ML4CC: Lecture 9

Sit with your discussion groups (same as last time)

Assignments reminder

Keep doing your PMIRO+Q

Your project assignment was due today

Keep working on your projects!

Summary of last paper

P - Want to recommend more sustainable products to people

M - Use evolutionary algorithms to identify recommended baskets that meet a lot of constraints

I - applying this method for sustainable product recommendation; use of a neural network-based method

R - the methods all find different optimal solutions that would results in reduced environmental impact if accepted

O - would people actually accept this suggestions?

Climate Change in the News

A few seconds later, a device resembling a snow maker began to rumble, then produced a great and deafening hiss. A fine mist of tiny aerosol particles shot from its mouth, traveling hundreds of feet through the air.

It was the first outdoor test in the United States of technology designed to brighten clouds and bounce some of the sun's rays back into space, a way of temporarily cooling a planet that is now dangerously overheating. The scientists wanted to see whether the machine that took years to create could consistently spray the right size salt aerosols through the open air, outside of a lab.

And yet, the idea of interfering with nature is so contentious, organizers of Tuesday's test kept the details tightly held, concerned that critics would try to stop them. Although the Biden administration is funding research into different climate interventions, including marine cloud brightening, the White House distanced itself from the California study, sending a statement to The New York Times that read: "The U.S. government is not involved in the Solar Radiation Modification (SRM) experiment taking place in Alameda, CA, or anywhere else."

David Santillo, a senior scientist at Greenpeace International, is deeply skeptical of proposals to modify solar radiation. If marine cloud brightening were used at a scale that could cool the planet, the consequences would be hard to predict, or even to measure, he said.

Paper 8 Discussion

SustainGym: A Benchmark Suite of Reinforcement Learning for Sustainability Applications

Christopher Yeh, Victor Li, Rajeev Datta, Yisong Yue, Adam Wierman
California Institute of Technology
Department of Computing and Mathematical Sciences
cyeh,vhli,rdatta,yyue,adamw@caltech.edu

Abstract

The lack of standardized benchmarks for reinforcement learning (RL) in sustainability applications has made it difficult to both track progress on specific domains and identify bottlenecks for researchers to focus their efforts on. In this paper, we present SustainGym, a suite of two environments designed to test the performance of RL algorithms on realistic sustainability tasks. The first environment simulates the problem of scheduling decisions for a fleet of electric vehicle (EV) charging stations, and the second environment simulates decisions for a battery storage system bidding in an electricity market. We describe the structure and features of the environments and show that standard RL algorithms have significant room for improving performance. We discuss current challenges in introducing RL to real-world sustainability tasks, including physical constraints and distribution shift.

Attendance

Select one person from the group to go to this Google Doc and write down the names of all people present in the group (remember to mark who took attendance!). **If someone is virtual, mark it with a V**.

https://docs.google.com/document/d/1PKhw9E2lJpAnFrFO88DOc2rscZFVlcIv47QNa5h1sGs/edit?usp=sharing (link is in Brightspace under Syllabus content)

Discussion Question 1

What is meant by “toy model” and why is it important to test methods not just on toy models?

Toy models are highly-simplified versions of real problems

Toy models can be useful to prototype a method and understand how it works. But they are limited because they don't usually capture what makes real problems challenging such as:

- Sparse feedback
- Long-term dependencies
- Complex observations
- Complex action spaces
- Uncertainty in the world
- Data needs

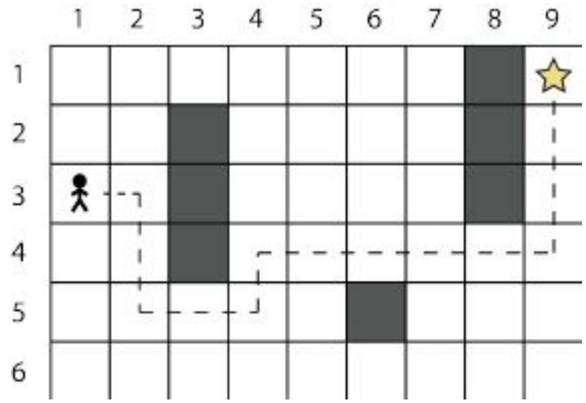

Discussion Question 2

Explain figure 1 in your own words. How did behavior change during the pandemic?

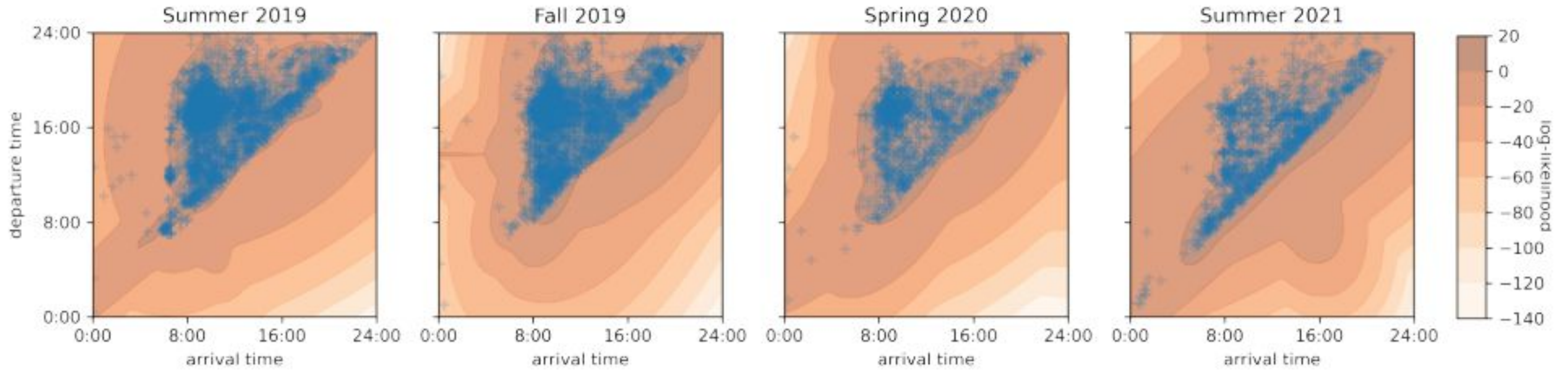

Figure 1: EV arrival vs. departure times for the Caltech EV charging network. Historical data is in blue, and log-likelihood contours from a 30-component GMM are in orange. The distribution of EV arrival and departure times changed noticeably between 2019 (pre-COVID) to 2020.

Far fewer people charged their cars during the pandemic; they also came at odd times of day.

Discussion Question 3

Explain MOER in your own words. What do you think might contribute to the MOER value changing over time?

MOER: how “dirty” electricity is at each time point

actions of an agent as a result of changes in electricity consumption. Our environments use data on California’s historical marginal operating emissions rate (MOER, in $kgCO_2/kWh$), which is the increase in CO_2 emissions per increase in energy demand. The MOER at time t is denoted $m_t \in \mathbb{R}_+$,

MOER can change based on the availability of different energy sources.

Discussion Question 4

What are the “actions” the agent can take in this reinforcement learning problem?

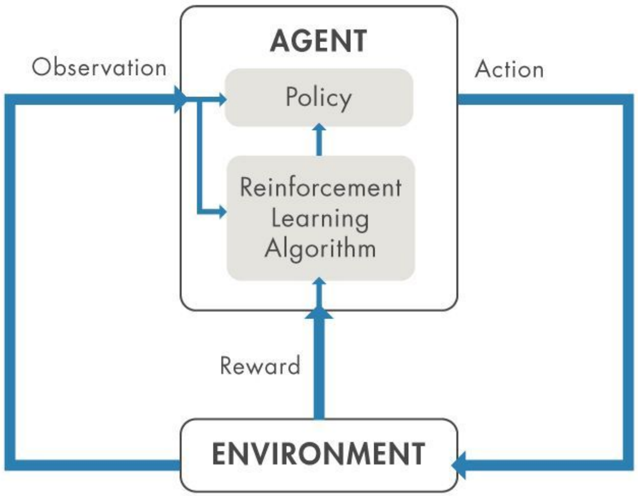

The agent controls the charging speed for all cars

and departure events are discretized to these 5-minute intervals ($T = 288$, $\tau = 5/60$ hours). At every time step, an agent decides the charging rates (a.k.a. "pilot signals") for each EVSE to be executed for the duration of that time step. That is, a single agent simultaneously controls all n EVSEs.

These actions are constrained by the charging infrastructure

Action Space. The action space is continuous $a(t) \in [0,1]^n$, representing the pilot signal normalized by the maximum signal allowed M (in amps) for each EVSE. Physical infrastructure in a charging network constrain the set $\mathcal{A}_t$ of feasible actions at each time step t [18]. Furthermore, the EVSEs only support discrete pilot signals, so $\mathcal{A}_t$ is nonconvex. To satisfy these physical constraints, `EVChargingEnv` can project (A1) an agent's action $a(t)$ into the convex hull of $\mathcal{A}_t$ and round it to the nearest allowed pilot signal, resulting in final normalized pilot signals $\tilde{a}(t)$. ACNSim processes $\tilde{a}(t)$ and returns the actual charging rate $M\bar{a} \in \mathbb{R}^n_+$ (in amps) delivered at each EVSE, as well as the remaining demand $e_i(t+1)$.

Discussion Question 5

How is the reward calculated in this RL setting?

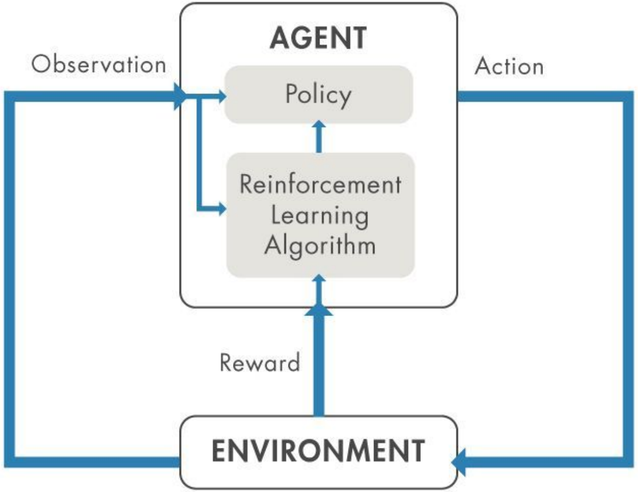

Get reward by doing a lot of charging in a way that doesn't violate physical limitations and also uses low-emissions energy.

Reward Function. The reward function is a sum of three components: $r(t) = p(t) - c_V(t) - c_C(t)$. The profit term $p(t)$ aims to maximize energy delivered to the EVs. The constraint violation cost $c_V(t)$ aims to reduce physical constraint violations and encourage the agent's action $a(t)$ to be in $\mathcal{A}_t$. Finally, the CO_2 emissions cost $c_C(t)$, which is a function of the MOER m_t and charging action, aims to reduce emissions by encouraging the agent to charge EVs when the MOER is low.

How does this relate to last week's paper?

Discussion Question 6

What are the observations given to the agent? How might they help it achieve high reward?

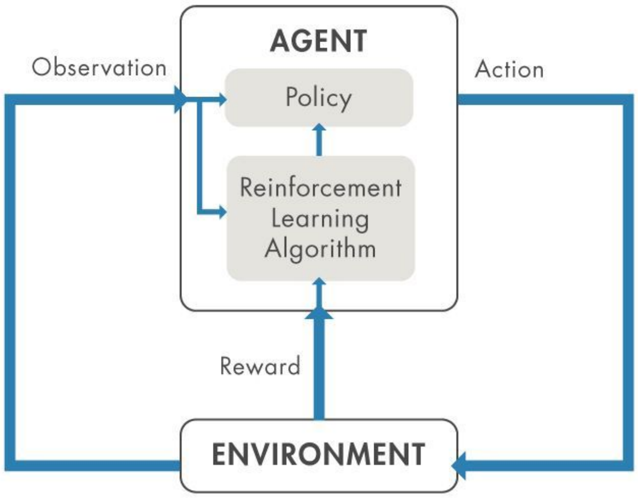

Observations

Observation Space. An observation at time t is $s(t) = (t, d, e, m_{t-1}, \hat{m}_{t:t+k-1|t})$. $t \in \mathbb{Z}_+$ is the fraction of day between 0 and 1, inclusive. $d \in \mathbb{Z}^n$ is estimated remaining duration of each EV (in # of time steps). $e \in \mathbb{R}^n_+$ is remaining energy demand of each EV (in kWh). If no EV is charging at EVSE i, then $d_i = 0$ and $e_i = 0$. If an EV charging at EVSE i has exceeded the user-specified estimated departure time, then d_i becomes negative, while e_i may still be nonzero.

Time of day can be used to predict arrival and leaving times

Estimated remaining duration indicates how much charging time is left

Remaining energy is needed to plan how to spread out charging over time

MOER provides info on how GHG-intensive energy will be

Discussion Question 7

What kind of out of distribution test is performed on this model? What are the results?

Training on pre-pandemic data and testing during the pandemic

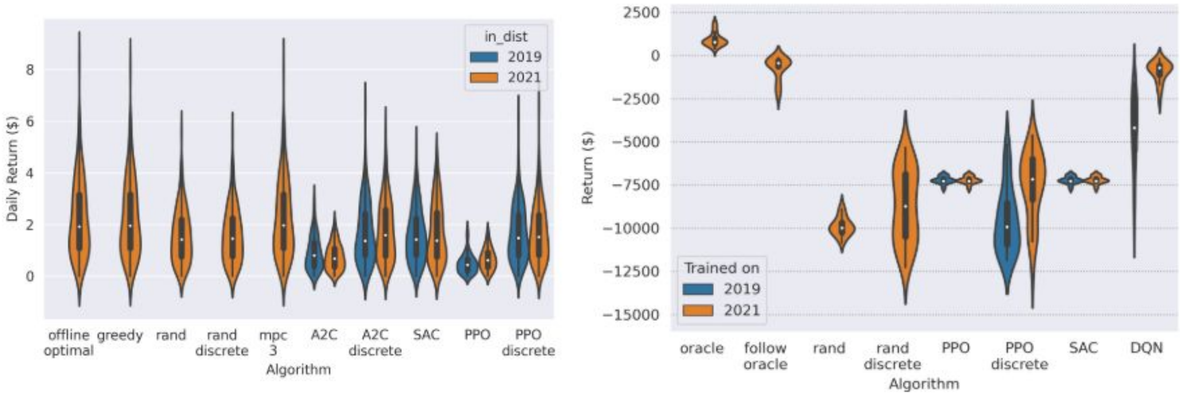

Figure 2: Returns from controllers evaluated on (left) EVChargingEnv and (right) ElectricityMarketEnv. For EVChargingEnv, algorithms are tested on actual data from Summer 2021 with discrete actions. A2C, SAC, and PPO were trained on artificial data sampled from GMM models fitted to Summer 2019 ("out dist") or Summer 2021 ("in dist"). For ElectricityMarketEnv, algorithms were evaluated on data from May 2021 with all rewards delayed until the terminal step. Thus, models trained on May 2019 data are out-of-distribution.

These RL models perform poorly on both...

Discussion Question 8

Share what questions you wrote in your PMIRO+Q and decide as a group what you’d like to ask.

Update your PMIRO+Q

Submit a second file to the Brightspace assignment (don’t overwrite the original):

It should:

Update your PMIRO as needed

Answer your own Q

You can be talking with your group during this!

15 min break

Lecture

Climate Content: Power grid and alternative energy sources

Machine Learning: Graph neural networks

How does the power grid work?

https://www.wsj.com/video/how-does-the-us-power-grid-work/1671AA83-D0D2-4C75-913C-B381341159F4.html

Power grid

Generators: sources of energy entering the grid, such as power plants

Consumers: users of energy like homes and commercial buildings

Transmission lines: connections that carry electricity between generators and consumers

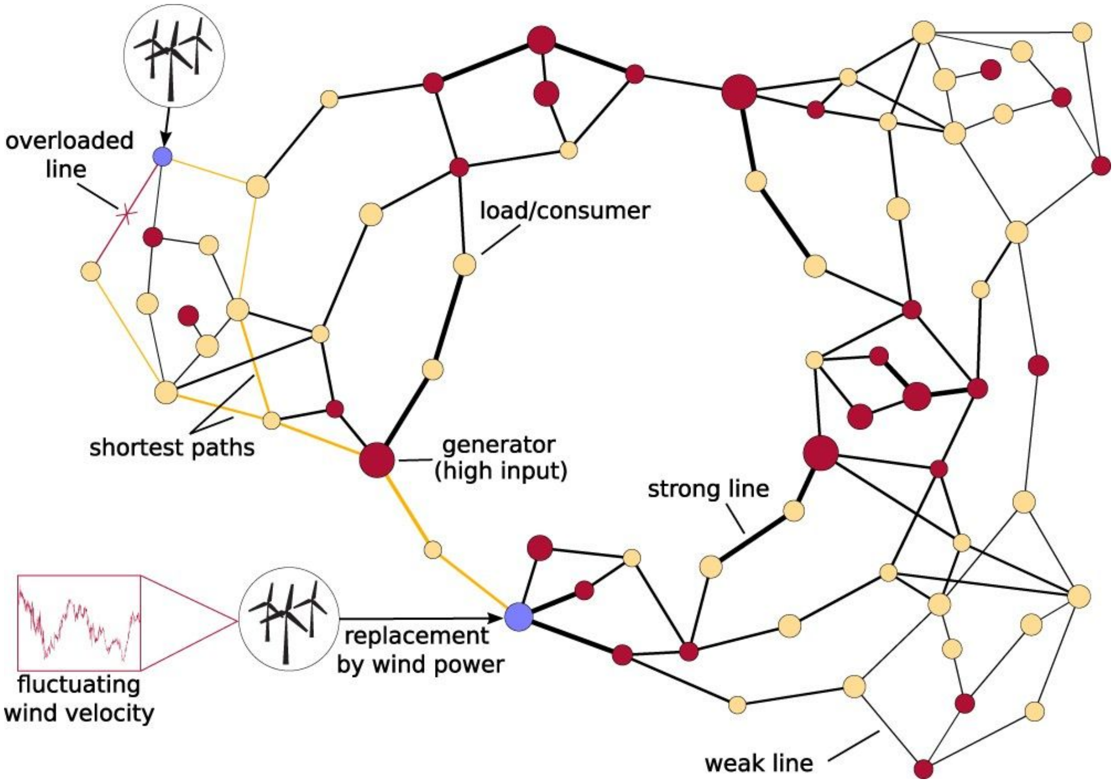

Ways to generate energy

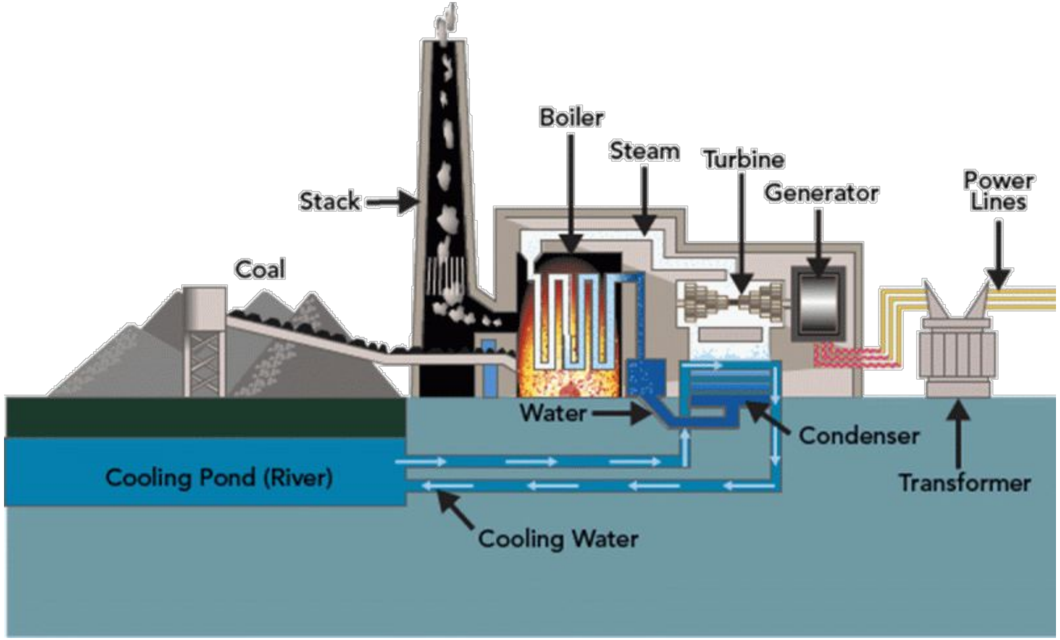

Let's Talk Science

Wind turbines

Convert kinetic energy of air into electric power using a generator.

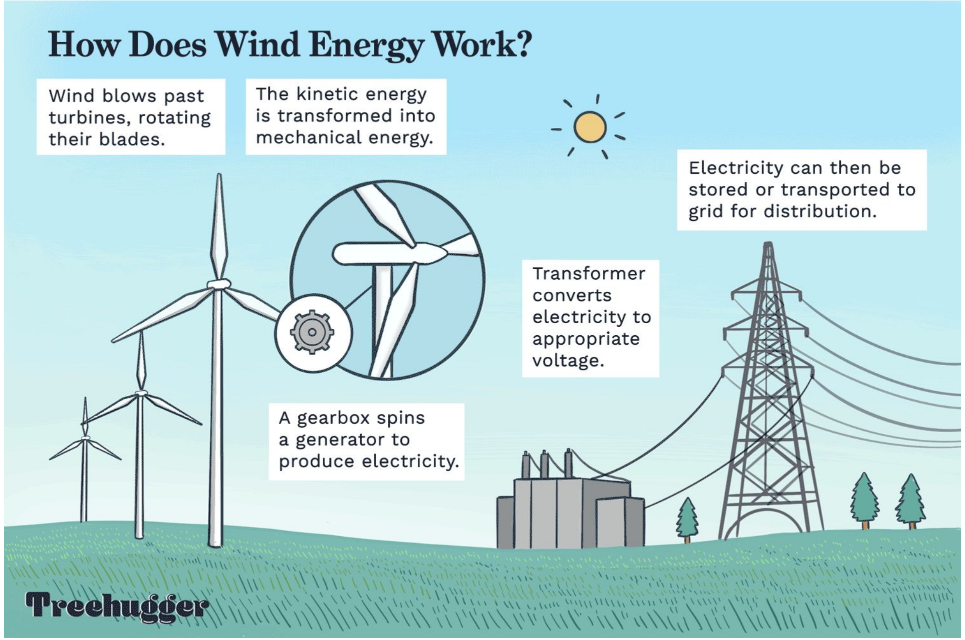

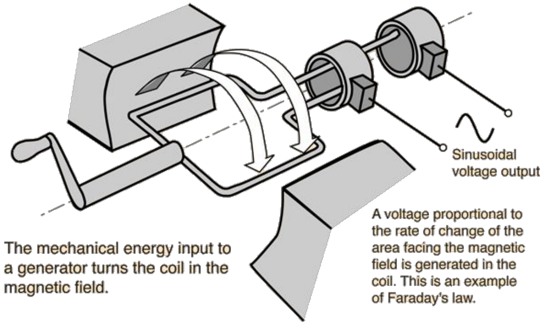

http://hyperphysics.phy-astr.gsu.edu

Hydropower

Hydropower plants are one of the oldest mechanisms used to produce power due to their simplistic mechanisms.

Very efficient: reaching up to 95% efficiency for large scale and 85% in small scale applications.

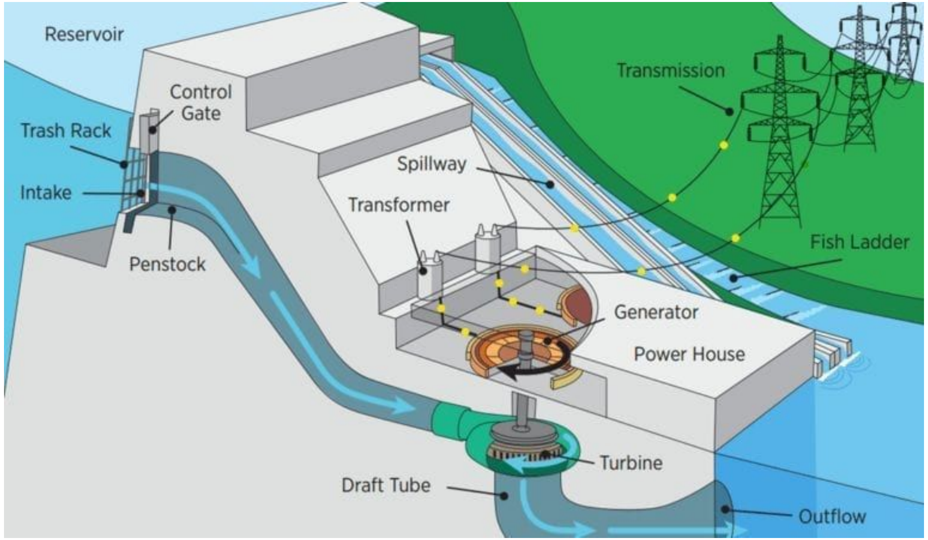

Nuclear Power

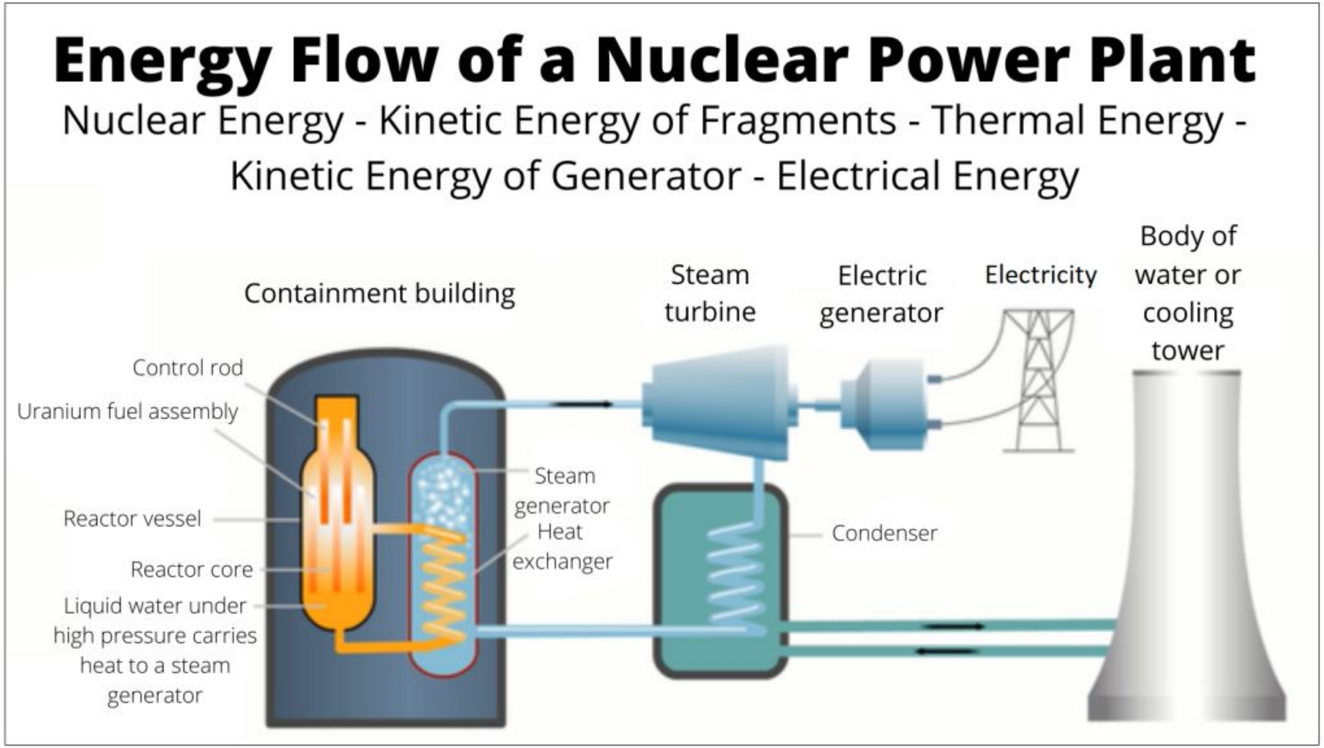

Photovoltaic solar panels

A semiconductor layer converts the Sun's energy into useful electricity through a process called the photovoltaic effect .

On either side of the semiconductor is a layer of conducting material which "collects" the electricity produced

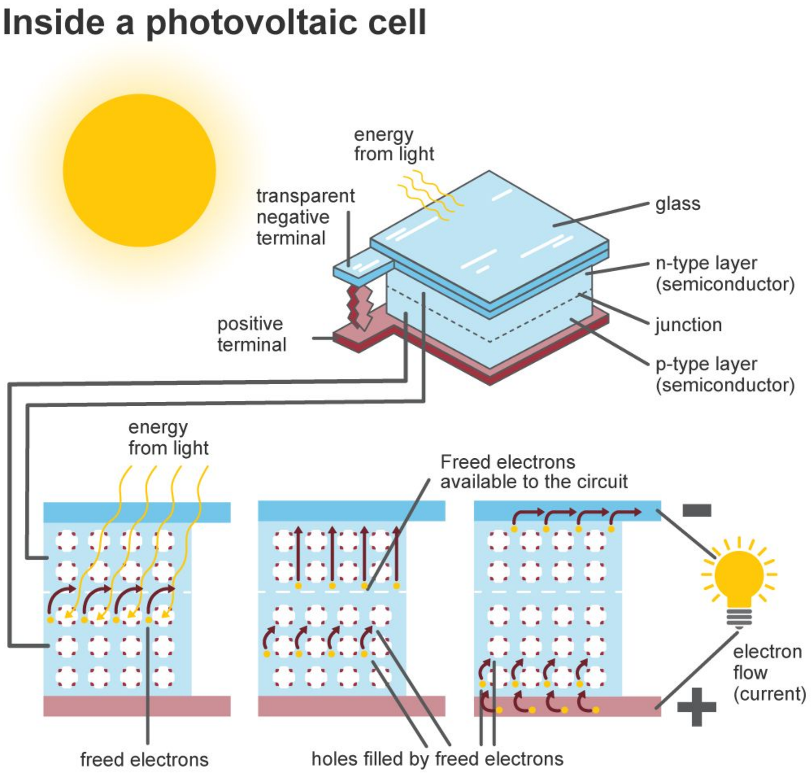

Source: U.S. Energy Information Administration

https://energyeducation.ca/encyclopedia/Photovoltaic_cell

Ways to generate energy

Non-renewable energy relies on limited resources (like really old dead animals).

Renewable energy relies on sources that can be regenerated by existing natural forces.

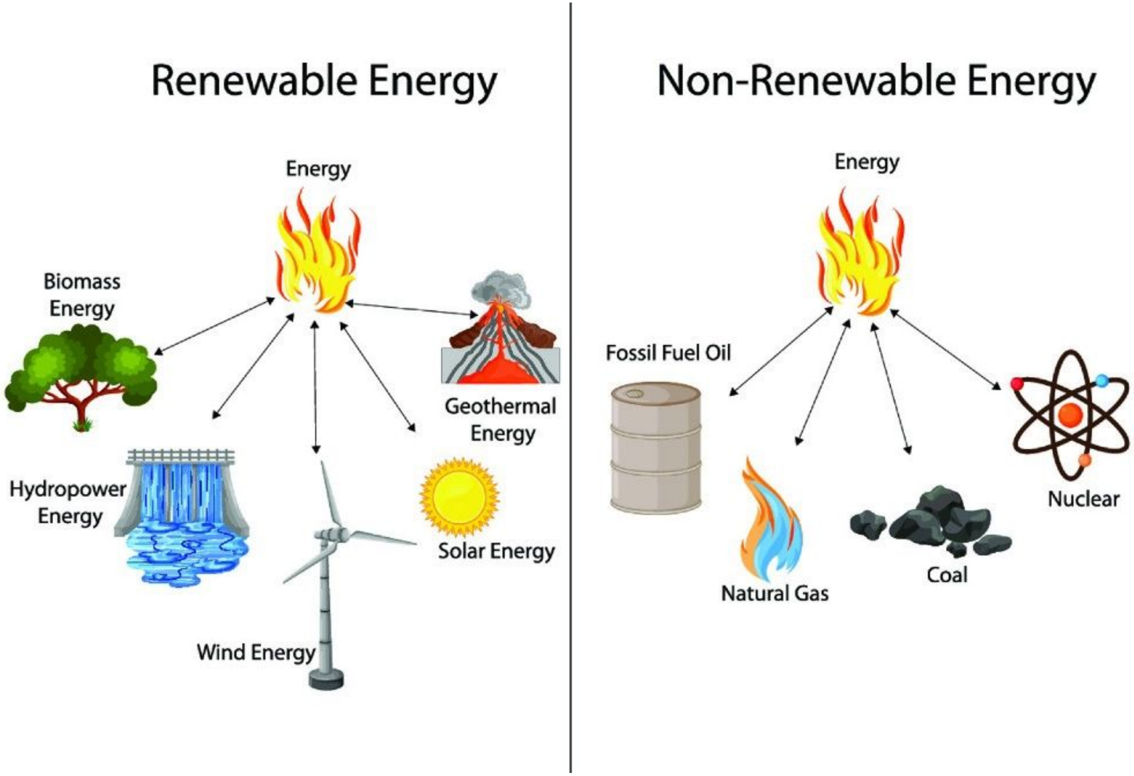

GHGs from different energy sources

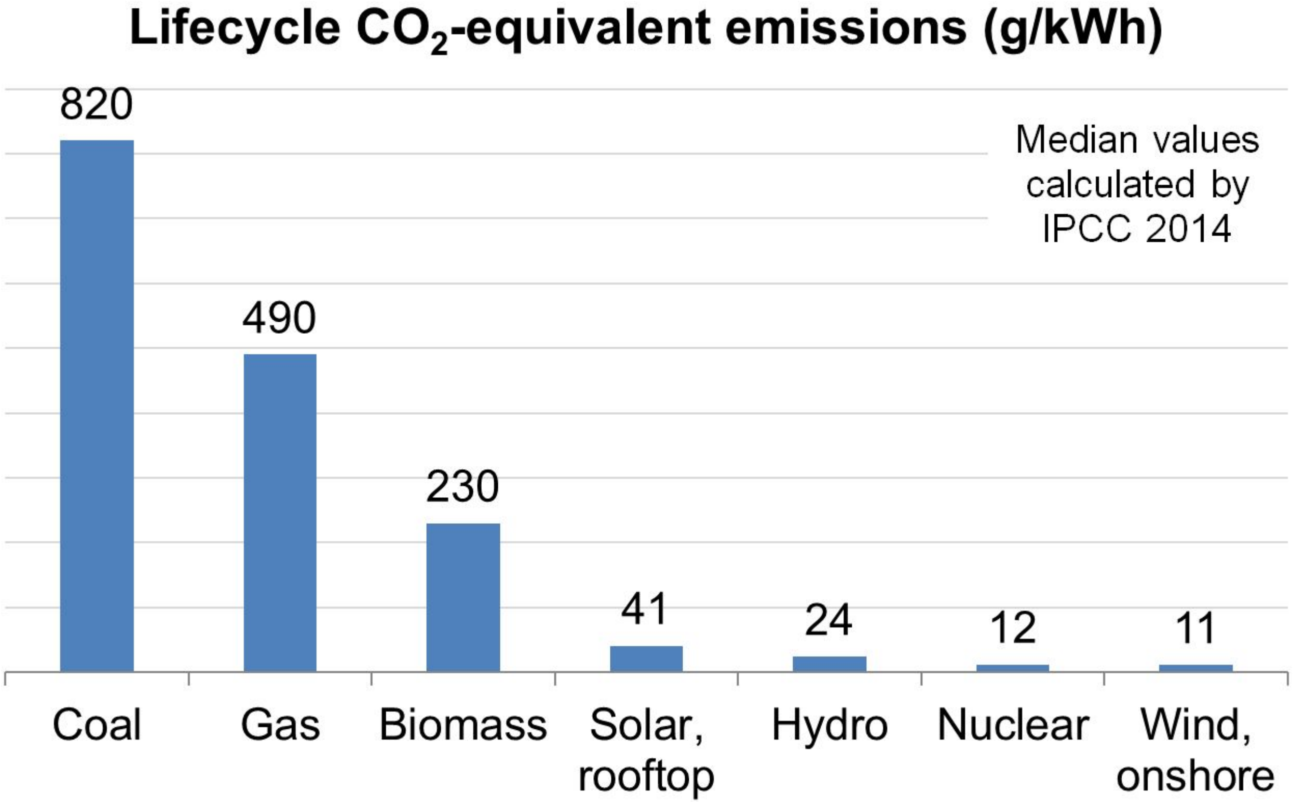

Direct emissions come from the generation of power itself.

Indirect emissions come from the production and maintenance of power plants.

Renewable is not the same as low-emission

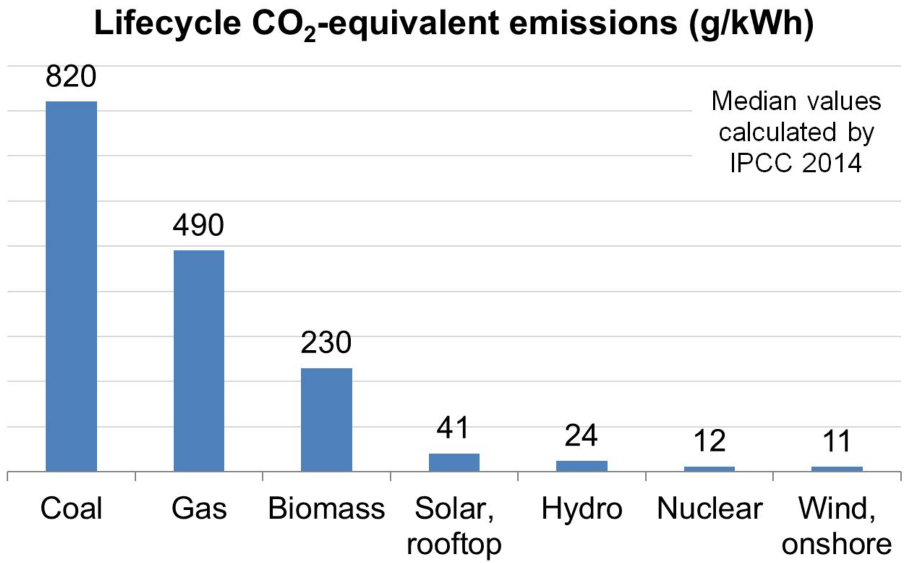

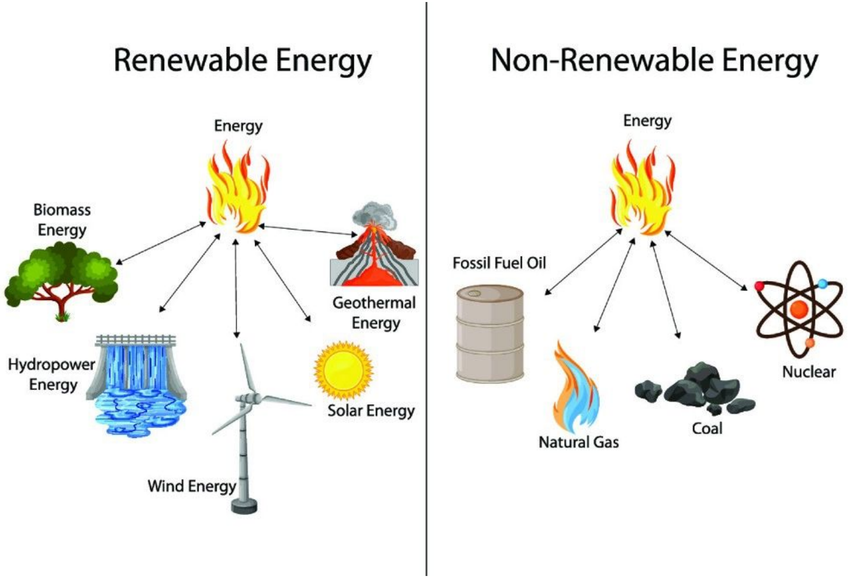

Labels

“Green” is not a well-specified term.

People debate if nuclear is green, and natural gas advocates have lobbied to label it as green.

Ohio Gov. Declares Natural Gas 'Green Energy.' It Doesn’t Work Like That.

A fossil fuel is a fossil fuel.

BY TIM NEWCOMB PUBLISHED: JAN 10, 2023

Tanaonte // Getty Images

- A new Ohio bill that calls natural gas a “green energy” opens state lands to oil and gas drilling.
- Natural gas is defined as a fossil fuel, albeit cleaner than some counterparts.
- Environment protectors say the new bill threatens Ohio state land to additional drilling.

popular mechanics

Where does electricity come from in the US?

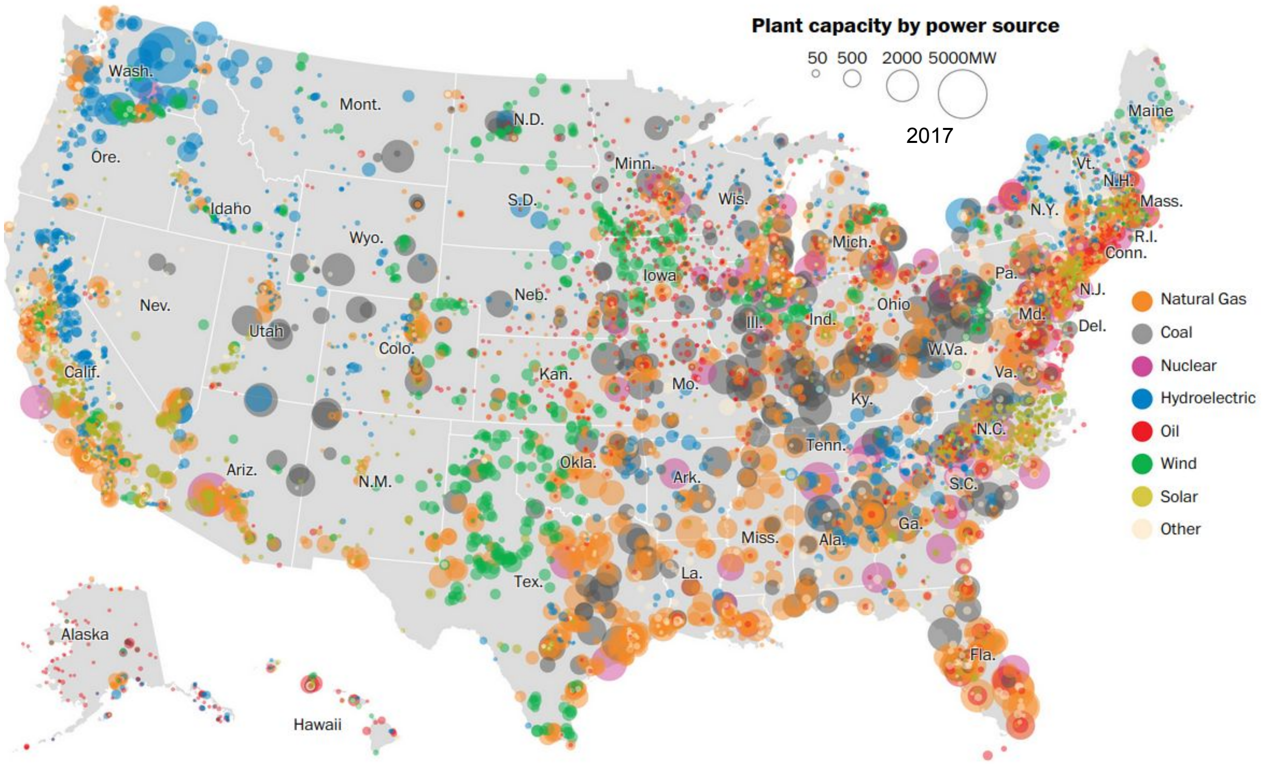

https://www.washingtonpost.com/graphics/national/power-plants/

Natural gas has expanded due to fracking.

Coal is more popular in the East.

Nuclear has a high power:space ratio, but is used unevenly across states.

Hydroelectric requires the right environmental factors.

Oil is only the leading source in Hawaii

Wind is best in the plain states.

Solar is predominant in Southwest and certain Eastern states.

Where does electricity come from globally?

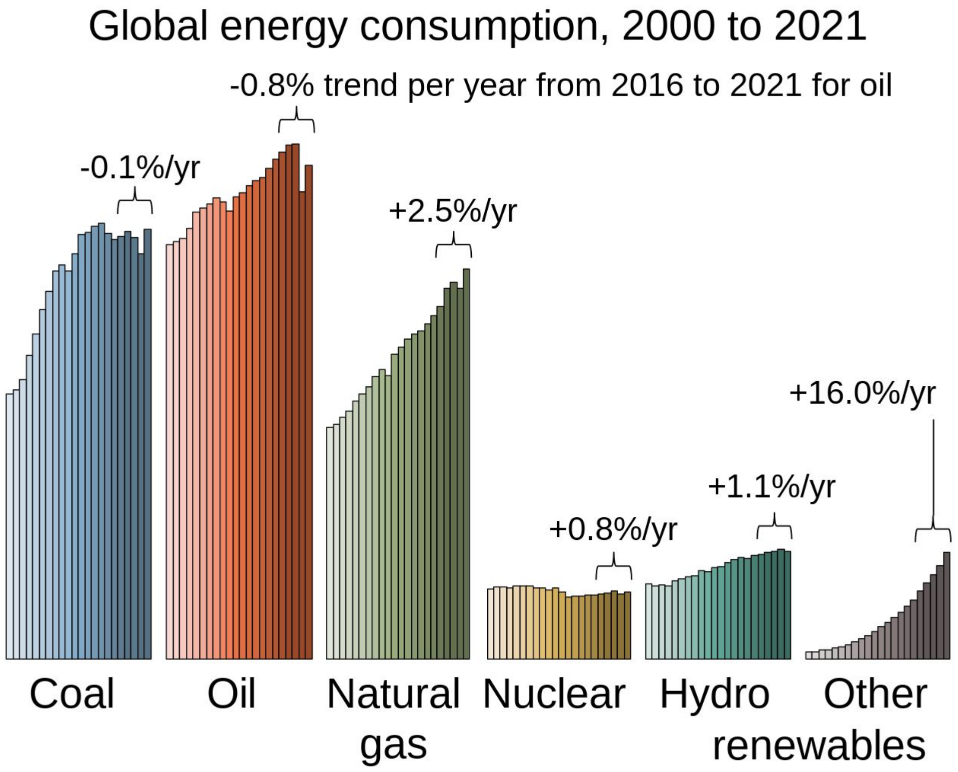

Figure 5.2 Electricity generation in the European Union by country and source, 2012

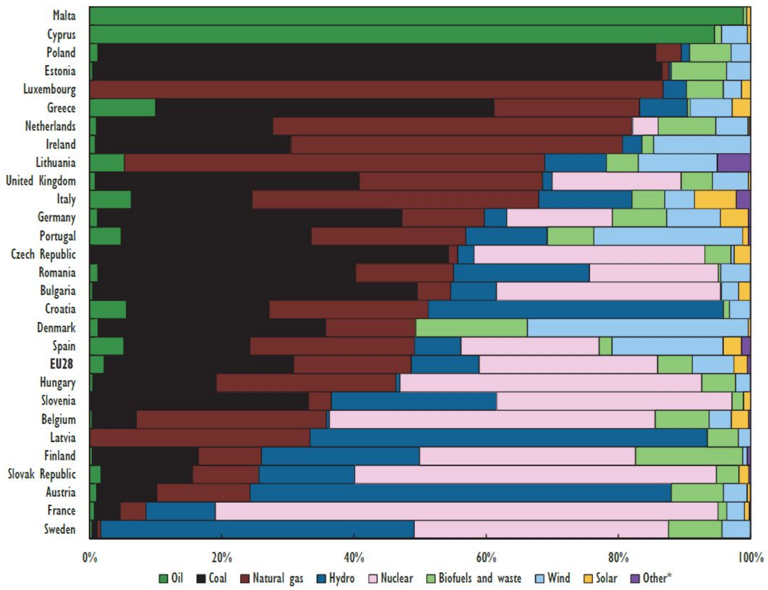

* *Other* includes geothermal, peat and ambient heat production.

Sources: IEA (2014a), *Energy Balances of OECD Countries*, OECD/IEA, Paris; IEA (2014c), *Energy Statistics of Non-OECD Countries*, OECD/IEA, Paris.

The need to produce more energy from clean sources

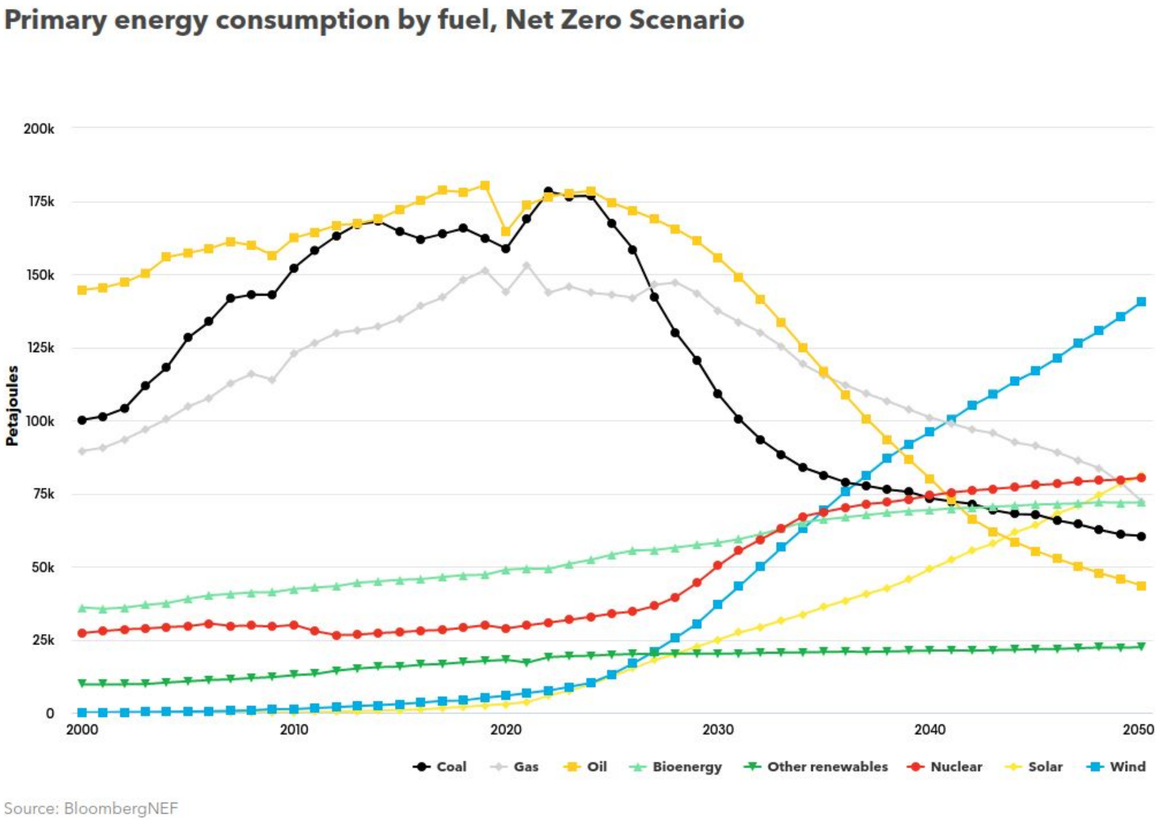

Problem: Not all energy sources are “dispatchable”

Wind, solar, and hydro power can depend on weather conditions and therefore can’t be relied on in times of higher energy demand

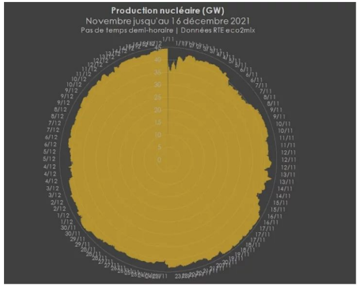

Nuclear production

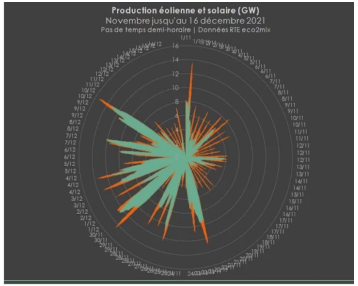

Wind and solar production

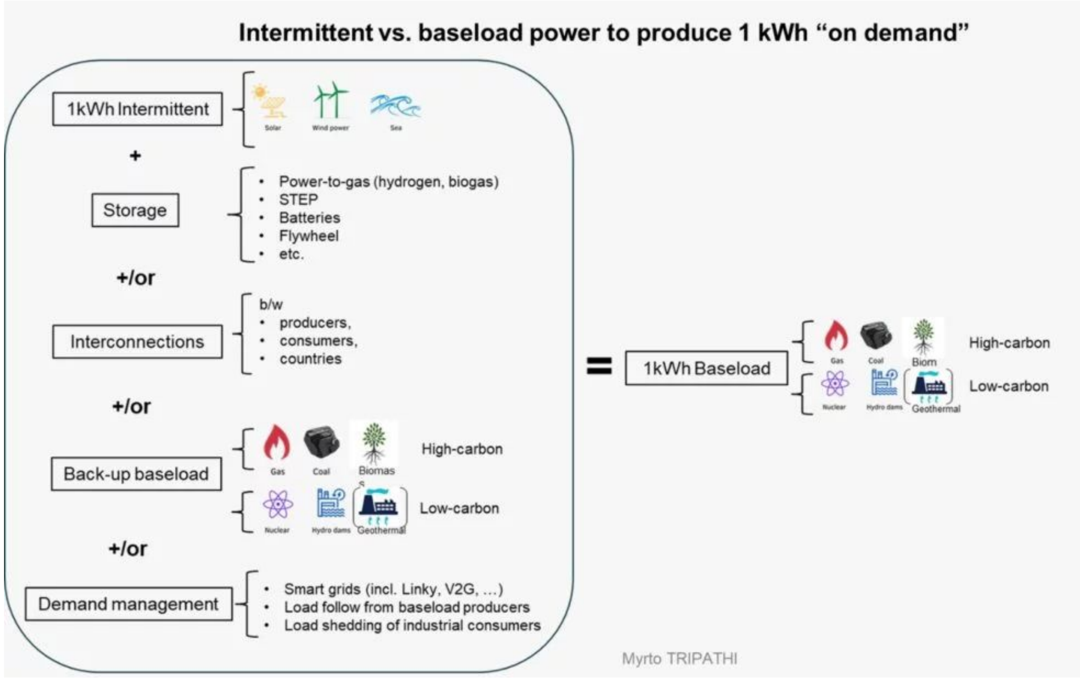

Grid Balancing

Supply needs to equal demand on a second-by-second basis. Errors of 1% in the frequency of generated AC currents can cause problems

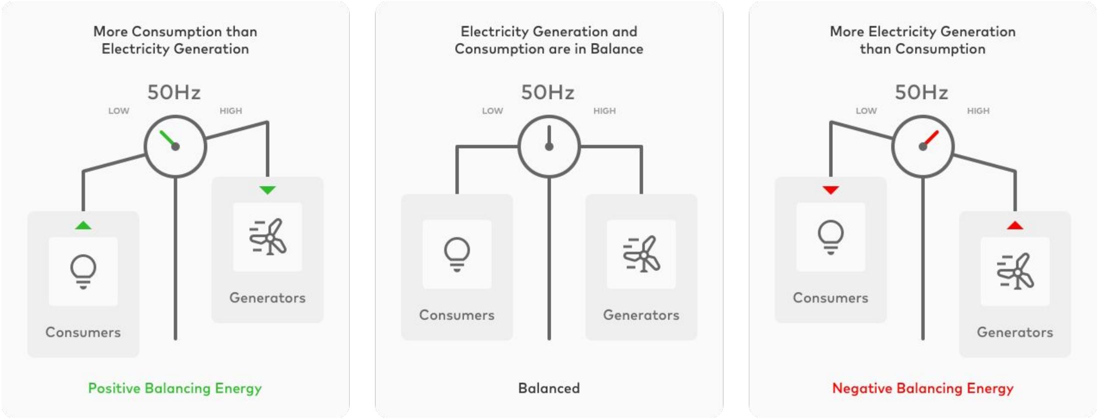

A 2012 report from the Federal Energy Regulatory Commission estimated that poor grid balancing may cost billions of dollars and release unnecessary emissions

Blackouts

Damage from excess voltage

Migodi

Climate change (and our response) impacts demand

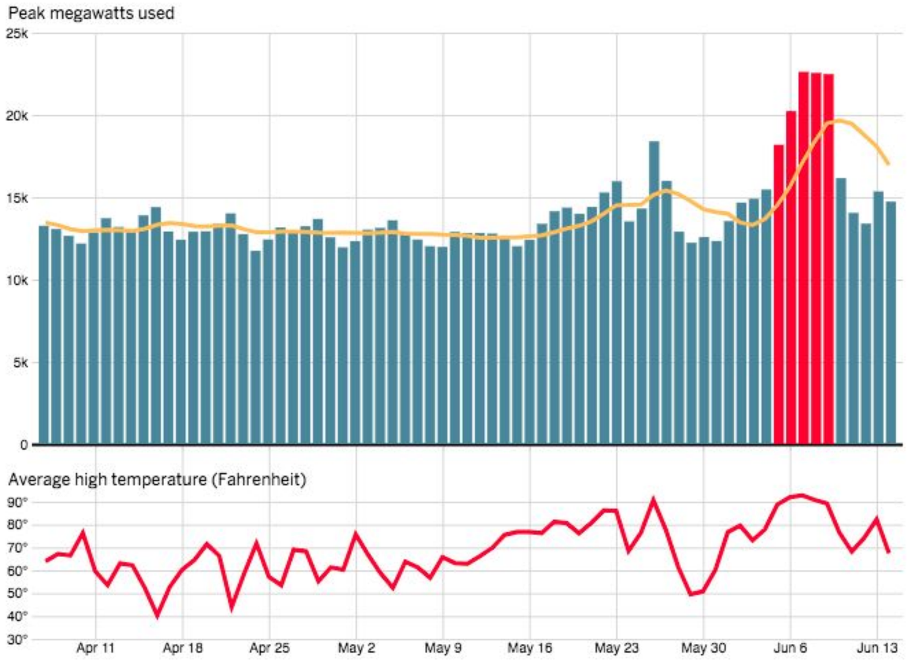

Boston Globe

Extreme weather events cause energy demand spikes

Climate change (and our response) impacts demand

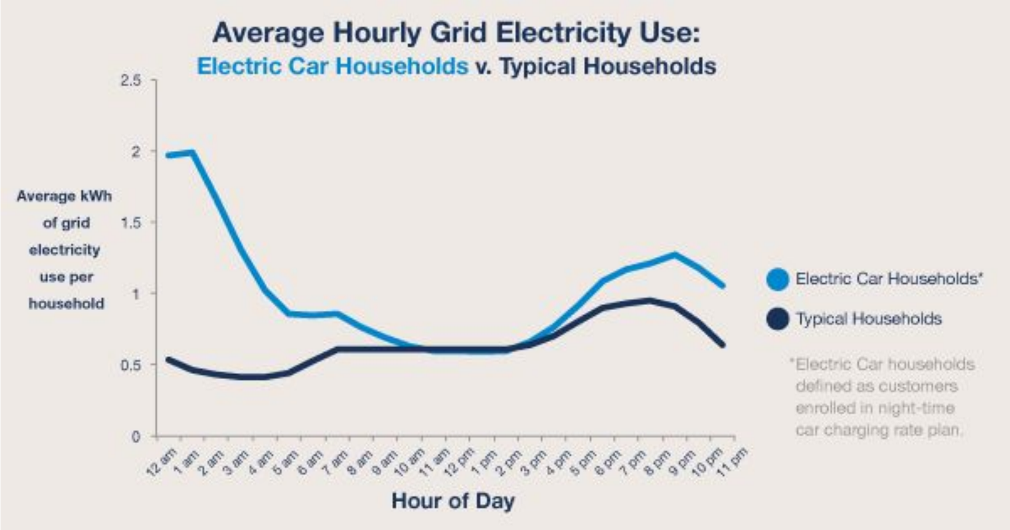

Increased electrification of devices will increase power demand.

World Economic Forum

How climate change will impact power

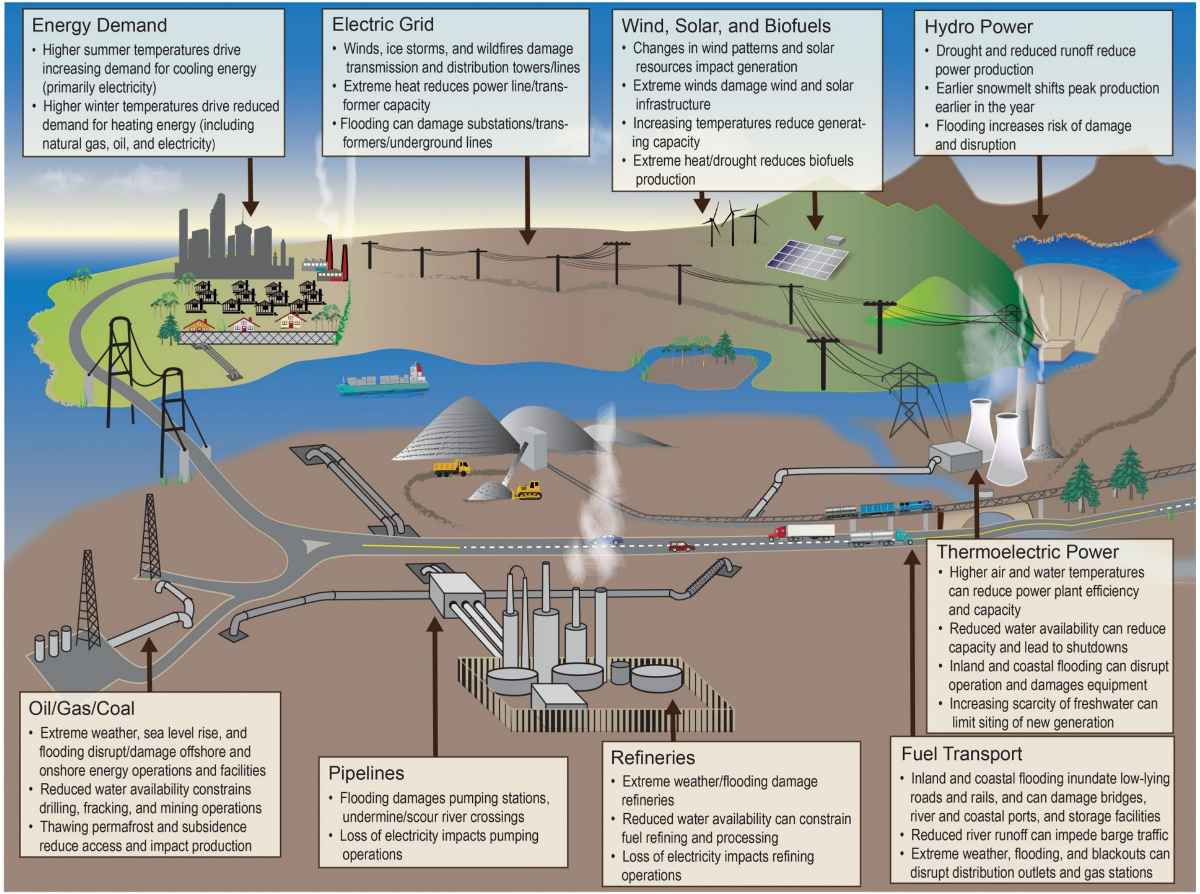

How is the grid balanced?

Solving the problem of “optimal power flow”: what energy should go where.

The objective of OPF is to find a steady state operating point that ***minimizes the cost of electric power generation while obeying physical constraints and meeting demand.***

Formulation	Variables	Number of variables	Number of equations
BIM	$V_i = (v_i e^{j\delta_i})$ $S_g^G = (p_g^G + jq_g^G)$ $S_l^L = (p_l^L + jq_l^L)$	$N + G + L$ $(2N + 2G + 2L)$	N $(2N)$
BFM	$V_i = (v_i e^{j\delta_i})$ $S_g^G = (p_g^G + jq_g^G)$ $S_l^L = (p_l^L + jq_l^L)$ $I_{ij}^s = (i_{ij}^s e^{j\gamma_{ij}^s})$ $S_{ij} = (p_{ij} + jq_{ij})$ $S_{ji} = (p_{ji} + jq_{ji})$	$N + G + L + 3E$ $(2N + 2G + 2L + 6E)$	$N + 3E$ $(2N + 6E)$

This is a really hard computational problem that scales with the size of the grid...and it needs to be solved every 3-5 minutes!

Can we use machine learning to help speed up OPF calculations?

Graph Neural Networks

The power grid can be represented as a network, or "graph"

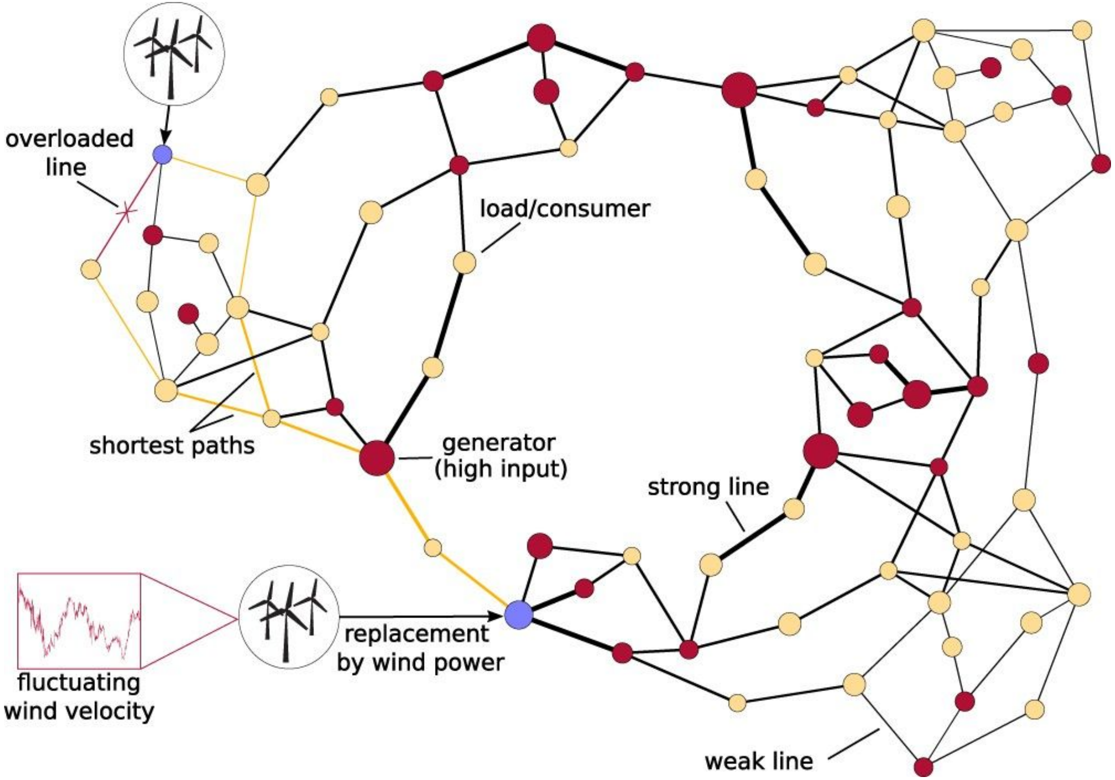

Graphs

Graphs are made of nodes that are connected via edges

"Topology" refers to the overall shape of the network, which is defined by an "adjacency matrix" (a matrix indicating which nodes have edges between them)

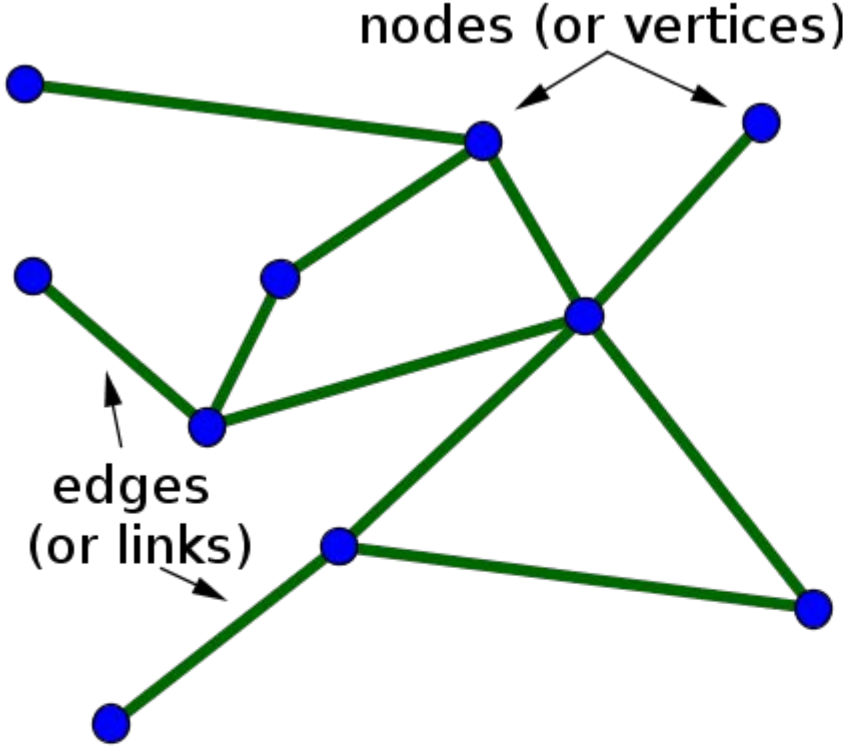

Graph Neural Networks (GNNs)

GNNs are artificial neural networks that can take *graphs* as input.

The graphs are represented by their adjacency matrix and any values needed to provide information about each node or edge.

The neural network learns how to combine information across nodes using a *message passing* algorithm.

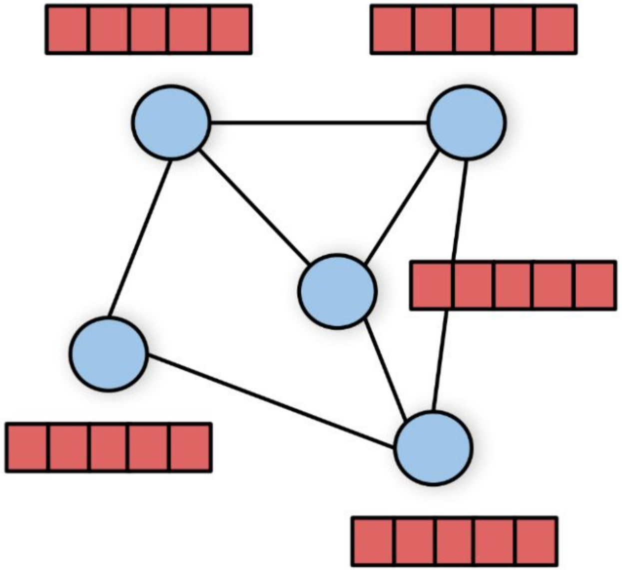

Message passing

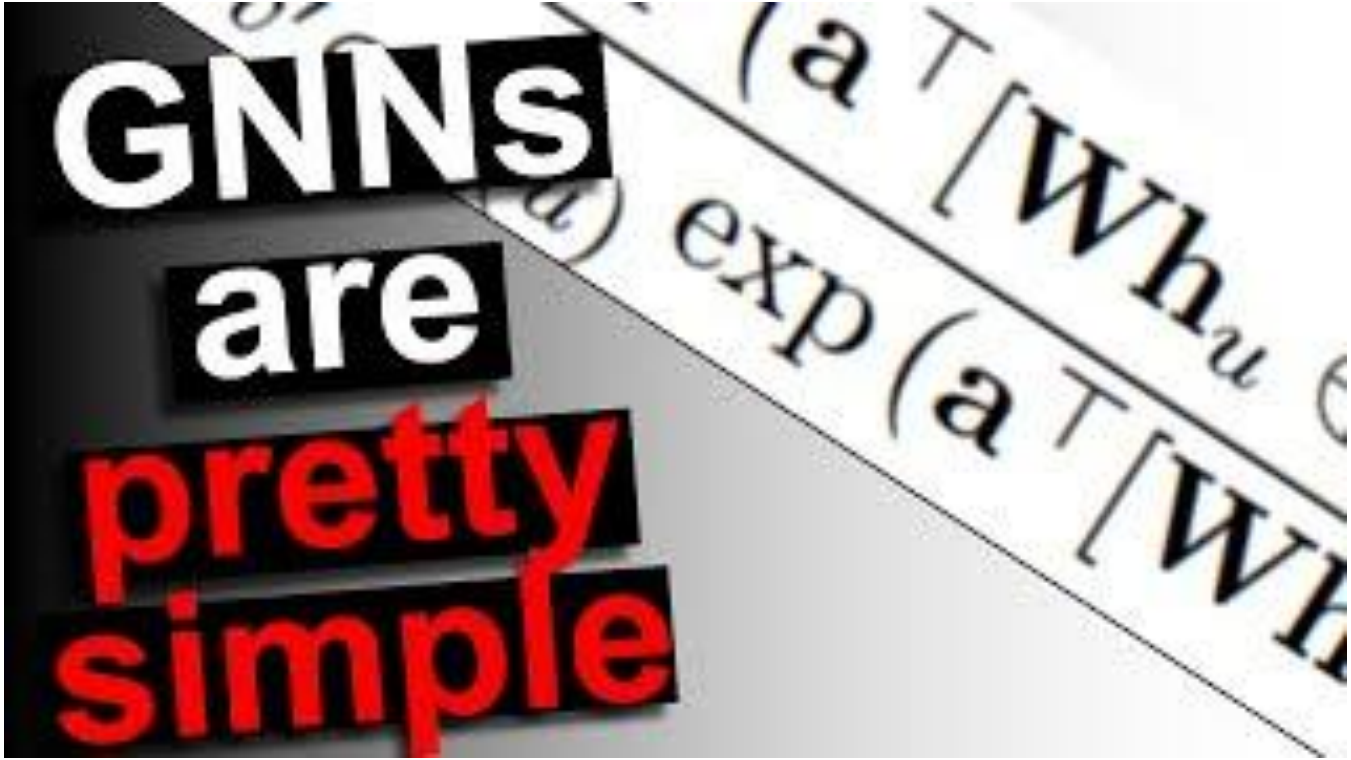

For your reading:

Bus = node in the power grid network

Load = something that consumes electricity

Line rating limit = the maximum power a line can safely conduct

IEEE provides simulated grid data based on real US grid properties.