding digital colonialism: E.g., smart meters, analysis of remote sensing data

Accounting for potential “dual use”

Outline

Wrap-up: Overview of ML for power & energy, and important considerations

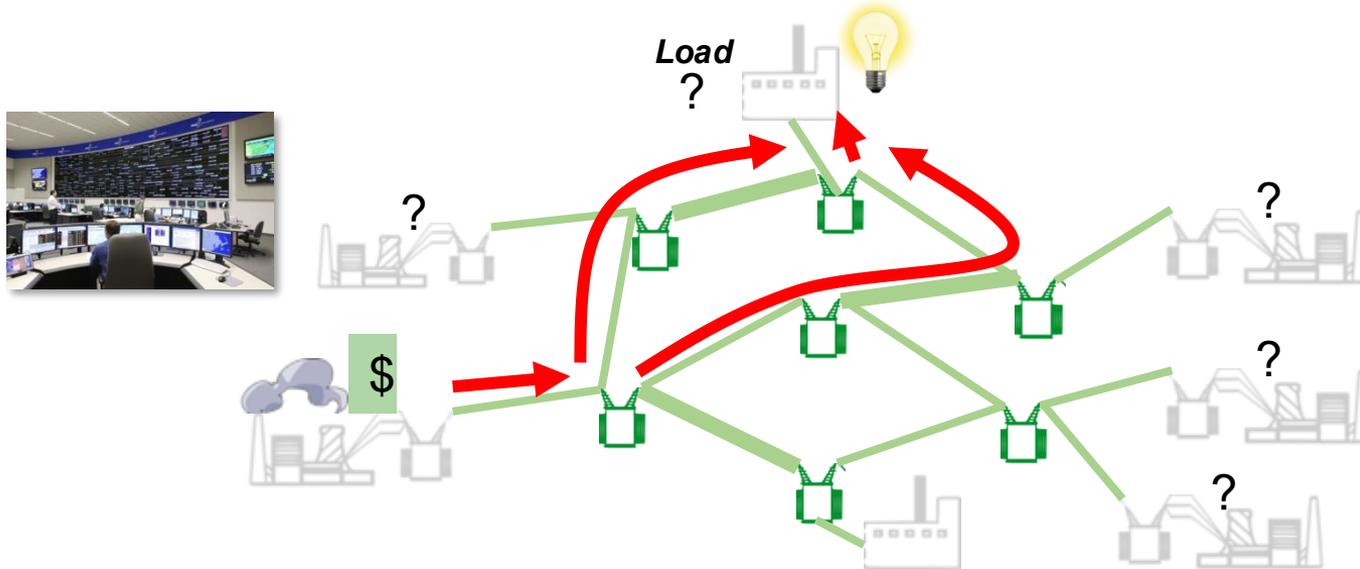
Background: AC grids and power flow

Overview of situational awareness applications, challenges, and methods

- State estimation
- Predictive maintenance
- Near-term forecasting

Challenge highlight: Distribution shift

Grid as a network of generators and loads



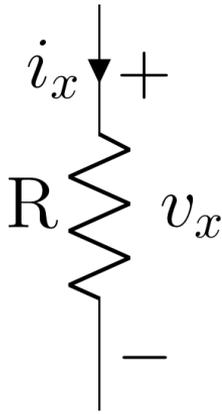
Generators and loads are at **buses (nodes)**

Buses are connected by **power lines** (edges)

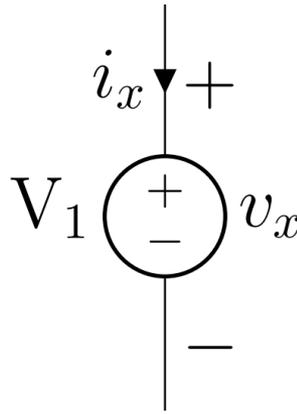
DC circuit: Basic behavior & components

Circuit behavior is described by:

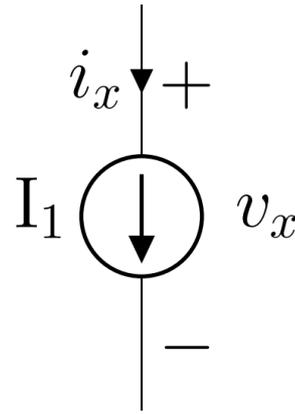
- Voltage (V): Difference in electrical potential between two points. Described **across** a component.
- Current (A): Rate of charge flowing **through** a component.



Ohm's Law: $v_x = i_x R$



$v_x = V_1$

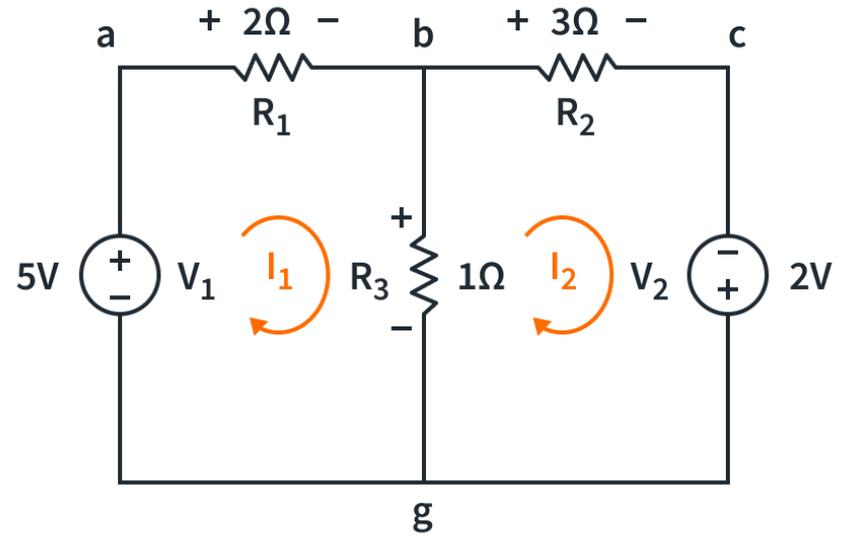


$i_x = I_1$

Kirchoff's voltage and current laws

Kirchoff's Voltage Law (KVL): The sum of voltages around any closed loop is zero

Kirchoff's Current Law (KCL): The sum of current flowing into a node must equal the sum of current flowing out of the node.



Alternating current (AC)

Direct current (DC): Polarity remains the same

- Potential always positive on one side, negative on the other
- Current always flows in same direction

Alternating current (AC): Polarity reverses/oscillates

- Voltage and current directions reverse at *AC frequency* (60 Hz in US, 50 Hz in many other places)

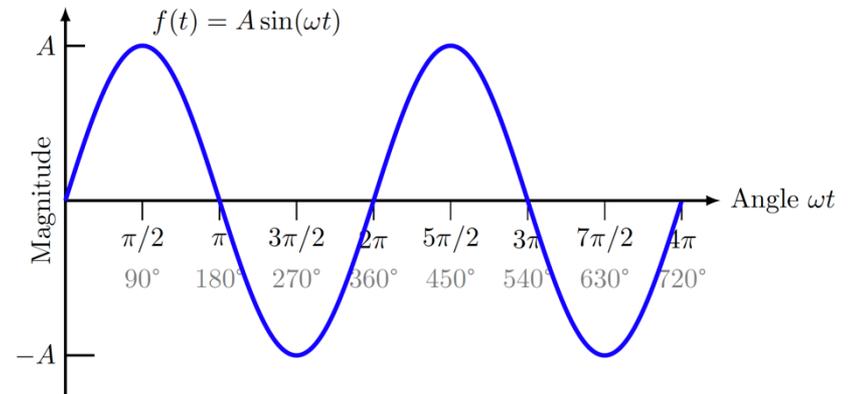


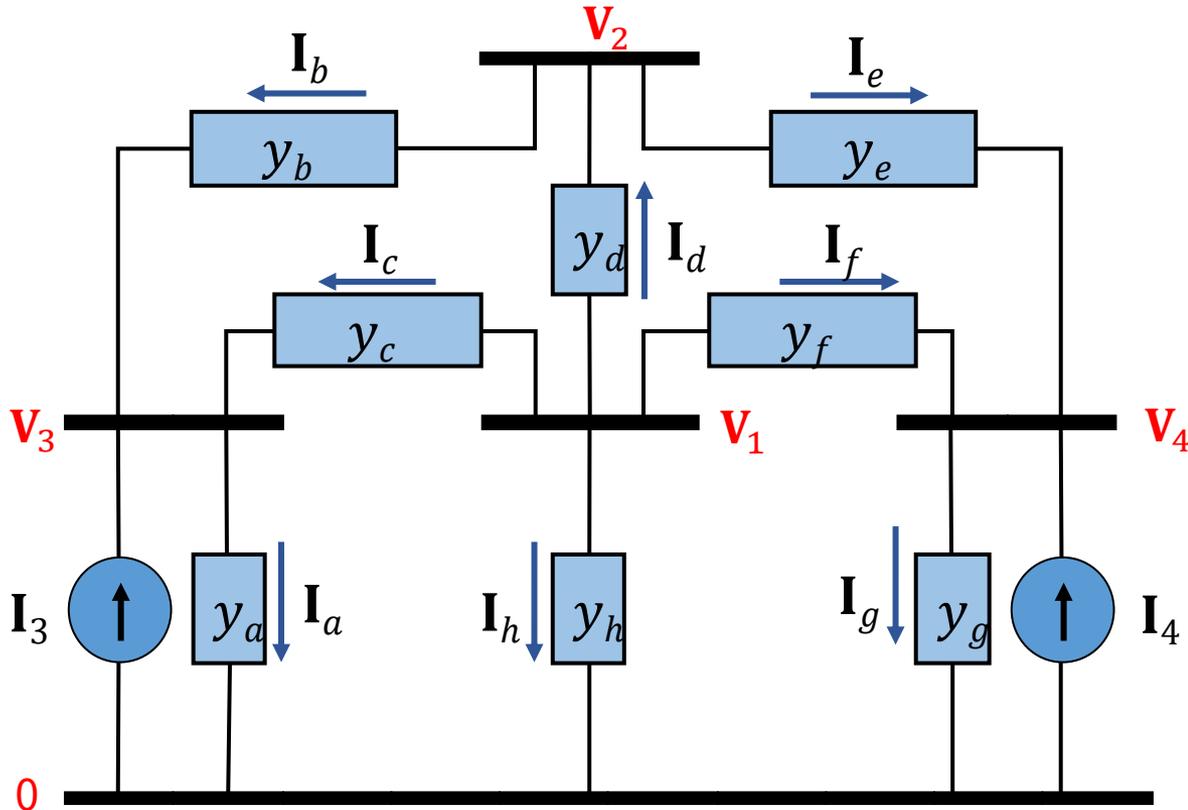
Figure copyright 2024 Alexandra von Meier

Quantities and relationships in AC vs. DC

Quantity	DC System	AC System
Current (A)	$i(t) = I_{\max}$	$i(t) = I_{\max} \cos(\omega t + \phi_I)$
Voltage (V)	$v(t) = V_{\max}$	$v(t) = V_{\max} \cos(\omega t + \phi_V)$
Impedance (Ω)	R	$Z = R + jX$
Admittance (S)	$G = 1/R$	$Y = 1/Z = G + jB$ (note: $G \neq 1/R$ in general)
Power (VA, W, VAR)	P	$S = P + jQ$

Equation	DC System	AC System
Power formula	$P = IV$	$p(t) = i(t)v(t)$ $S = VI^*$ [phasor notation]
Ohm's Law	$V = IR$ or $I = V/R$	$V = IZ$ or $I = VY$ [phasor notation]

Analyzing large-scale AC networks



Ohm's Law

$$I_a = y_a(V_3 - 0)$$

$$I_b = y_b(V_2 - V_3)$$

$$I_c = y_c(V_1 - V_3)$$

$$I_d = y_d(V_1 - V_2)$$

$$I_e = y_e(V_2 - V_4)$$

$$I_f = y_f(V_1 - V_4)$$

$$I_g = y_g(V_4 - 0)$$

$$I_h = y_h(V_1 - 0)$$

KCL

$$I_c + I_d + I_f + I_h = 0$$

$$I_b - I_d + I_e = 0$$

$$I_a - I_b - I_c = I_3$$

$$I_g - I_e - I_f = I_4$$

→ Combine, rearrange, and put into matrix form

Nodal voltage equations

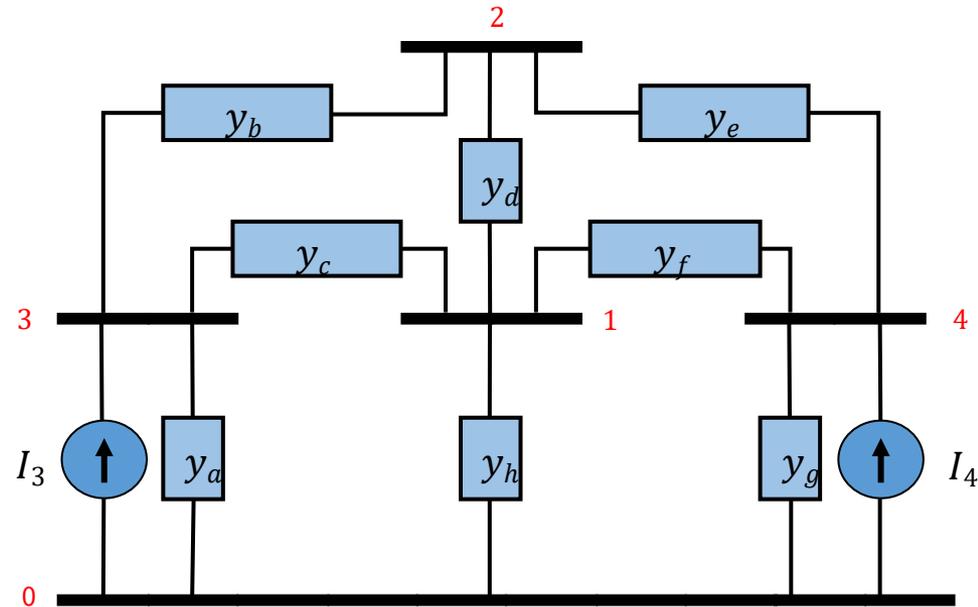
$$\underbrace{\begin{pmatrix} y_c + y_d + y_f + y_h & -y_d & -y_c & -y_f \\ -y_d & y_d + y_b + y_e & -y_b & -y_e \\ -y_c & -y_b & y_a + y_b + y_c & 0 \\ -y_f & -y_e & 0 & y_g + y_e + y_f \end{pmatrix}}_{\text{Admittance matrix } Y} \underbrace{\begin{pmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \end{pmatrix}}_{\text{Unknown nodal voltages}} = \underbrace{\begin{pmatrix} 0 \\ 0 \\ I_3 \\ I_4 \end{pmatrix}}_{\text{Source currents } I}$$

$YV = I$

$$\begin{pmatrix} y_c + y_d + y_f + y_h & -y_d & -y_c & -y_f \\ -y_d & y_d + y_b + y_e & -y_b & -y_e \\ -y_c & -y_b & y_a + y_b + y_c & 0 \\ -y_f & -y_e & 0 & y_g + y_e + y_f \end{pmatrix}$$

Structure of the admittance matrix:

- Rows and columns of Y correspond to nodes of the network
- Diagonal element (i, i) :
 - Sum of the admittances connected to node i
- Off-diagonal element (i, j) :
 - Minus admittance between nodes i and j
- Given the network data, we can build the admittance matrix directly



Nodal power balance

$$I = YV$$

Row k of this matrix equation expresses KCL at node k

Current injected at node k $\mathbf{I}_k = \sum_{m=1}^N Y_{km} \mathbf{V}_m$ Sum of the currents flowing in the branches connected to node k

N : number of nodes in the network

$$\mathbf{I}_k^* = \sum_{m=1}^N Y_{km}^* \mathbf{V}_m^*$$

Complex power injected at node k $\mathbf{S}_k = \mathbf{V}_k \mathbf{I}_k^* = \mathbf{V}_k \sum_{m=1}^N Y_{km}^* \mathbf{V}_m^*$ Sum of the complex powers flowing in the branches connected to node k

AC power flow equations

$$\mathbf{S}_k = \mathbf{V}_k \sum_{m=1}^N Y_{km}^* \mathbf{V}_m^* = V_k e^{j\varphi_k} \sum_{m=1}^N |Y_{km}| e^{-j\delta_{km}} V_m e^{-j\varphi_m}$$

Since $S_k = P_k + jQ_k$,

$$P_k = \sum_{m=1}^N V_k V_m |Y_{km}| \cos(\varphi_k - \varphi_m - \delta_{km})$$

$$Q_k = \sum_{m=1}^N V_k V_m |Y_{km}| \sin(\varphi_k - \varphi_m - \delta_{km})$$

for $k = 1, \dots, N$

P_k : Active power injection at bus k

Q_k : Reactive power injection at bus k

N : Number of buses in the network

V_k : Magnitude of the voltage at node k

V_m : Magnitude of the voltage at node m

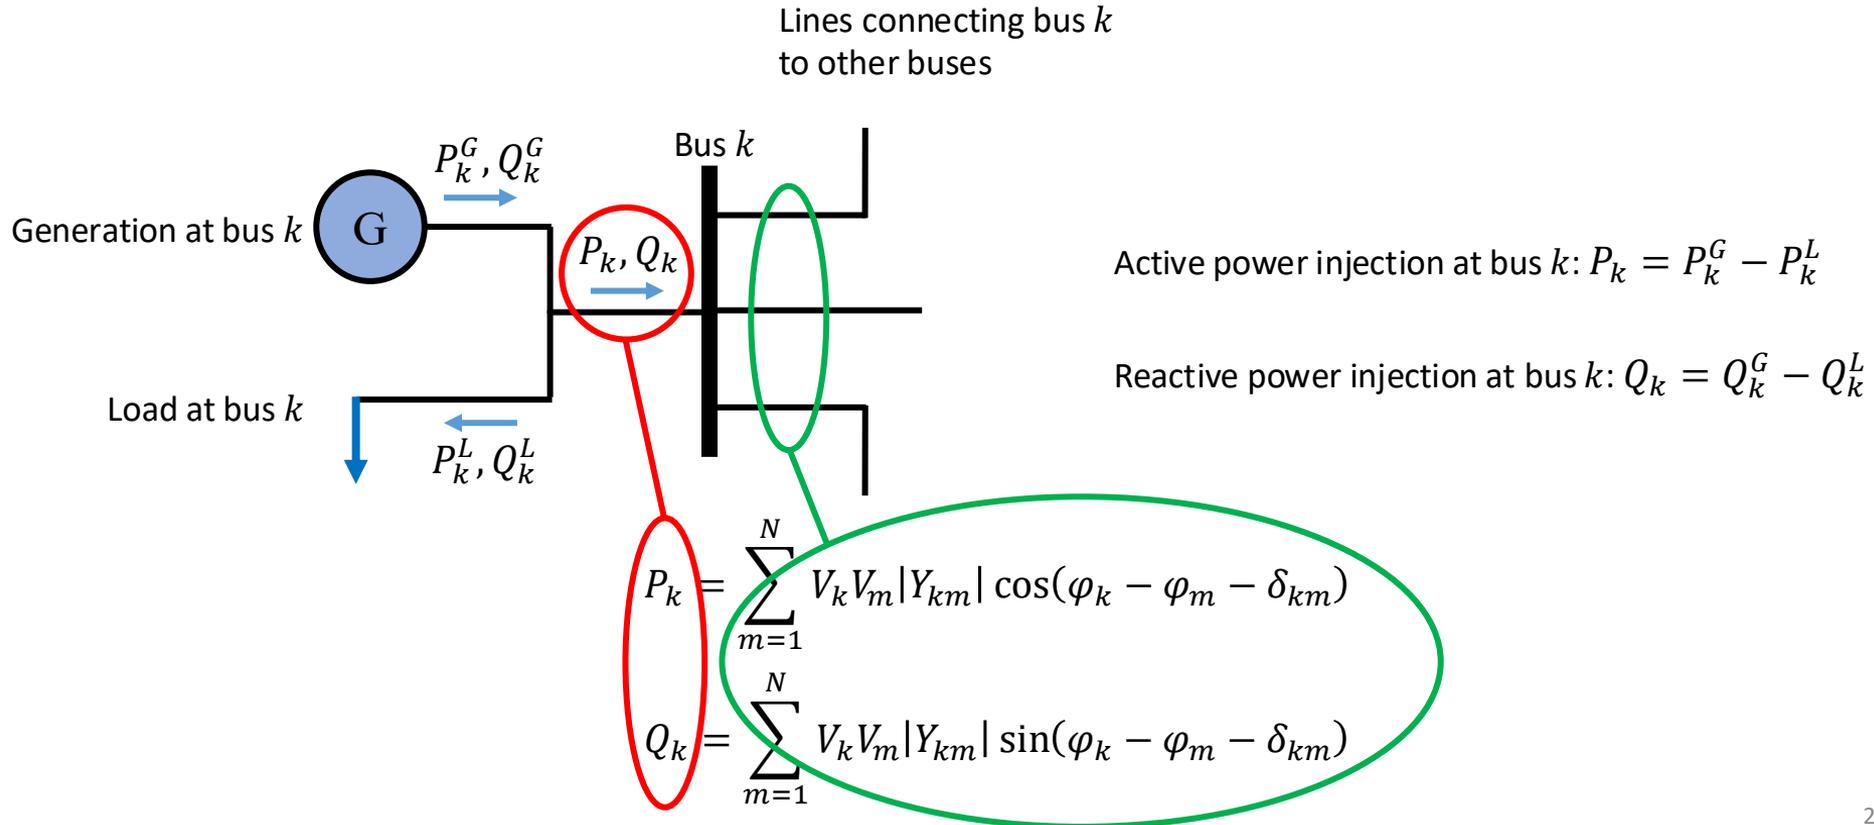
φ_k : Phase of the voltage at node k

φ_m : Phase of the voltage at node m

$|Y_{km}|$: Magnitude of (k, m) term of matrix Y

δ_{km} : Phase of (k, m) term of matrix Y

Physical interpretation of the power flow equations



Outline

Wrap-up: Overview of ML for power & energy, and important considerations

Background: AC grids and power flow

Overview of situational awareness applications, challenges, and methods

- **State estimation**
- **Predictive maintenance**
- **Near-term forecasting**

Challenge highlight: Distribution shift

State estimation

What is the **state** of the system (i.e., voltage magnitudes and angles at all buses)?

Easiest version: Given “enough” high-quality measurements of the system variables, just solve the AC power flow equations

Bus type	P_k	Q_k	V_k	φ_k
Load (PQ)	X	X		
Generator (PV)	X		X	
Slack / reference*			X	X

Challenge: Even “easy version” is more computationally expensive than desirable

⇒ Power flow approximations
Power flow linearizations
Learning surrogate models

Even more challenges in real systems

- Sensor noise (non-Gaussian) or data attacks
- Low observability (not “enough” measurements)
- Inconsistent data availability

⇒ Structural priors (e.g., “low rank”)
Ordinary least squares approaches
Physics-informed ML

Predictive maintenance

Monitor equipment health in real-time to predict potential failures before they occur

Things to monitor:

- Generators
- Batteries
- Lines, transformers, and switches
- Natural gas pipelines

Data modalities:

- Data from voltage/current sensors
- Visual data (e.g., from drones)
- Acoustic data
- Infrared data (thermal imaging)

Why is this challenging?

- Rare event detection is hard!
- Difficult to create models that are generalizable
- Shared challenges with state estimation (sensor noise, observability, inconsistent data)

Predictive maintenance: Supervised learning approaches

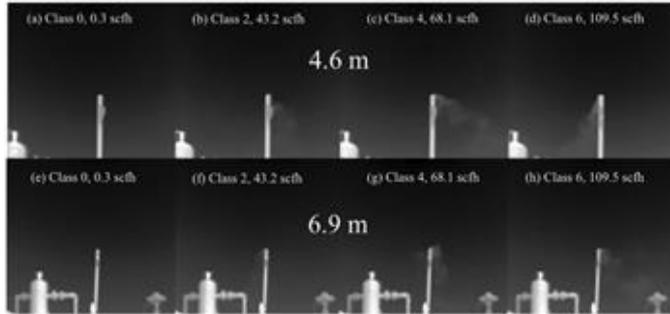


Image from: Wang, Jingfan, et al. "VideoGasNet: Deep learning for natural gas methane leak classification using an infrared camera." *Energy* 238 (2022): 121516.

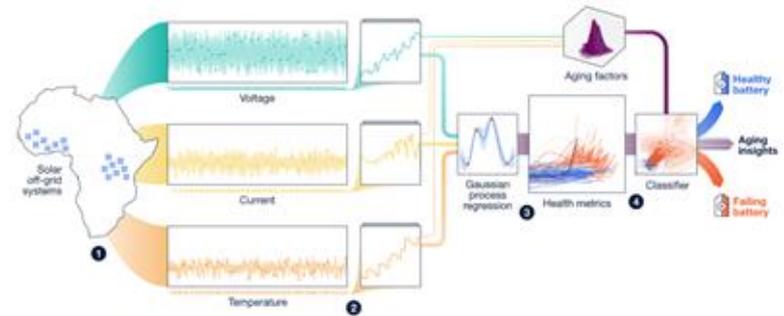


Figure from: Aitio, Antti, and David A. Howey. "Predicting battery end of life from solar off-grid system field data using machine learning." *Joule* 5.12 (2021): 3204-3220.



Figure from: [EPRI Journal](#)

Predictive maintenance: Non-supervised learning?

Self-supervised approaches: Methods that generate their own supervisory signals from unlabeled input

- E.g., Supervising on the structure of the problem (“mismatches” in power flow equations, fluid dynamics, other physical equations)
- E.g., *Contrastive learning* (minimize distance between generated “similar pairs” and maximize distance between “dissimilar pairs”)

Unsupervised approaches: Methods that focus on learning underlying structure in the unlabeled data

- E.g., Anomaly detection via autoencoders (learn to compress & reconstruct normal data, then flag data with high reconstruction data as anomalous)

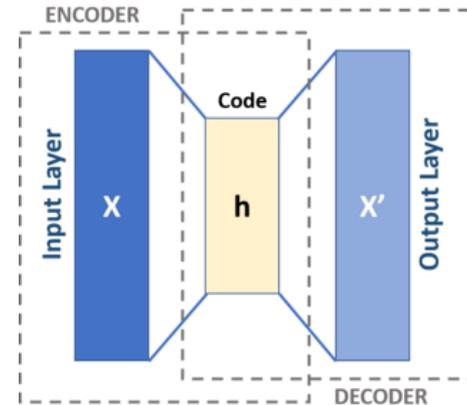


Image source: Wikipedia

Near-term forecasting

Predict demand, renewable energy production, prices, emissions, etc. in the near term (e.g., minutes, hours, or days ahead)

Data modalities (opportunity for multi-modal approaches):

- Historical time series data
- Physical model outputs (e.g., weather data)
- Satellite or aerial imagery (e.g., of clouds moving overhead)

Challenges:

- Creating predictions that are *probabilistic* or with *quantified uncertainty*
- Ensuring predictions are *interpretable, explainable, and/or auditable*
- Dealing with missing data
- Calibration with needs of decision-making process
- Distribution shift [particularly in areas with load growth]

Case study: Nowcasting in the UK

Project by Open Climate Fix & National Grid ESO

Demand: Used Temporal Fusion Transformer to reduce error by 2-3x for 30-min- and 48-hr-ahead national demand forecasts [CRDK+2021]

Solar PV: Used time series data, satellite data, and numerical weather predictions to reduce error by ~3x of 2-hr-ahead forecasts [K2022]

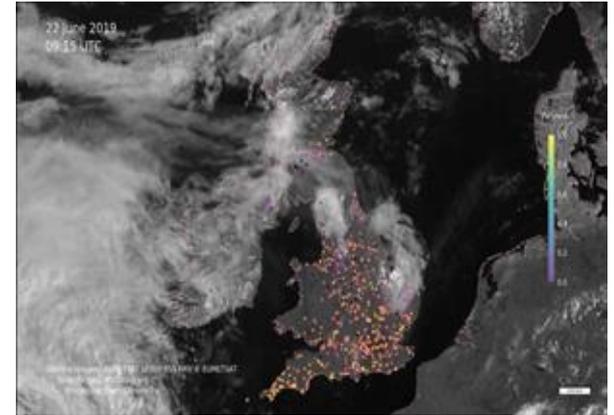


Image from Open Climate Fix. “Solar PV power and clouds over UK in January 2019.” YouTube video (2019). [Link here.](#)

[CRDK+2021] Clutton-Brock, Peter, Rolnick, David, Donti, Priya L., Kaack, Lynn, et al. Climate Change and AI: Recommendations for Government Action. Online (2021). [Link here.](#)

[K2022] Krasuka, Kasia. “Six Months into the Nowcasting Project Our Results Are Highly Promising” (2022). [Link here.](#)

Outline

Wrap-up: Overview of ML for power & energy, and important considerations

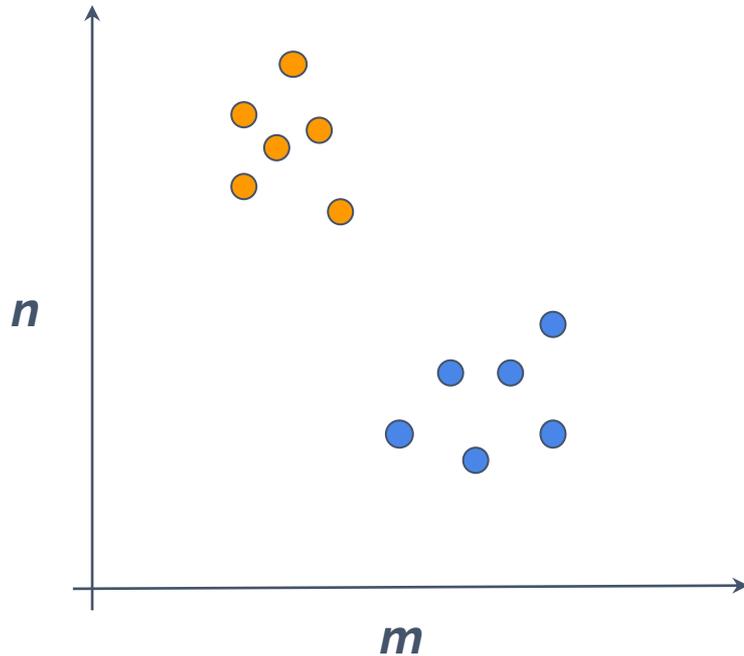
Background: AC grids and power flow

Overview of situational awareness applications, challenges, and methods

- State estimation
- Predictive maintenance
- Near-term forecasting

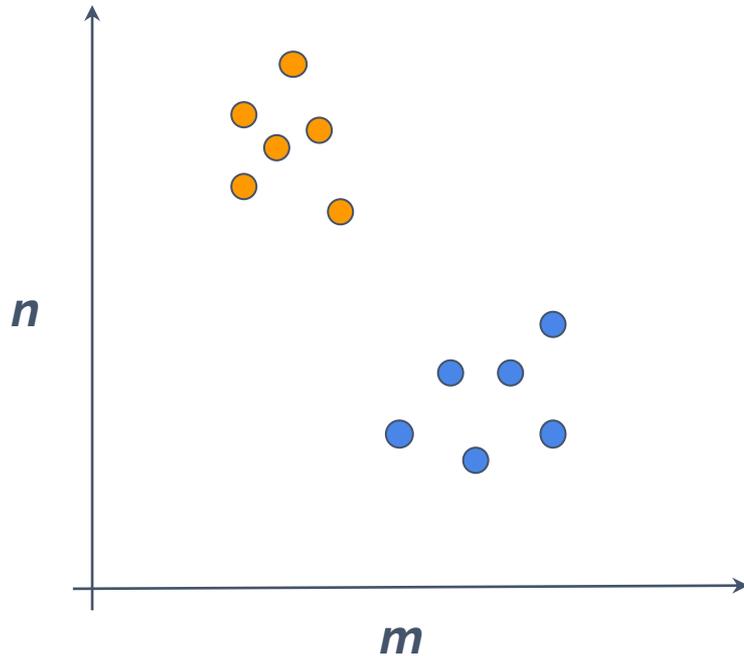
Challenge highlight: Distribution shift

Challenge highlight: Distribution shift



Let's consider a simple case: learning to categorize points on a plane.

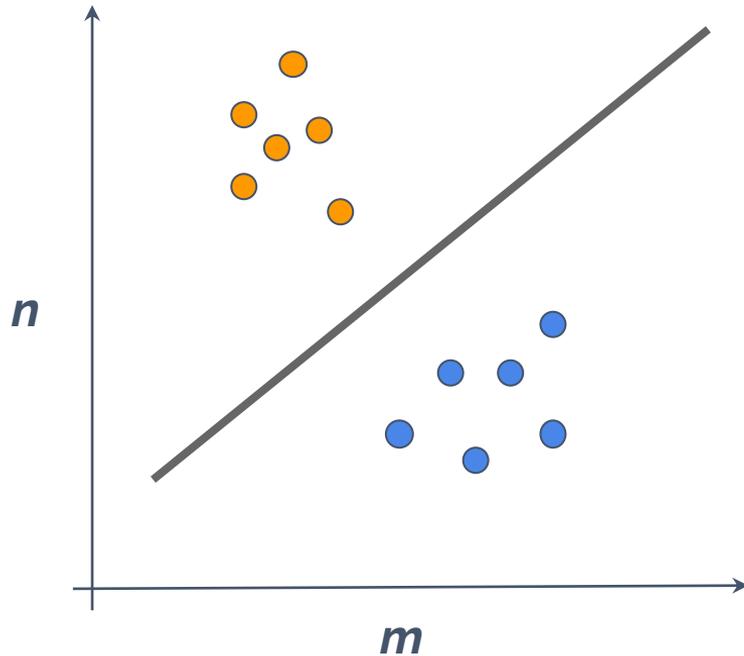
Challenge highlight: Distribution shift



Let's consider a simple case: learning to categorize points on a plane.

We assume we have **representative** data - aka **IID** - with labels

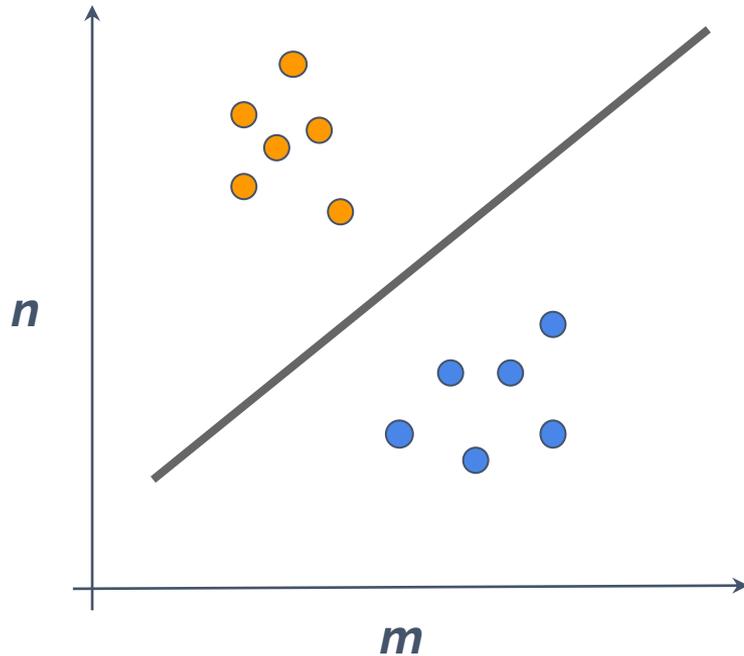
Challenge highlight: Distribution shift



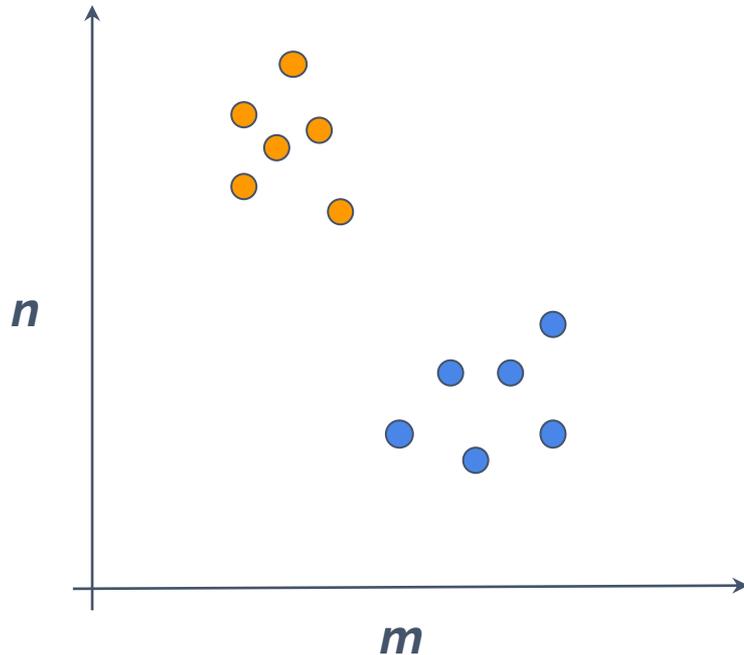
Using our labeled data, we learn this linear classifier

Challenge highlight: Distribution shift

How do we know if our model works?



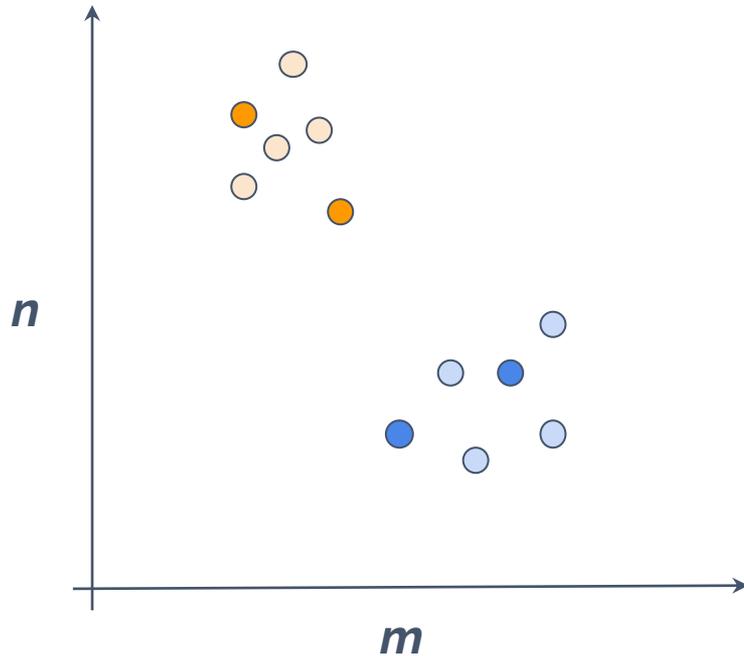
Challenge highlight: Distribution shift



How do we know if our model works?

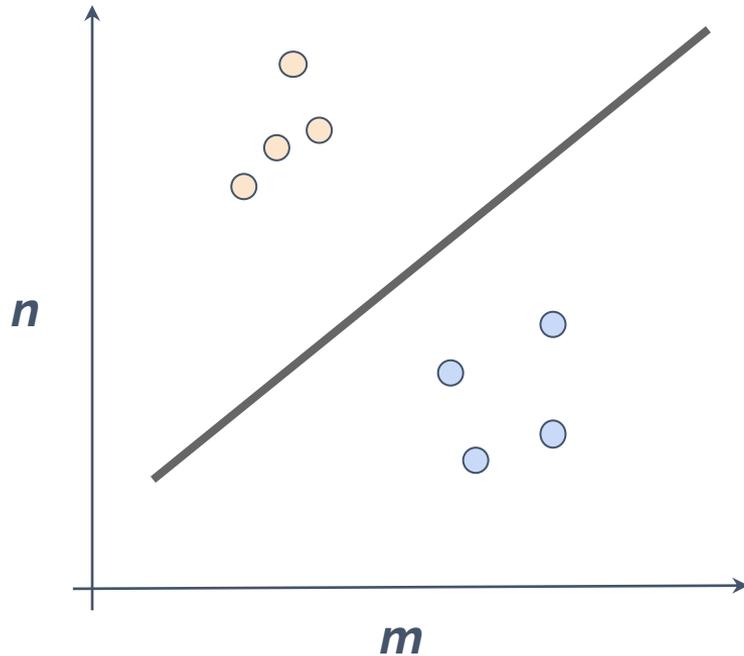
- We could get more data, label it, and test our model on the new data
- We can *pretend* we don't have some of the data during training and test the model on that data (most common)

Challenge highlight: Distribution shift



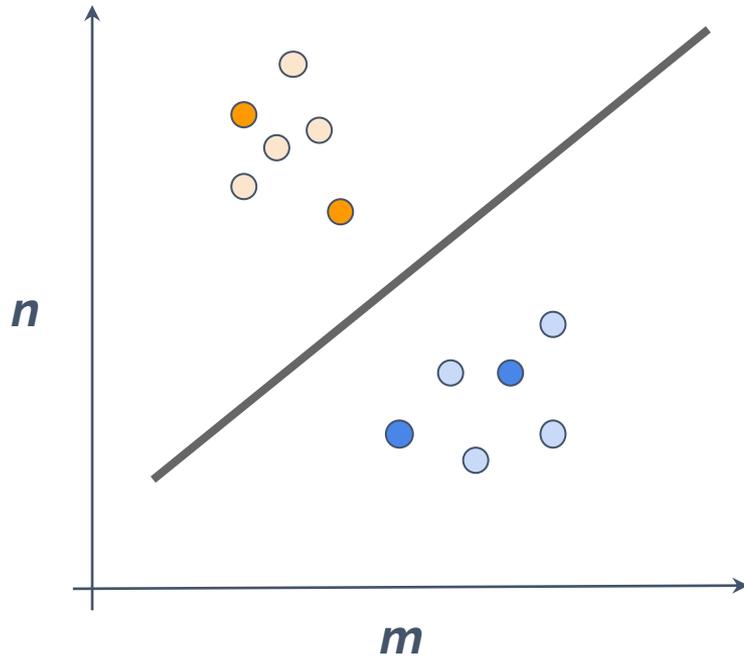
Since the data is assumed to be IID, we take a random subset to use as evaluation data

Challenge highlight: Distribution shift



Now we would learn our model using only the “training data”

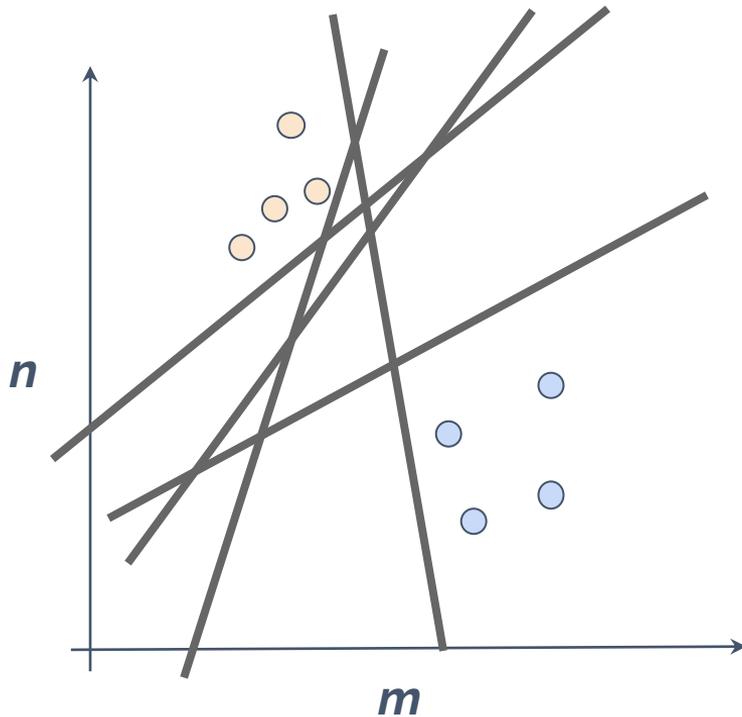
Challenge highlight: Distribution shift



Now we would learn our model using only the “training data”

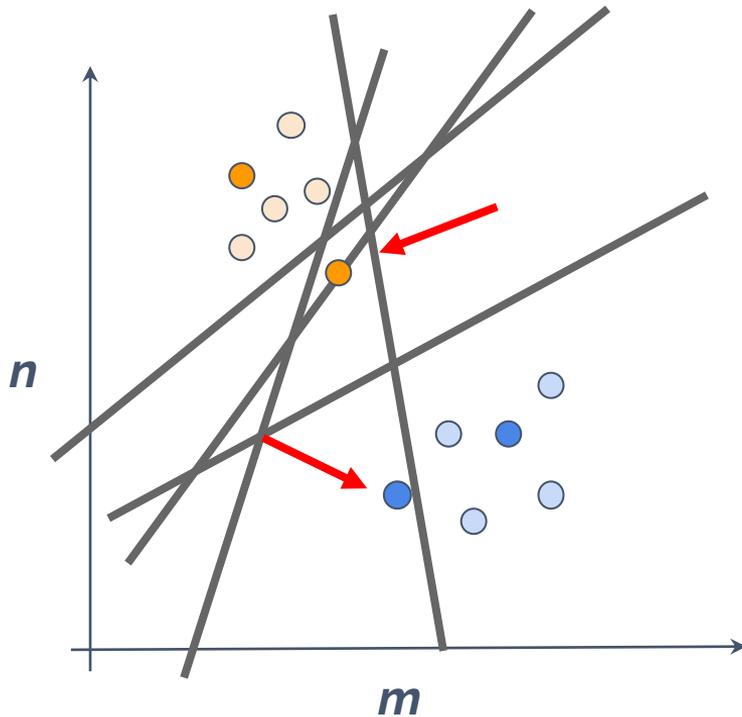
And then see if our model is a good classifier by checking our predictions on the “validation data” are correct.

Challenge highlight: Distribution shift



Note that there are lots of possible linear classifiers that would perfectly solve the training set

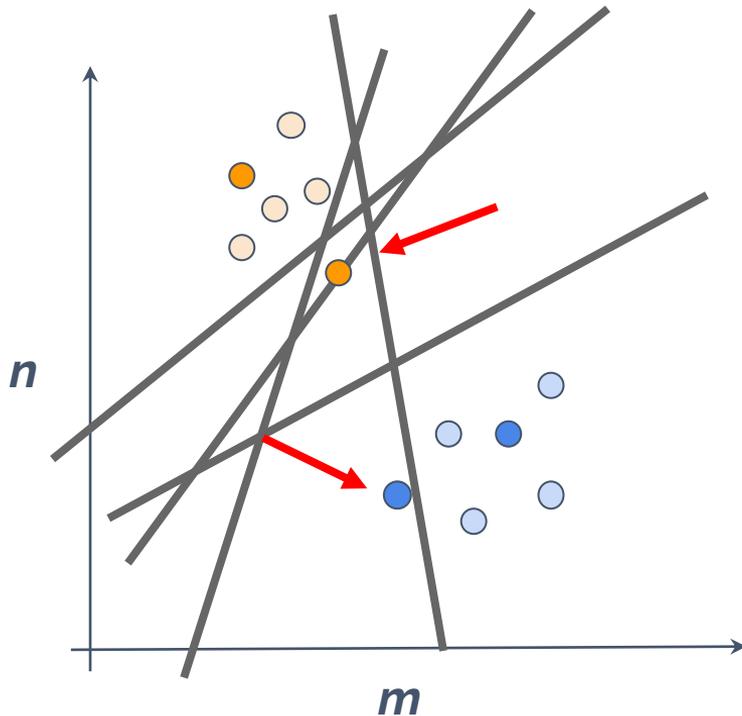
Challenge highlight: Distribution shift



Note that there are lots of possible linear classifiers that would perfectly solve the training set

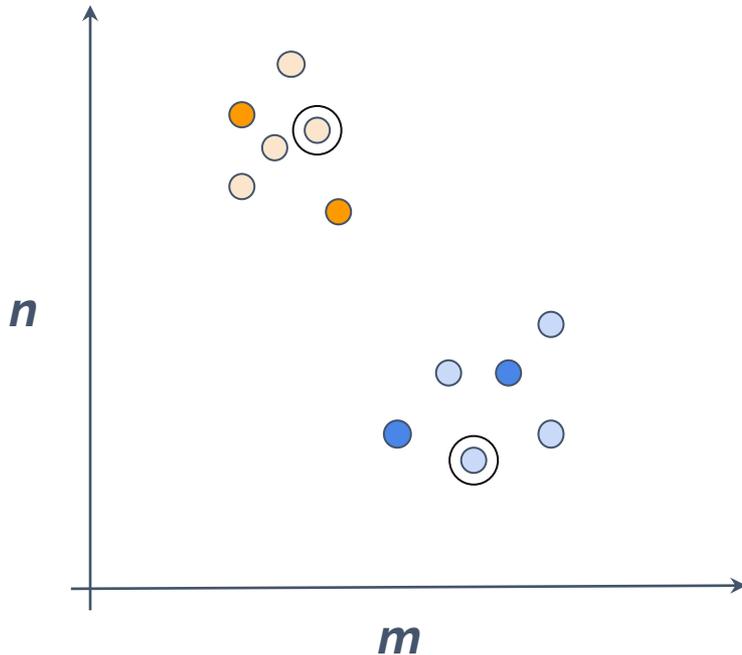
But not all of these would work well on the validation data

Challenge highlight: Distribution shift



If we decide which classifier to use based on validation data performance, then that means we're using some of our data for ***model selection***.

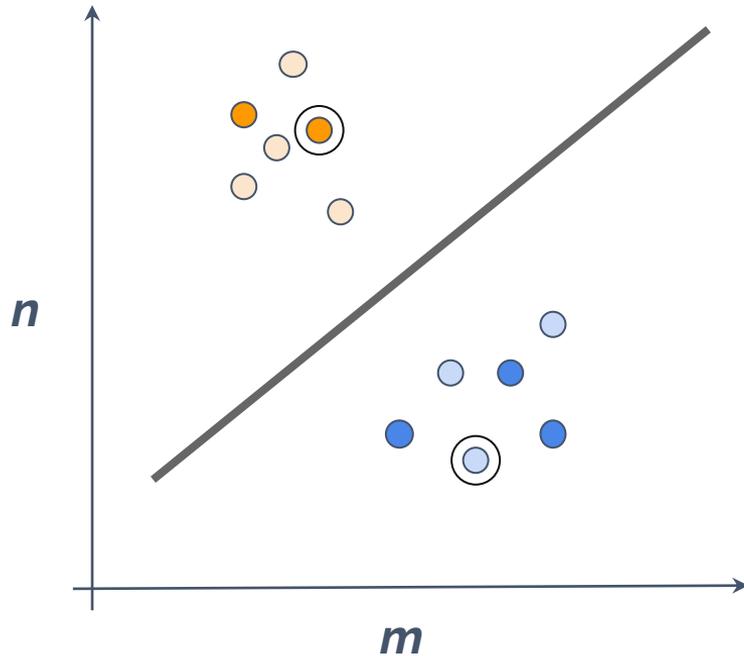
Challenge highlight: Distribution shift



This means that our validation dataset isn't truly unseen, so to maintain the integrity of our model evaluation, we further subsample the dataset to include a "test set," which we shouldn't ever look at until we've chosen our final model.

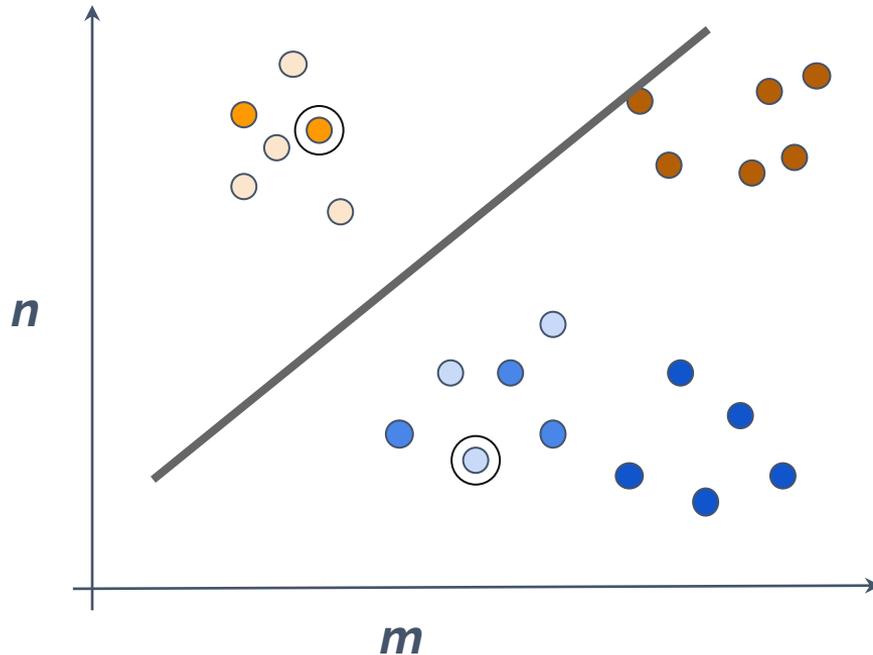
(note, ML researchers/subfields are **VERY BAD AT THIS**)

Challenge highlight: Distribution shift



What if our initial assumption was wrong and the labeled data we started with wasn't representative?

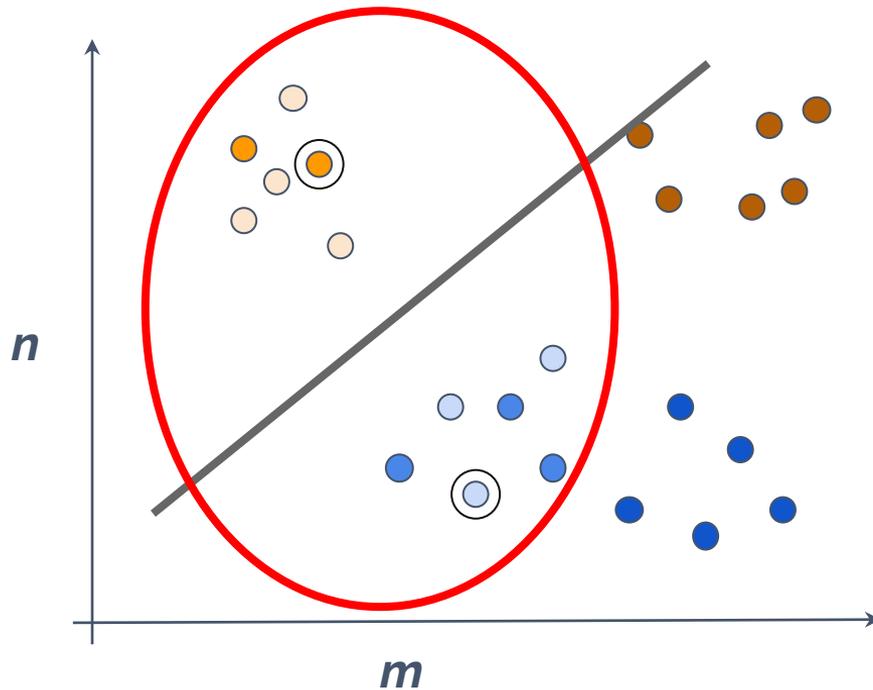
Challenge highlight: Distribution shift



Now say we try to use our model, and the data we use it on is sampled with larger values of m .

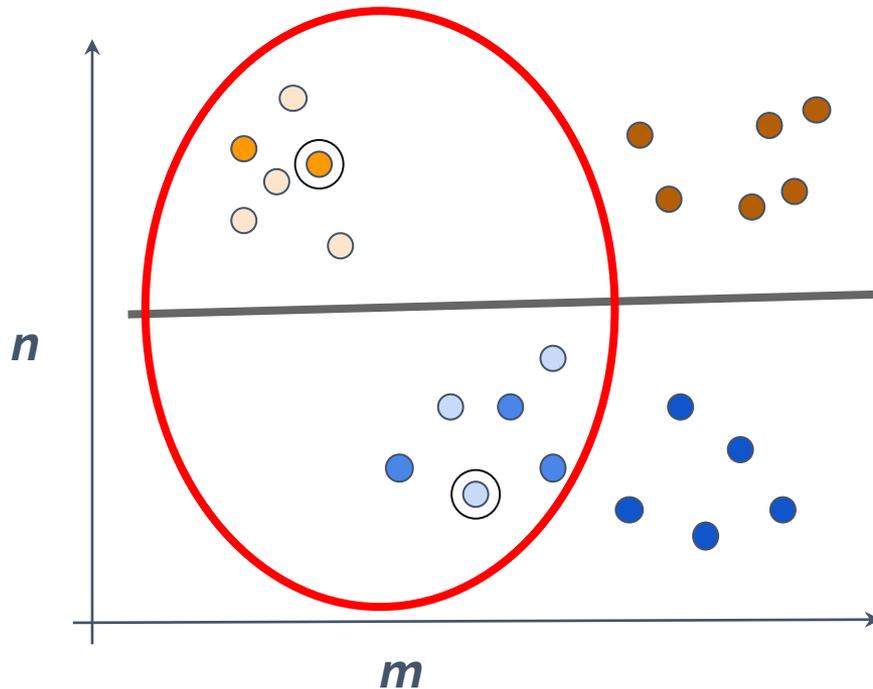
It turns out our model is **dead wrong**, it would classify all these points as blue.

Challenge highlight: Distribution shift



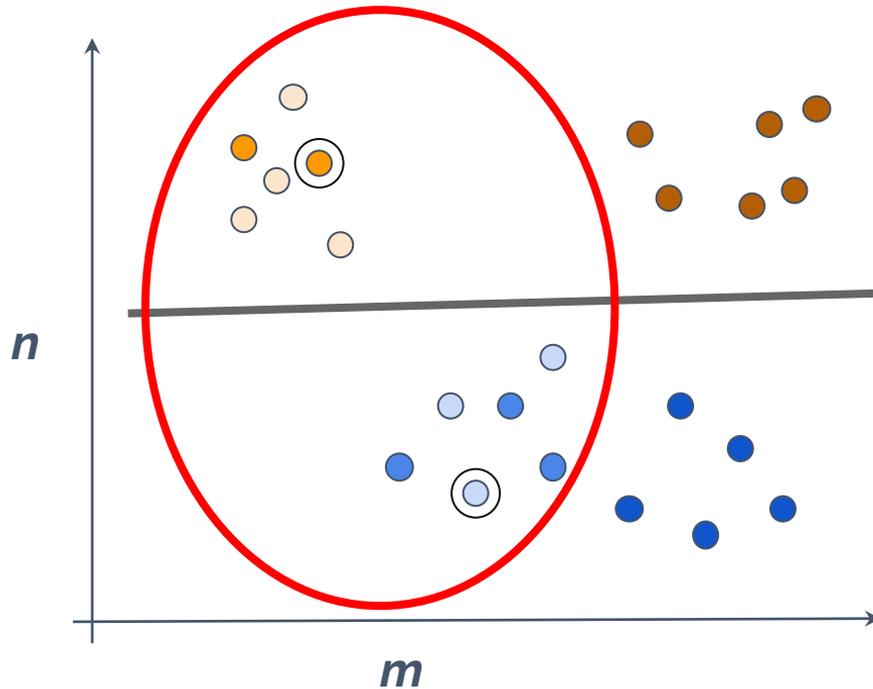
The “domain” of the original data, fell within the red circle, and the new data was *out of domain* (**OOD**)

Challenge highlight: Distribution shift



Ideally we would learn a classifier that would work on **all possible data**, but that's really hard if we've never seen data outside the red circle before

Challenge highlight: Distribution shift



In ML, this challenge is often called “**distribution shift**” and the goal for models that handle it well is “**generalization**” or “**robustness**”

Distribution shift example: Electricity use during COVID-19 pandemic

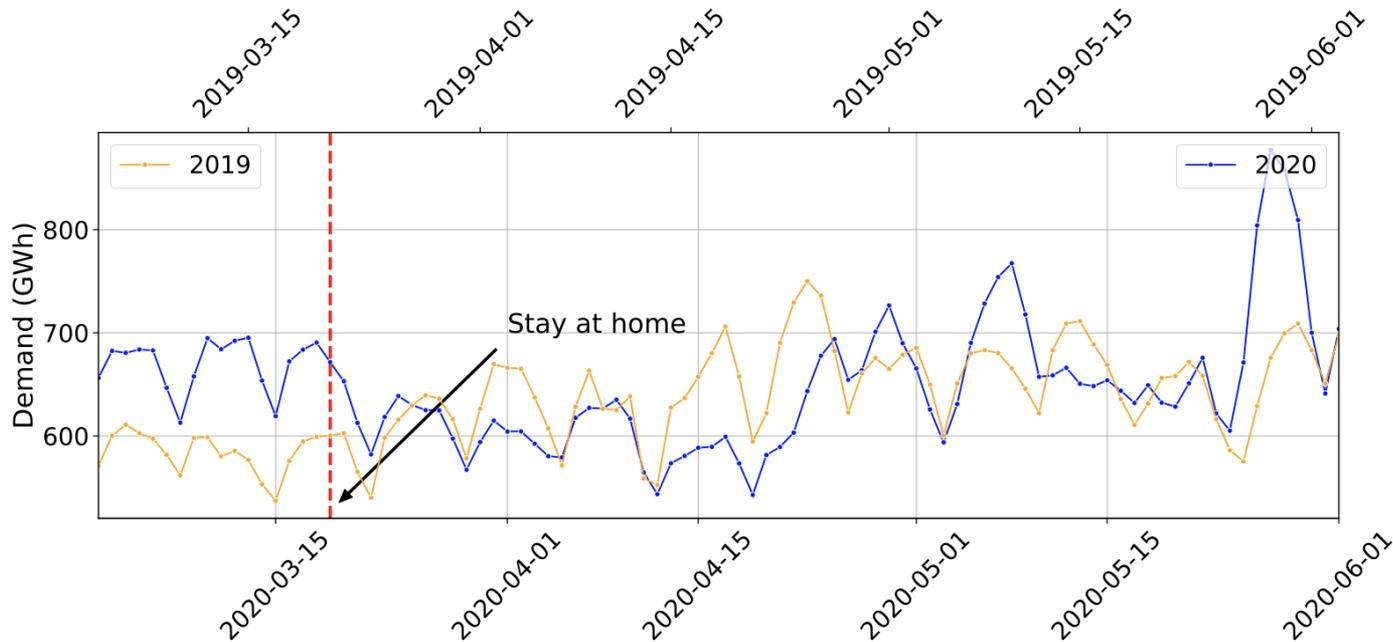


FIGURE 2. Daily Electricity Demand for California for March-May 2019 and 2020 (The first Monday of March for both years are vertically aligned).

Distribution shift example: Long-term load changes

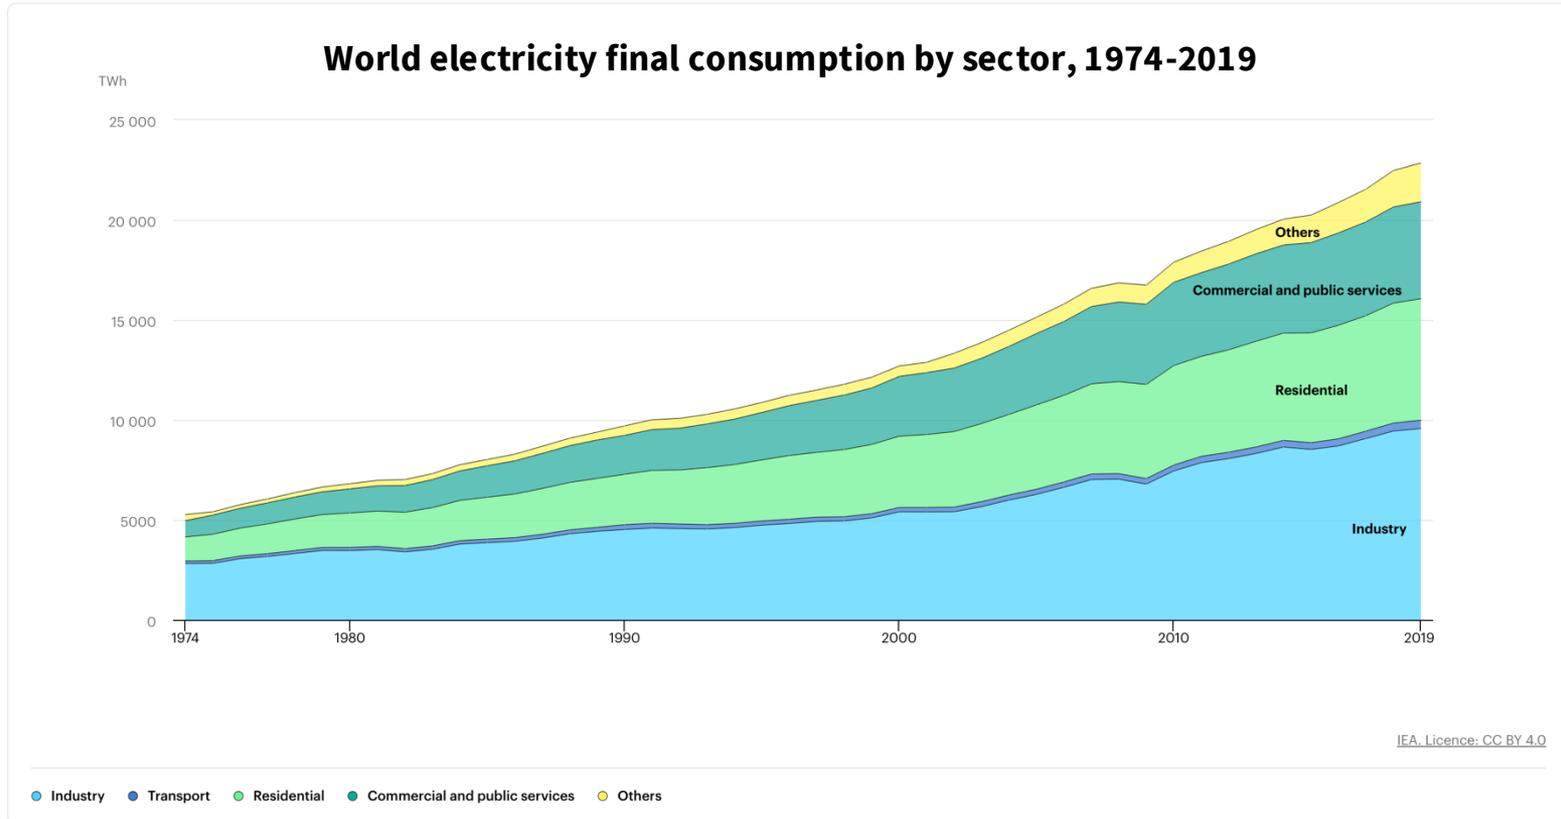


Figure source: IEA (2021), World electricity final consumption by sector, 1974-2019, IEA, Paris <https://www.iea.org/data-and-statistics/charts/world-electricity-final-consumption-by-sector-1974-2019>, Licence: CC BY 4.0

Distribution shift example: Short-term load growth

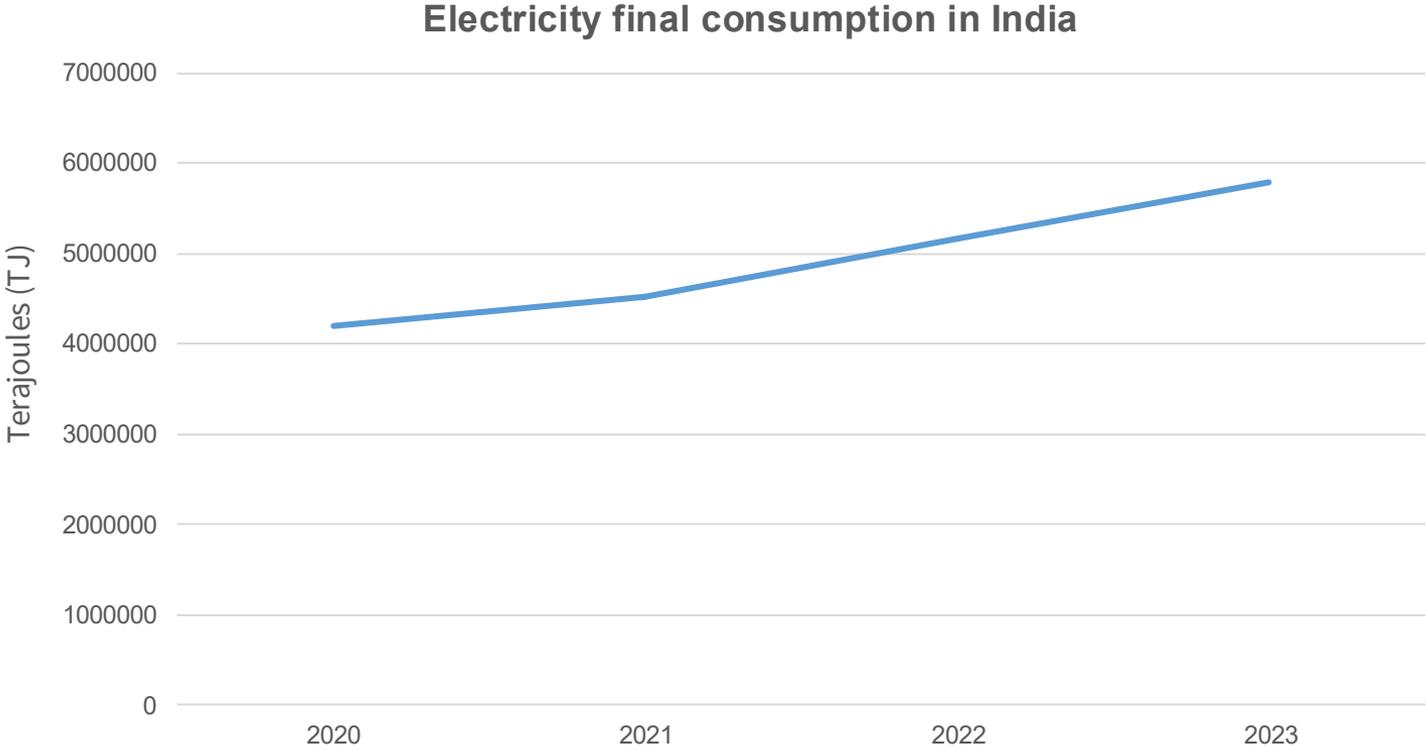


Chart data from: <https://www.iea.org/countries/india/electricity>

Distribution shift example: Differing sensor characteristics in different locations

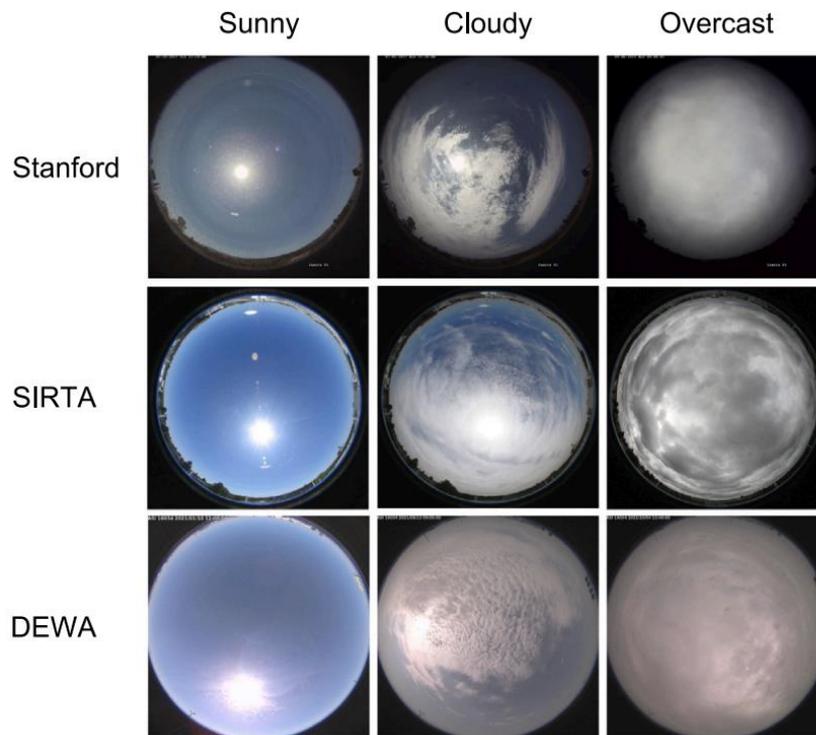


Fig. 3.1. Sky image examples from three locations with different weather conditions.

Some mechanisms for dealing with distribution shift

Data

- Collect in-domain data
- Mimic in-domain data (data augmentation, synthetic data)

ML models:

- Encourage some form of invariance across domains
- Distributionally robust optimization

Priors and hybrid techniques:

- Embed or incorporate structural knowledge into the ML model
- Use ML alongside other models (e.g., physical models or economic models)

Other: Using non-ML techniques (ML is not always the right fit!)

Takeaways

Physical modeling of power grids

- Often viewed as a network of generators and loads, at *buses* connected by *lines*
- Alternating current (AC) systems with complex-valued voltages, currents, etc.
- AC power flow equations capture nodal power balance

 **Need for situational awareness:** State estimation, predictive maintenance, and near-term forecasting

 **Application-related challenges** incl. sensor noise, low observability, inconsistent data availability, missing data, difficulty of rare event detection, & distribution shift

 **Important on-the-ground requirements** include uncertainty quantification, interpretability, generalization, robustness, & calibration with decision-making needs