ML4CC: Lecture 9

Sit with your **new** discussion groups!

Assignments reminder

Keep doing your PMIRO+Q

Your project assignment is due April 4th by 8am.

Summary of last paper

P - Want to know what climate change content in a tweet makes it engaging

M - Use an LLM plus topic clustering to get a tweet’s topic and use that (plus metadata) to predict which of two tweets is more engaging

I - predicting engagement specifically for climate topics

R - the model performs above chance and on par with a human

O - Is this a good metric of engagement? Would writing tweets based on these topics actually help?

Climate Change in the News

NVIDIA.

Watch NVIDIA CEO Jensen Huang's GTC keynote to catch all the announcements a

Newsroom NVIDIA in Brief Exec Bios NVIDIA Blog Podcast Media Assets In the

Press Release

NVIDIA Announces Earth Climate Digital Twin

The Weather Company, Central Weather Administration of Taiwan Among First to Adopt New Earth-2 Cloud APIs, Using AI to Speed Creation of High-Resolution Simulations and Visualization of Global Climate, Weather at Groundbreaking 2-Kilometer Scale

March 18, 2024

Groundbreaking Generative AI for Climate Tech

Earth-2's APIs offer AI models and employ a new NVIDIA generative AI model called CorrDiff, using state-of-the-art diffusion modeling, that generates 12.5x higher resolution images than current numerical models 1,000x faster and 3,000x more energy efficiently. It corrects inaccuracies of coarse-resolution forecasts and synthesizes metrics critical to decision-making. CorrDiff is a first-of-its-kind generative AI model to deliver super-resolution, synthesize new metrics of interest to stakeholders, and learn the physics of fine-scale local weather from high-resolution datasets.

The Central Weather Administration of Taiwan plans to use these diffusion models to forecast more precise locations of typhoon landfall. When a typhoon warning is launched, the priority is to minimize casualties by carrying out early evacuations based on quality information generated by relevant agencies, including Taiwan's National Science and Technology Center for Disaster Reduction (NCDR). In the last decade, the death toll due to typhoons has fallen.

"Taiwan is a critical component of the global supply chain, and flooding risk analysis and evacuation preparedness are core to our mandate," said Chia-Ping Cheng, administrator of CWA.

With more than 136 typhoons striking the island since 2000, using Earth-2 to mitigate these impacts is key to improving the quality and resolution of disaster informatics, NCDR said.

Paper 7 Discussion

Multi-Objective Optimization for Value-Sensitive and Sustainable Basket Recommendations

Thomas Asikis
asikist@ethz.ch
asikis@soziologie.uzh.ch

[1]Chair of Computational Social Science, ETH Zurich, Zurich, Switzerland
[2]Game Theory, UZH Zurich, Zurich, Switzerland

Abstract

Sustainable consumption aims to minimize the environmental and societal impact of the use of services and products. Over-consumption of services and products leads to potential natural resource exhaustion and societal inequalities as access to goods and services becomes more challenging. In everyday life, a person can simply achieve more sustainable purchases by drastically changing their lifestyle choices and potentially going against their personal values or wishes. Conversely, achieving sustainable consumption while accounting for personal values is a more complex task as potential trade-offs arise when trying to satisfy environmental and personal goals. This article focuses on value-sensitive design of recommender systems, which enable consumers to improve the sustainability of their purchases while respecting personal and societal values. Value-sensitive recommendations for sustainable consumption are formalized as a multi-objective optimization problem, where each objective represents different sustainability goals and personal values. Novel and existing multi-objective algorithms calculate solutions to this problem. The solutions are proposed as personalized sustainable basket recommendations to consumers. These recommendations are evaluated on a synthetic dataset, which comprises three established real-world datasets from relevant scientific and organizational reports. The synthetic dataset contains quantitative data on product prices, nutritional values, and environmental impact metrics, such as greenhouse gas emissions and water footprint. The recommended baskets are highly similar to consumer purchased baskets and aligned with both sustainability goals and personal values relevant to health, expenditure, and taste. Even when consumers would accept only a fraction of recommendations, a considerable reduction of environmental impact is observed.

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

Did you do anything fun over spring break?

Discussion Question 2

What is the “intended basket”? What does it represent and how is it represented mathematically? Why is it called “intended”?

Intended Basket

Value-sensitive sustainable recommendations are formalized as a multi-objective optimization problem of selecting combinations of discrete quantities over $N = 132$ distinct products. First, an intended basket is defined as the purchased weekly basket, i.e. a vector of non-negative integer product quantities $\boldsymbol{x}^*_{k,q} \in \mathbb{N}_0^N$ for a specific household k at week q. In a real-world application,

Example of a real set of items that were purchased from a household in a week, according to the dataset. Called “intended” here because we are using it as the starting point for our recommendations.

A vector with 132 entries, corresponding to the number of each item purchased

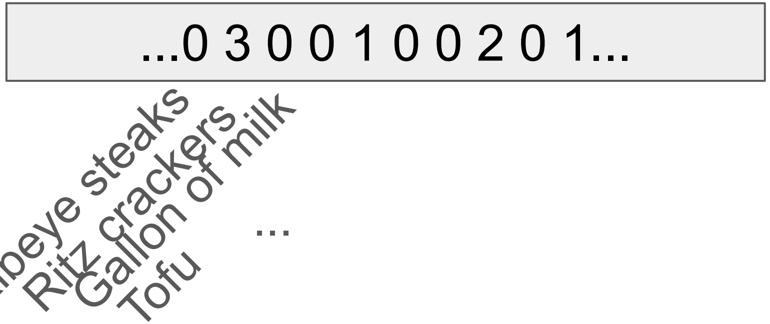

Discussion Question 3

Explain what table 2 is showing in your own words

The multiple objectives

This table shows all the things we are trying to optimize our recommendation for

j	Feature	Unit	Scope	Target
1	Cosine similarity	-	Personal	Max.
2	Cost	Dollars ($)	Personal	Min.
3	Energy	kilo Calories (kCal)	Personal	Pres.
4	Protein	grams (g)	Personal	Pres.
5	Fat	grams (g)	Personal	Pres.
6	GHG emissions	CO_2 kg eq.	Environment	Min.
7	Acidification pollution	SO_2 kg eq.	Environment	Min.
8	Eutrophication pollution	$PO_4{}^{-3}$ kg eq.	Environment	Min.
9	Land use	m^2	Environment	Min.
10	Water usage	L	Environment	Min.
11	Stressed water usage	L	Environment	Min.

TABLE 2: Features relevant to objectives for basket selection. The first column shows the index of each feature, which also coincides with their position in the ordered set C. Feature, Unit, and Scope columns give a brief overview of the objectives and finally the Target column describes whether the goal of the optimization is to minimize, maximize, or preserve the intended basket value.

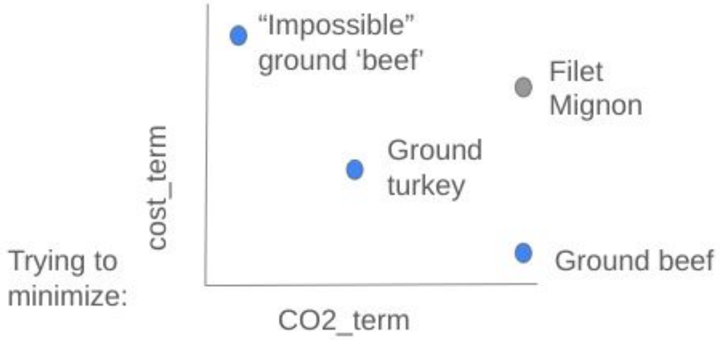

The multiple objectives

This table shows all the things we are trying to optimize our recommendation for

j	Feature	Unit	Scope	Target
1	Cosine similarity	-	Personal	Max.
2	Cost	Dollars ($)	Personal	Min.
3	Energy	kilo Calories (kCal)	Personal	Pres.
4	Protein	grams (g)	Personal	Pres.
5	Fat	grams (g)	Personal	Pres.
6	GHG emissions	CO_2 kg eq.	Environment	Min.
7	Acidification pollution	SO_2 kg eq.	Environment	Min.
8	Eutrophication pollution	$PO_4{}^{-3}$ kg eq.	Environment	Min.
9	Land use	m^2	Environment	Min.
10	Water usage	L	Environment	Min.
11	Stressed water usage	L	Environment	Min.

TABLE 2: Features relevant to objectives for basket selection. The first column shows the index of each feature, which also coincides with their position in the ordered set C. Feature, Unit, and Scope columns give a brief overview of the objectives and finally the Target column describes whether the goal of the optimization is to minimize, maximize, or preserve the intended basket value.

Whether we want to minimize, maximize, or preserve (i.e. not change) this value

The multiple objectives

This table shows all the things we are trying to optimize our recommendation for

j	Feature	Unit	Scope	Target
1	Cosine similarity	-	Personal	Max.
2	Cost	Dollars ($)	Personal	Min.
3	Energy	kilo Calories (kCal)	Personal	Pres.
4	Protein	grams (g)	Personal	Pres.
5	Fat	grams (g)	Personal	Pres.
6	GHG emissions	CO_2 kg eq.	Environment	Min.
7	Acidification pollution	SO_2 kg eq.	Environment	Min.
8	Eutrophication pollution	$PO_4{}^{-3}$ kg eq.	Environment	Min.
9	Land use	m^2	Environment	Min.
10	Water usage	L	Environment	Min.
11	Stressed water usage	L	Environment	Min.

TABLE 2: Features relevant to objectives for basket selection. The first column shows the index of each feature, which also coincides with their position in the ordered set C. Feature, Unit, and Scope columns give a brief overview of the objectives and finally the Target column describes whether the goal of the optimization is to minimize, maximize, or preserve the intended basket value.

The cosine similarity simply encourages the model to make as few changes as possible

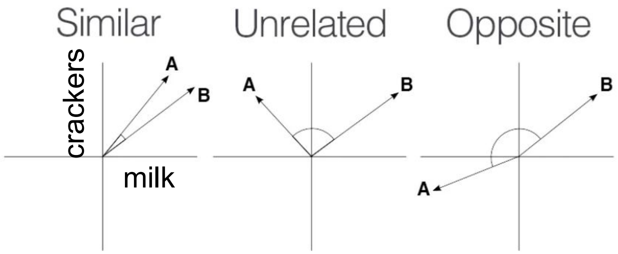

Discussion Question 4

Explain why the intended basket is one solution that minimizes Eqn 5. Why are we aiming to minimize Eqn 5 when Table 2 says features 3-5 should be preserved?

Next, the nutritional values of a recommended basket are considered for optimization. For each unit of product i and nutritional product feature j the nutritional quantity per unit $c_{i,j}$ is calculated. Three nutritional features are denoted by indices $j \in \{3, 4, 5\}$. The health objective functions use the intended basket nutritional value as a baseline to evaluate the difference for each nutritional feature between recommended and intended baskets:

$$J_j\left(\boldsymbol{x}, \hat{\boldsymbol{x}}\right) = \left(1 - \rho_j(\boldsymbol{x}, \boldsymbol{x}^*)\right)^2 = \left(\frac{v_j(\boldsymbol{x}^*) - v_j(\boldsymbol{x})}{v_j(\boldsymbol{x}^*)}\right)^2, j \in \{3, 4, 5\}. \quad (5)$$

The intended basket is one solution that minimizes the Relation (5).

To preserve we must minimize change

$$J_j\left(\boldsymbol{x}, \hat{\boldsymbol{x}}\right) = \left(1 - \rho_j(\boldsymbol{x}, \boldsymbol{x}^*)\right)^2 = \left(\frac{v_j(\boldsymbol{x}^*) - v_j(\boldsymbol{x})}{v_j(\boldsymbol{x}^*)}\right)^2, j \in \{3, 4, 5\}. \quad (5)$$

The intended basket is one solution that minimizes the Relation (5).

The synthetic dataset provides coefficients $c_{i,j}$, which in this study are calculated based on the mean[1] values over all transactions in the dataset and describe the corresponding feature j quantity $c_{i,j}$ per unit for each product i. Therefore, for a basket $\boldsymbol{x}$, one can calculate the total quantity for a specific feature as

$$v_j(\boldsymbol{x}) = \sum_{i=1}^{N} c_{i,j} x_i. \quad (1)$$

When designing the objective functions and comparing recommendations among baselines, often the ratio of total feature quantities between two baskets $\boldsymbol{x}, \boldsymbol{x}'$

$$\rho_j(\boldsymbol{x}, \boldsymbol{x}') = \frac{v_j(\boldsymbol{x})}{v_j(\boldsymbol{x}')} \quad (2)$$

is used. In particular, most calculations related to environmental and personal objectives use the ratio of a recommendation towards the intended basket $\rho_j(\boldsymbol{x}, \boldsymbol{x}^*)$ for a specific feature j.

Equation 2 describes the ratio between nutrients in two baskets. If this is 1, the baskets have identical nutrients and the loss term in Eqn 5 is zero

Discussion Question 5

Explain this definition of "dominated" in your own terms

algorithm is to find a non-dominated set of baskets. A basket $\boldsymbol{x}$ dominates $\boldsymbol{x} \prec \boldsymbol{x}'$ another basket $\boldsymbol{x}'$ if $J_j(\boldsymbol{x}) \leq J_j(\boldsymbol{x}')$ for all $j \in C$ and $J_j(\boldsymbol{x}) < J_j(\boldsymbol{x}')$ holds at least for one j [13]. If no other basket dominates $\boldsymbol{x}$, then it is referred as non-dominated.

Dominated solutions

To dominate another solution, you must be better on at least one dimension while being worse on none.

algorithm is to find a non-dominated set of baskets. A basket $\boldsymbol{x}$ dominates $\boldsymbol{x} \prec \boldsymbol{x}'$ another basket $\boldsymbol{x}'$ if $J_j(\boldsymbol{x}) \leq J_j(\boldsymbol{x}')$ for all $j \in C$ and $J_j(\boldsymbol{x}) < J_j(\boldsymbol{x}')$ holds at least for one j [13]. If no other basket dominates $\boldsymbol{x}$, then it is referred as non-dominated.

Ground beef dominates Filet Mignon but not ground turkey (because it is worse on cost)

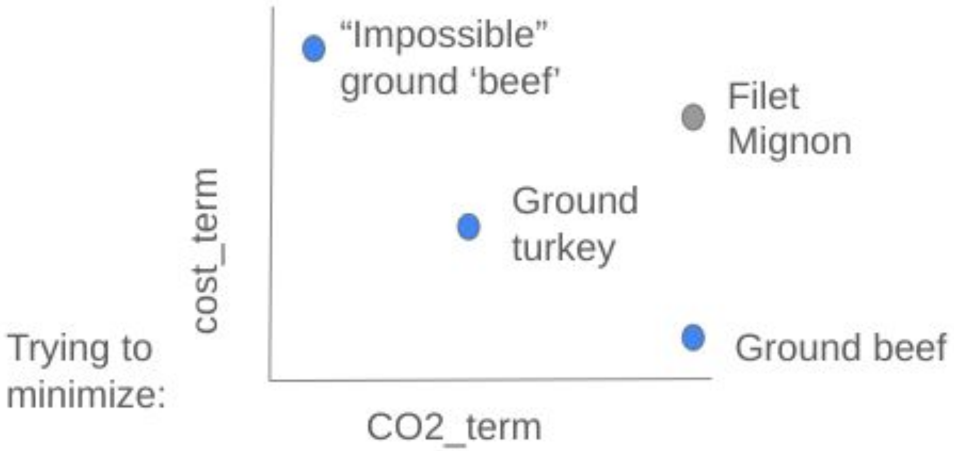

Discussion Question 6

What is a reference point in RNSGA-II? What problem does having a reference point solve?

Selecting the best solutions

Problem: The "front" of non-dominated solutions is still going to include some not very good solutions

3.0.1 RNSGA-II

A non-dominated sorting algorithm may produce a large number of non-dominated solutions that are not preferable, e.g. solutions that optimize a single objective very well and not the others. To keep the population size B per generation constant, a secondary selection operation needs to be performed. Random selection is often undesired in problems that have multiple objectives [12], and thus a more sophisticated technique is preferred. Some probabilistic evolutionary algorithms use a sorting operation to perform a secondary selection operation that guide the evolutionary processes towards preferred non-dominated solutions, e.g. non-dominated solutions that optimize specific combinations of the objectives very well. A typical example that will be used as a baseline in the current study is reference point NSGA-II, abbreviated as RNSGA-II [12], which uses reference directions to guide evolution towards preferred solutions. In brief, one or more reference points are selected to guide the evolution. A reference point $\hat{\zeta}$ is generated by providing a vector of preferred objective values to the system. Each candidate solution receives two ranks determined by the non-dominated sorting and a distance metric from each reference point, i.e. lower distance values receive lower ranks. Lower ranks are used to select the candidates for the next generation. This algorithm

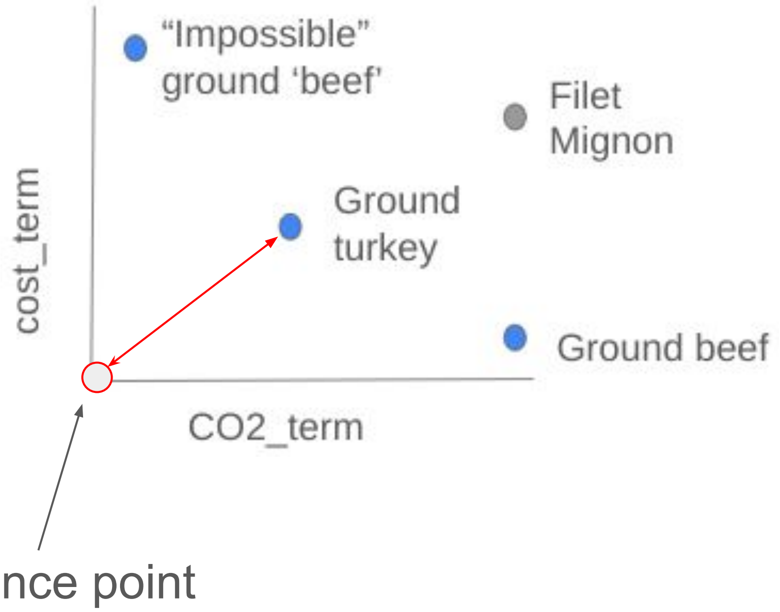

Discussion Question 7

Do the three methods tested here tend to find the same recommended baskets? How do you know?

Different methods find different solutions

4.1 Recommendation Comparison

First, the ability of baselines to produce non-dominated solutions for the problem is evaluated. Table 3 contains a comparison where all recommendations for an intended basket $\boldsymbol{x}^*$ from all methods are compared against each other and only the non-dominated solutions are kept across all methods. The ratio of total non-dominated solutions divided by total recommendations per method is calculated. All three baselines produce diverse non-dominated solutions, as they all achieve high mean ratio of non-dominated to total recommended baskets per intended basket. This indicates that the problem can be tackled effectively by all methods.

Model	Mean	Mean CI	Median	Median CI
G3A	0.980	(0.979, 0.981)	1.0	(1.0, 1.0)
MO-NES	0.948	(0.946, 0.949)	1.0	(1.0, 1.0)
RNSGA-II	0.986	(0.985, 0.986)	1.0	(1.0, 1.0)

If you put the best solutions from all three methods together (represented by different colors) they all contribute to the non-dominated front:

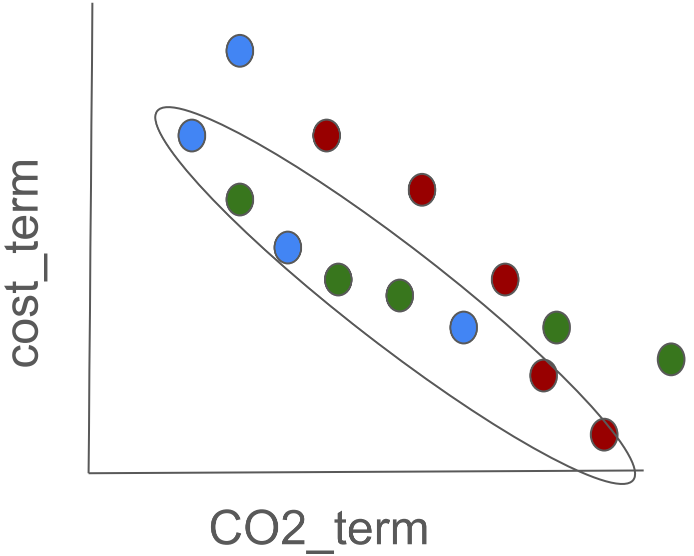

Different methods find different solutions

Different methods achieve different values for the different objective terms

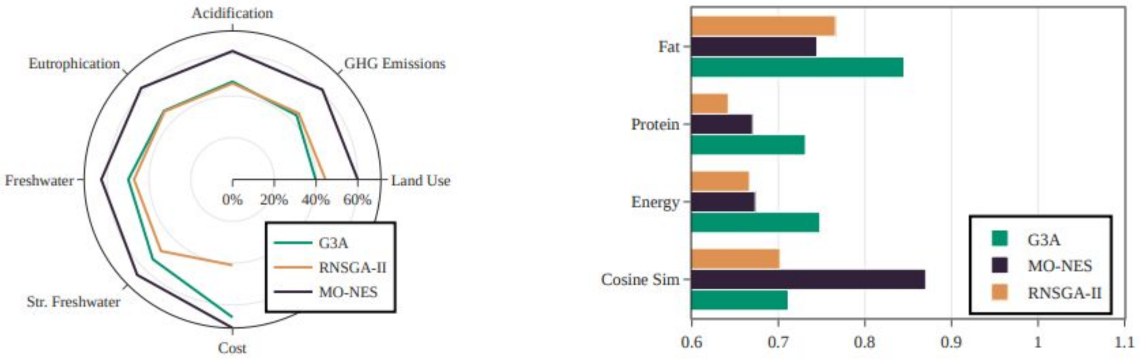

(a) Cost and Environmental Impact Ratios (lower values - points closer to center are preferred).

(b) Nutritional ratios and cosine similarity (higher values - longer bars are preferred).

FIGURE 2: A comparison of cosine similarity and the total emission, nutritional, and cost of a recommendation, as a ratio to the corresponding intended basket. For each baseline the mean ratio value over all recommendations that achieve cosine similarity higher than 0.5 and have all environmental ratios costs below 1.0 are considered.

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: Transportation

Machine Learning: Reinforcement learning

Transportation emissions make up 16% of global total

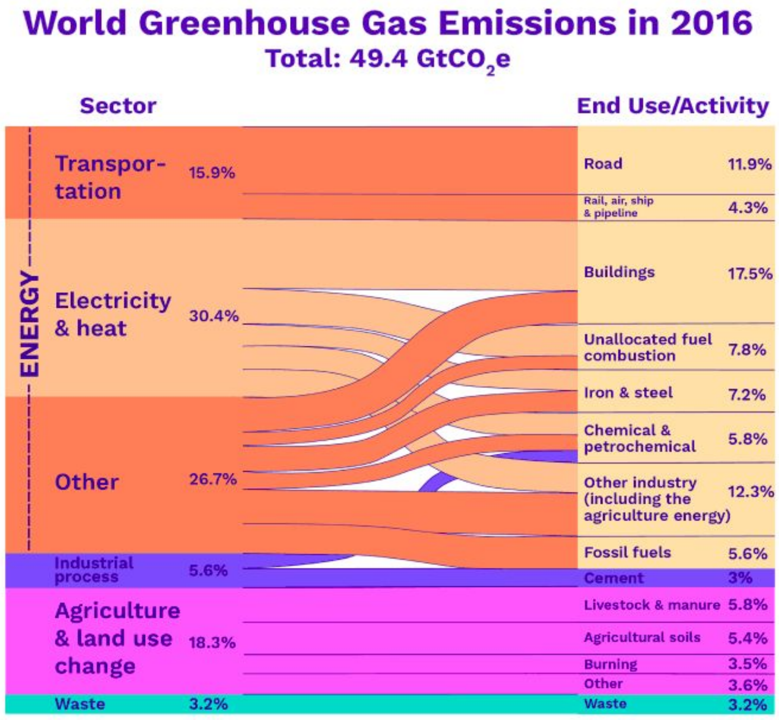

Source: Greenhouse gas emissions on Climate Watch. Available at: https://www.climatewatchdata.org

Breakdown of U.S. transport emissions

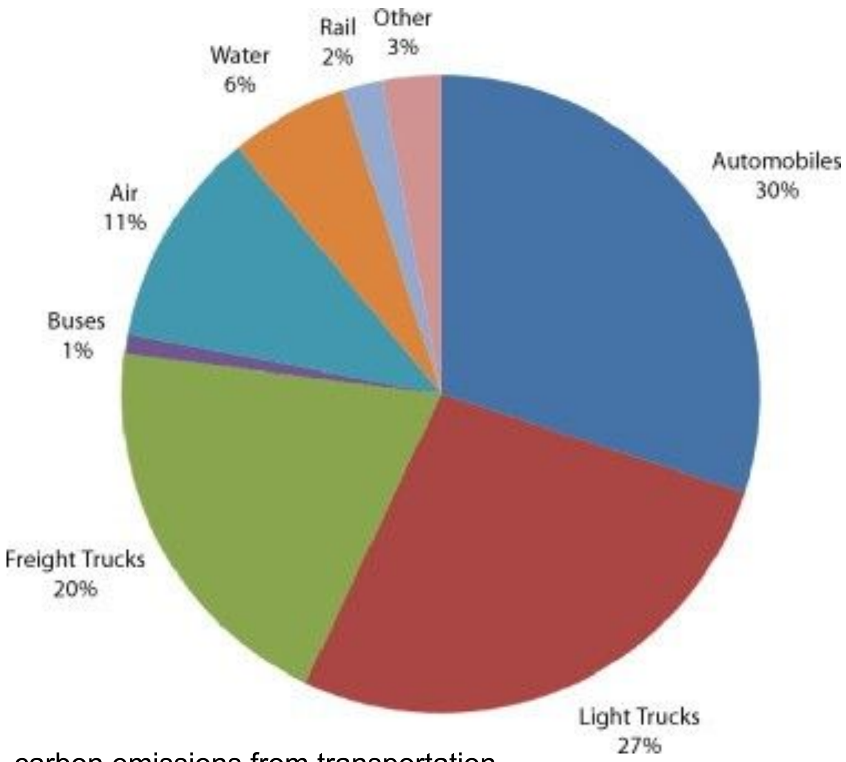

U.S. carbon emissions from transportation, 2005 (*Source*: EIA, 2007b).

The largest sector contributing to transportation emissions is personal car use. Aviation is the worst per distance emitter.

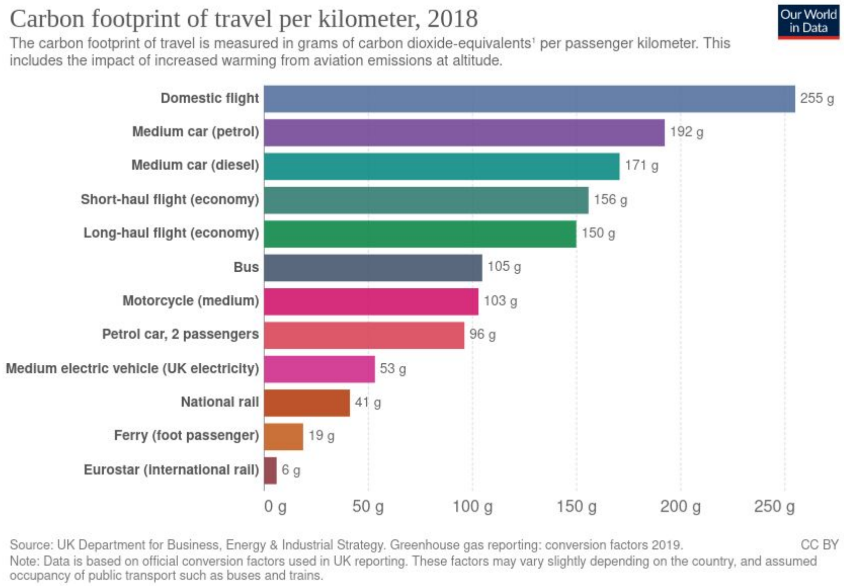

Breakdown of global transport emissions over time

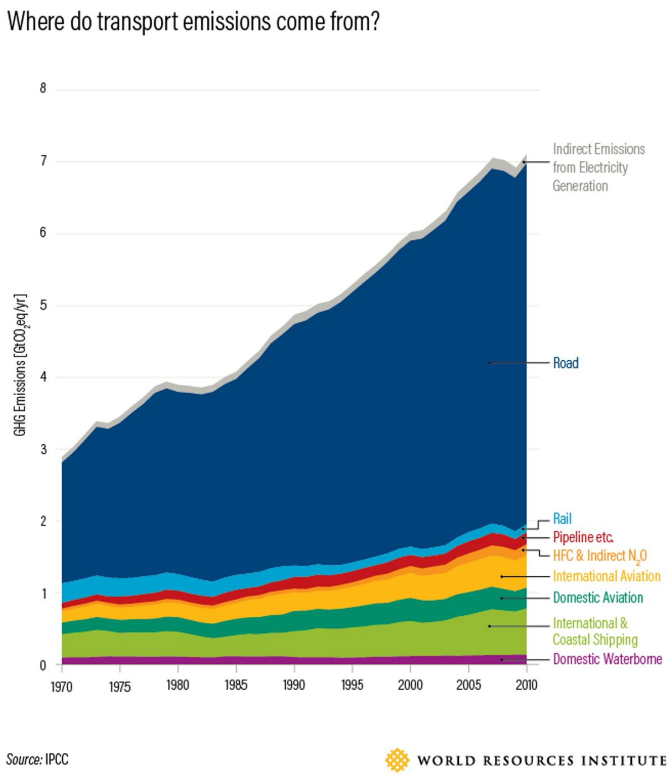

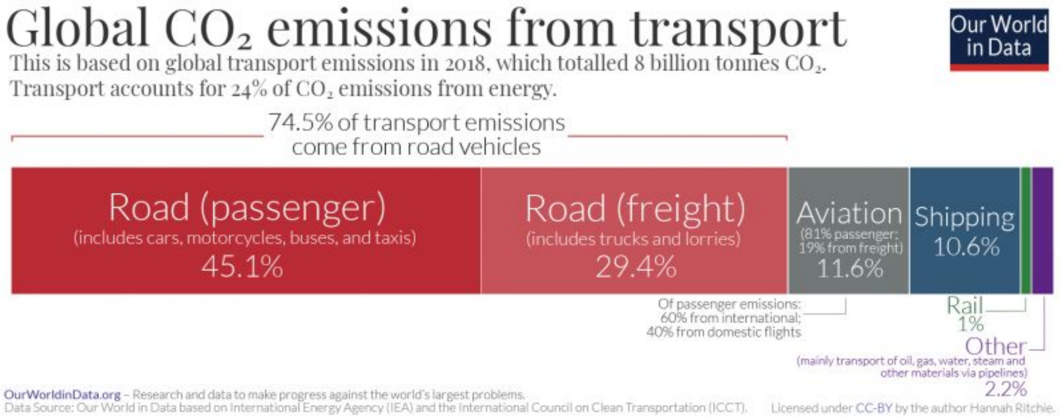

"Average annual greenhouse gas emissions growth between 2010 and 2019 slowed compared to the previous decade in energy supply (from 2.3% to 1.0%) and industry (from 3.4% to 1.4%), **but remained roughly constant at about 2% per year in the transport sector**." -IPCC

Transport emissions by country

The US is the biggest transportation emitter

Greenhouse gas emissions of the transportation sector worldwide in 2021, by select country (in million metric tons of carbon dioxide equivalent)

Worldwide; European Commission; EDGAR/JRC; Expert(s) (Crippa et al.); 2021

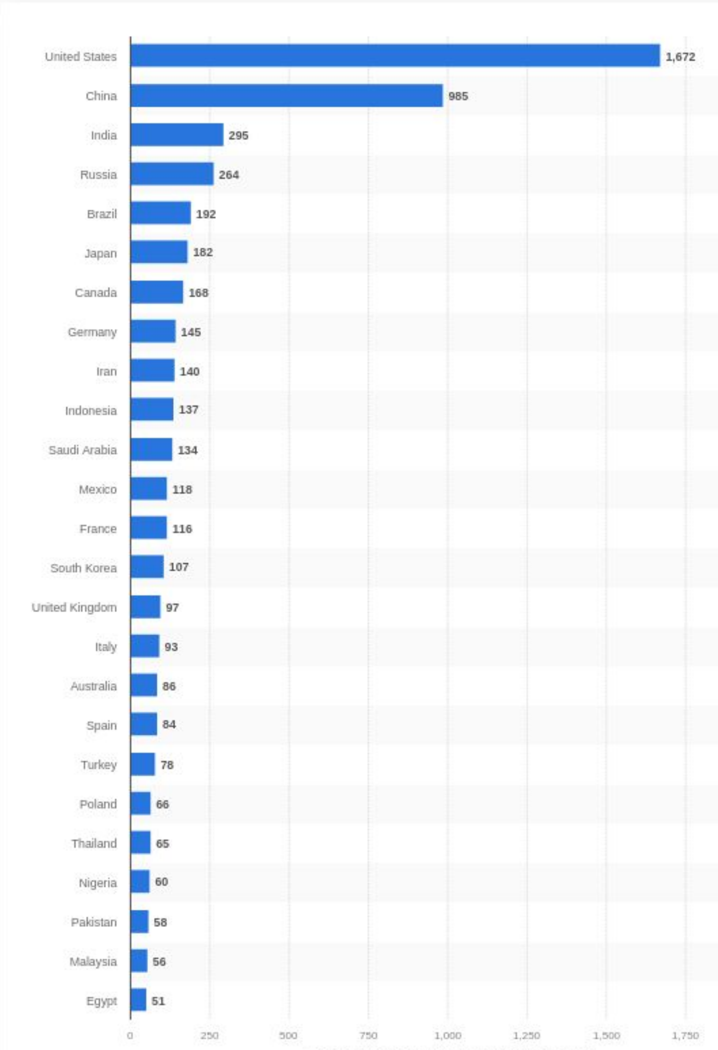

How are transportation emissions calculated?

How are transportation emissions calculated?

“In 1990, the Federal Clean Air Act was amended in an effort to greatly reduce air pollution. As a result, the Environmental Protection Agency devised a set of emissions standards to minimize the amount of hazardous air pollutants released by motor vehicles. This means your car may have to undergo periodic testing to ensure it’s within EPA standards and is limiting its negative impact on the environment.”

nationwide

How are transportation emissions calculated?

Urban cycle test:

- Accelerate to 9mph in four seconds
- Cruise at 9mph for eight seconds
- Brake to rest in five seconds
- Accelerate to 20mph over 12 seconds
- Cruise at 20mph for 24 seconds
- Brake to rest in 11 seconds
- Accelerate to 31mph over 26 seconds
- Cruise at 31mph for 12 seconds
- Brake to 22mph over eight seconds
- Cruise at 22mph for 13 seconds
- Brake to rest in 12 seconds

carthrottle.com

How are transportation emissions calculated?

Bottom up method:

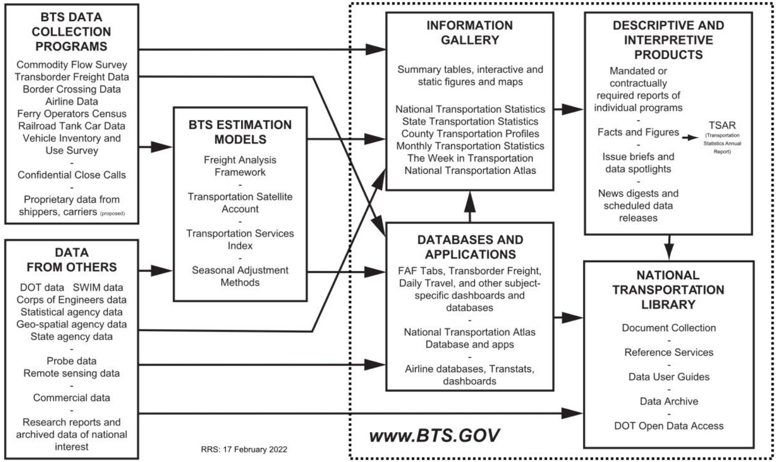

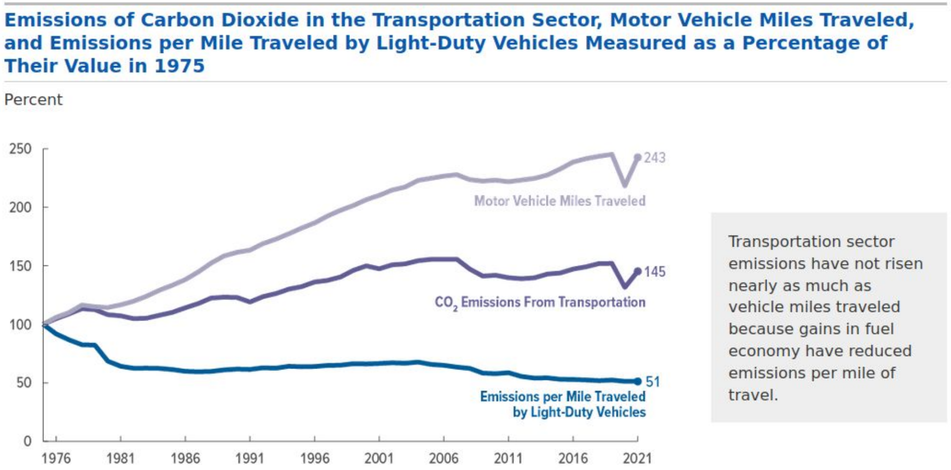

CBO

The need to reduce transportation emissions

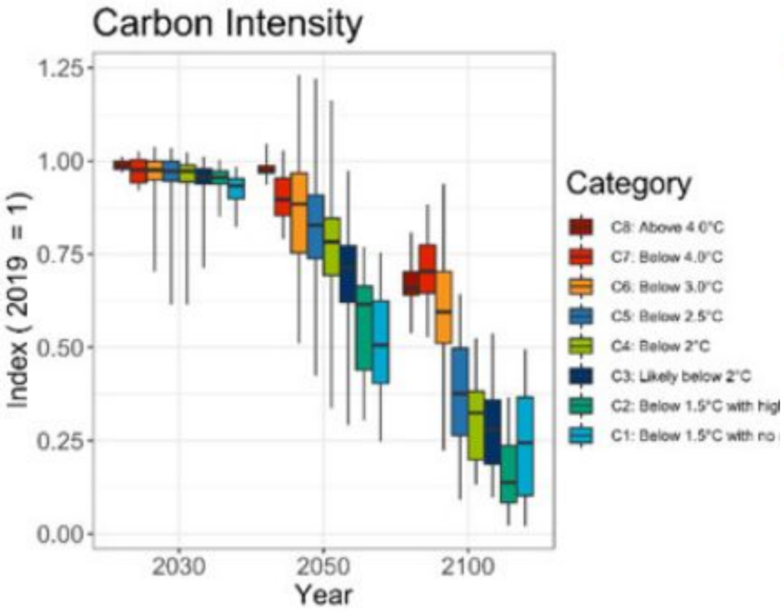

IPCC projections for different warming scenarios show the need for a reduction in how much carbon is emitted by transportation methods.

Figure 3.16 Global CO_2 emissions in transport by mode in the Sustainable Development Scenario, 2000-70

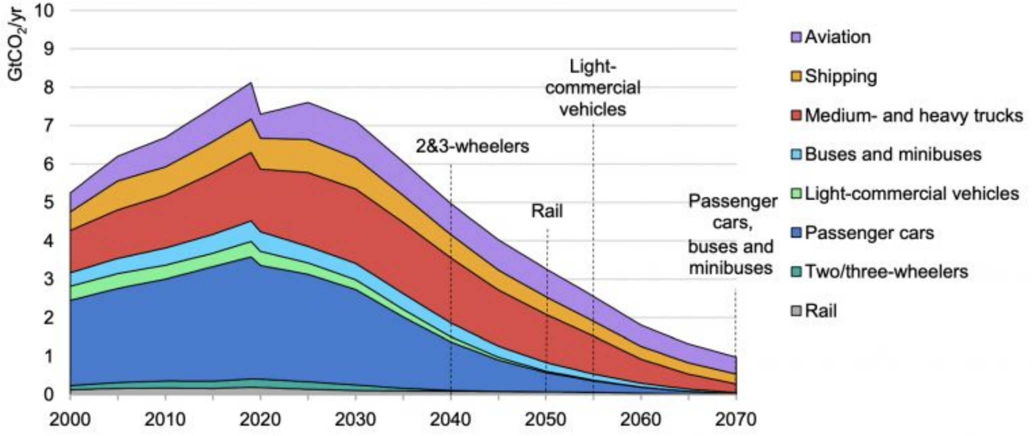

Notes: Dotted lines indicate the year in which various transport modes have largely stopped consuming fossil fuels and hence no longer contribute to direct emissions of CO_2 from fossil fuel combustion. Residual emissions in transport are compensated by negative emissions technologies, such as BECCS and DAC, in the power and other energy transformation sectors.

Reducing emissions from transportation

- Reduce amount of miles traveled
- Convert from high emissions forms of travel to low (aviation or single-person driving to land transport and public transportation)
- Lower emissions of fossil fuel-burning vehicles
- Switch to electric vehicles (with clean power grids)

E-bikes and scooters as alternatives to cars

The Weather Network

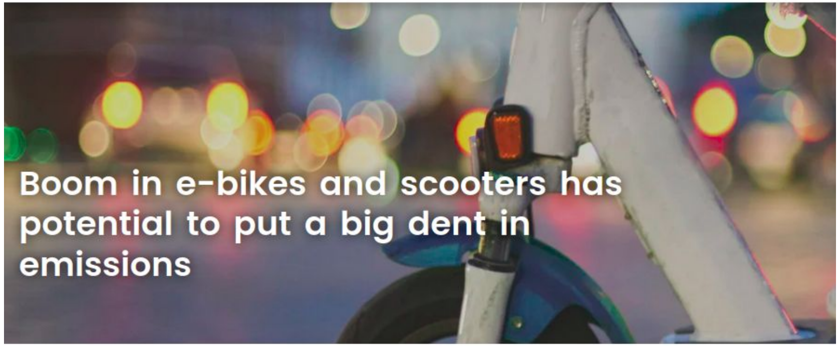

M.A. Jacquemain
Contributor

Published on Nov. 1, 2022, 12:43 PM

Amid rising gas prices this year, many commuters have made the switch to e-bikes and scooters, a move that has potential to drastically slash emissions in the process.

This new transportation trend has caught on in major Canadian cities, with the comparably affordable price of e-bikes and scooters seen by many customers as a way to offset historically high costs at the pumps.

Other converts have been won over by the flexibility e-bikes and scooters offer in terms of parking, storage, and navigability in traffic.

A study released earlier this year determined that the CO2 emissions produced by the energy required to run e-bikes average about 22 g/km, while gas-powered cars emit more than 250 g/km.

The research, which focused on e-bike use in England, determined that 24.4 million tons of CO2 emissions could be avoided by adopting a greater use of e-bikes, or a savings of as much as 750 kg of CO2 per person yearly.

E-bikes and scooters as alternatives to walking...

SHARED E-SCOOTERS MAY NOT BE SO CLIMATE FRIENDLY AFTER ALL

JANUARY 3RD, 2022

POSTED BY ETH ZURICH

"Shared e-scooters and e-bikes in the city of Zurich primarily replace more sustainable modes of transport—walking, public transport, and cycling. This means that they emit more carbon than the means of transport they replace," says Daniel Reck. (Credit: Getty Images)

Increasing fuel efficiency: car size

Size	Cost[2]	Characteristics[3]
Subcompact	32.2	4 cylinder Avg MPG = 32
Compact	42.3	4 cylinder Avg MPG = 23
Intermediate	46.9	6 cylinder Avg MPG = 20
Full-Size Vehicle	51.1	6 cylinder Avg MPG = 19
Compact Pickup	40.2	4 cylinder Avg MPG = 18
Full-Size Pickup	47.7	8 cylinder Avg MPG = 13
Compact Utility	45.6	4 cylinder Avg MPG = 15
Intermediate Utility	51.4	6 cylinder Avg MPG = 15
Full-size Utility	52.9	8 cylinder Avg MPG = 13
Mini-Van	50.7	6 cylinder Avg MPG = 17
Full-Size Van	52.0	6 cylinder Avg MPG = 13

Cents Per Mile[1]

FHA

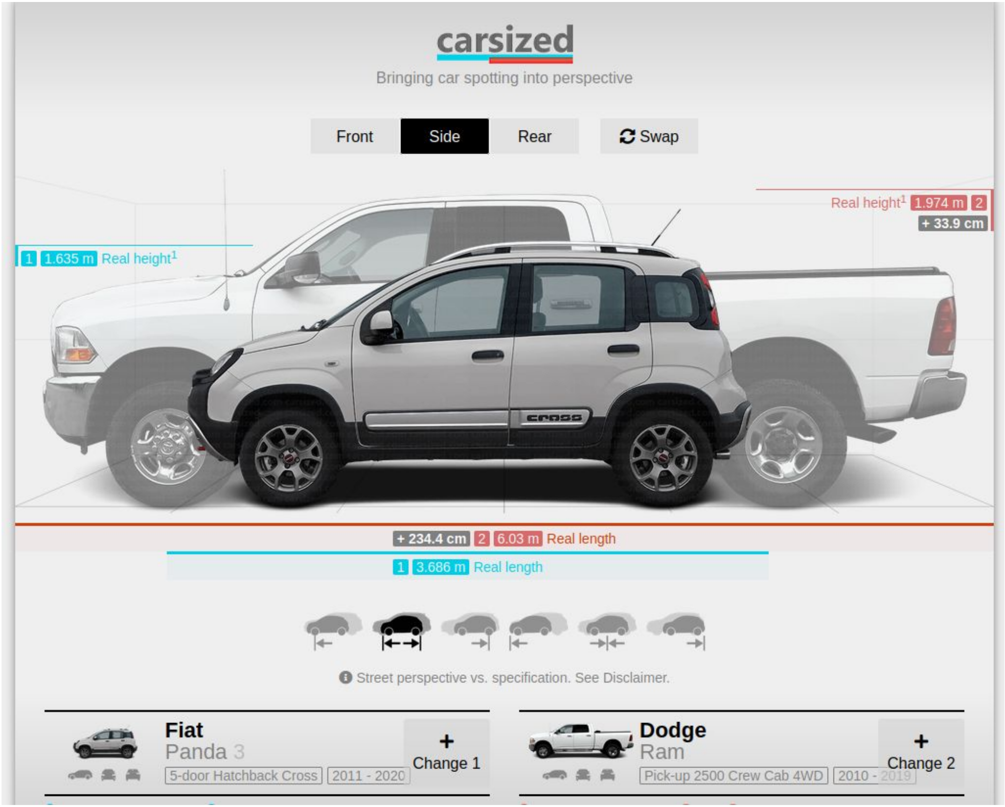

How different electric vehicles increase fuel efficiency

Electric vehicle adoption

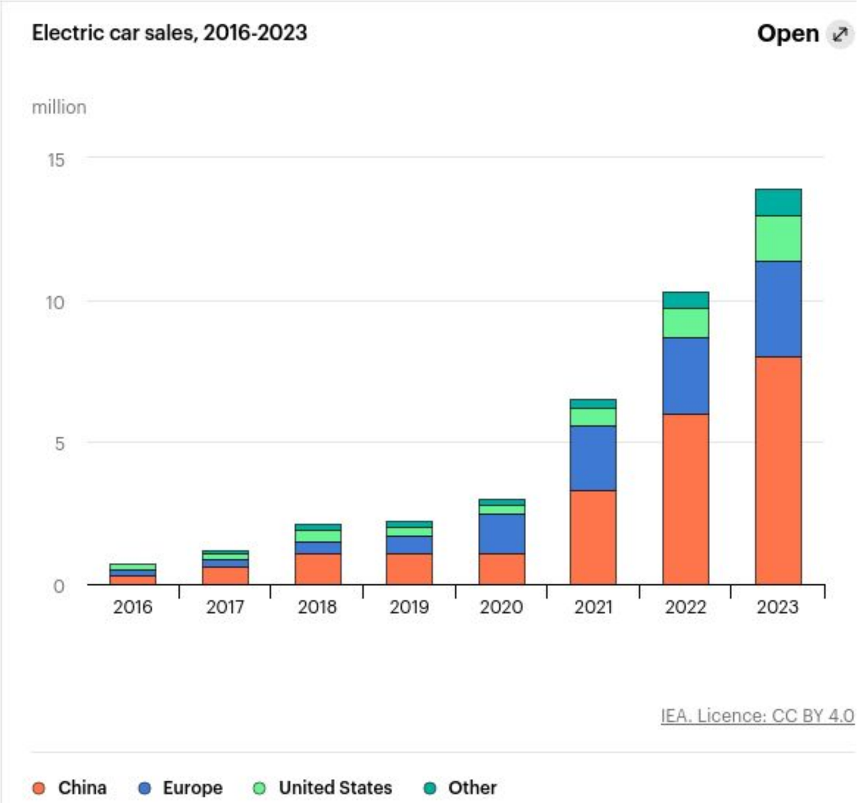

The switch to EVs is expected to avoid the use of over 5 million barrels of oil a day by 2030

Policies can encourage EV adoption

The New York Times

Biden Administration Announces Rule Aimed at Expanding Electric Vehicles

The regulation would essentially require automakers to sell more electric vehicles and hybrids by gradually tightening limits on tailpipe pollution.

Life Cycle Analysis (LCA)

Life Cycle Analysis refers to the process calculating emissions for a product based on the full supply, production, and disposal chain.

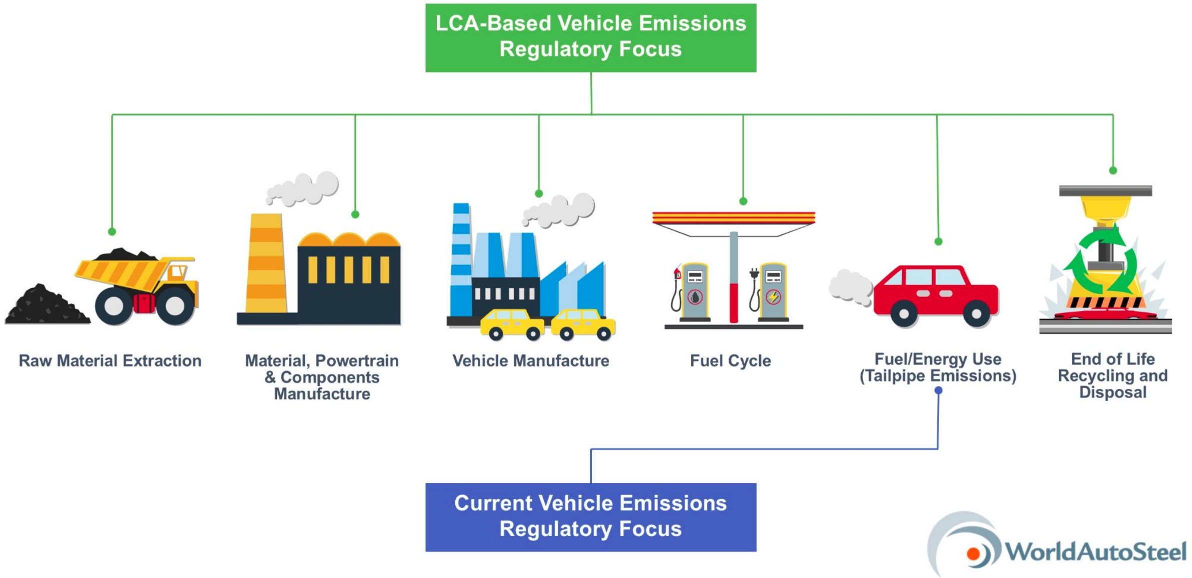

Life Cycle Analysis of Electric vs Gas Cars

ELSEVIER

Renewable and Sustainable Energy Reviews

Volume 159, May 2022, 112158

Total CO_2-equivalent life-cycle emissions from commercially available passenger cars

Johannes Buberger [a], Anton Kersten [a b], Manuel Kuder [a], Richard Eckerle [a], Thomas Weyh [a], Torbjörn Thiringer [b]

Show more

+ Add to Mendeley   Share   Cite

https://doi.org/10.1016/j.rser.2022.112158   Get rights and content

Highlights

- Quantification of the CO_2-equivalent greenhouse gas emissions of 790 different commercially available vehicle variants.
- The total life-cycle emissions of hybrid and electric vehicles are reduced by up to 89% compared to internal combustion engine vehicles.
- Modern battery recycling techniques can counterbalance the production emissions by about 60% to 65%.
- Vehicles powered by renewable fuels, such as compressed biogas, have a similar climate change impact as electric vehicles.

What powers the grid matters

The New York Times

How Green Are Electric Vehicles?

In short: Very green. But plug-in cars still have environmental effects. Here's a guide to the main issues and how they might be addressed.

Share full article

CLIMATE | How Green Are Electric Vehicles?

An all-electric Chevrolet Bolt, for instance, can be expected to produce 189 grams of carbon dioxide for every mile driven over its lifetime, on average. By contrast, a new gasoline-fueled Toyota Camry is estimated to produce 385 grams of carbon dioxide per mile. A new Ford F-150 pickup truck, which is even less fuel-efficient, produces 636 grams of carbon dioxide per mile.

But that's just an average. On the other hand, if the Bolt is charged up on a coal-heavy grid, such as those currently found in the Midwest, it can actually be a bit worse for the climate than a modern hybrid car like the Toyota Prius, which runs on gasoline but uses a battery to bolster its mileage. (The coal-powered Bolt would still beat the Camry and the F-150, however.)

"Coal tends to be the critical factor," said Jeremy Michalek, a professor of engineering at Carnegie Mellon University. "If you've got electric cars in Pittsburgh that are being plugged in at night and leading nearby coal plants to burn more coal to charge them, then the climate benefits won't be as great, and you can even get more air pollution."

Reinforcement Learning

Framing the problem

When using reinforcement learning to solve a control problem, the aim is to develop a **policy** that controls an **agent** to maximize **reward**.

The policy is a function that takes in a **state observation** and produces an **action**.

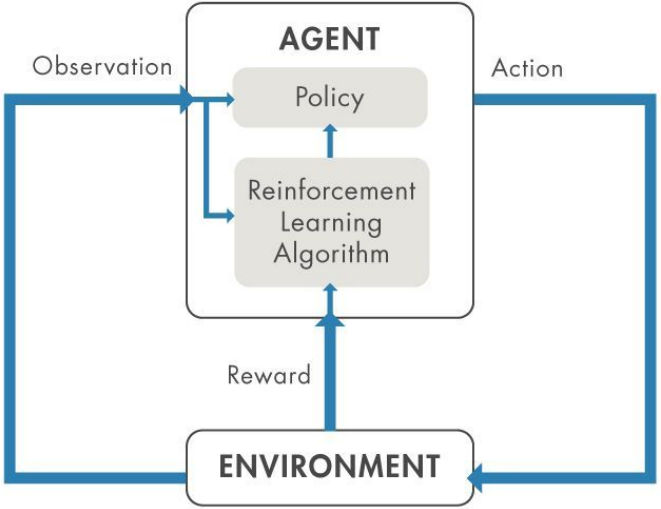

https://www.mathworks.com/company/newsletters/articles/reinforcement-learning-a-brief-guide.html

RL example

Agent = person

State = coordinates, walls, reward

Reward = star

Actions = up,down,left,right

Policy = what action to take when in a certain state

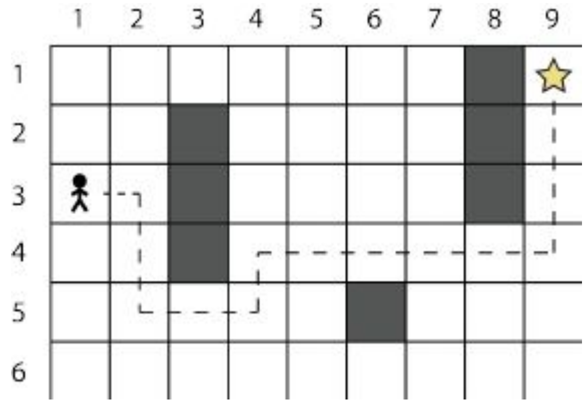

Reinforcement learning is hard because of...

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

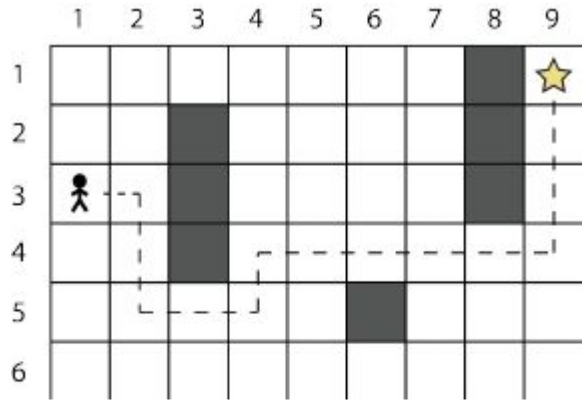

RL can be data intensive

RL agents can make a lot of mistakes while learning and need many iterations

They can also find unexpected solutions

How does RL work?

How do we decide the policy? That is, how do we learn how produce an action given a state observation?

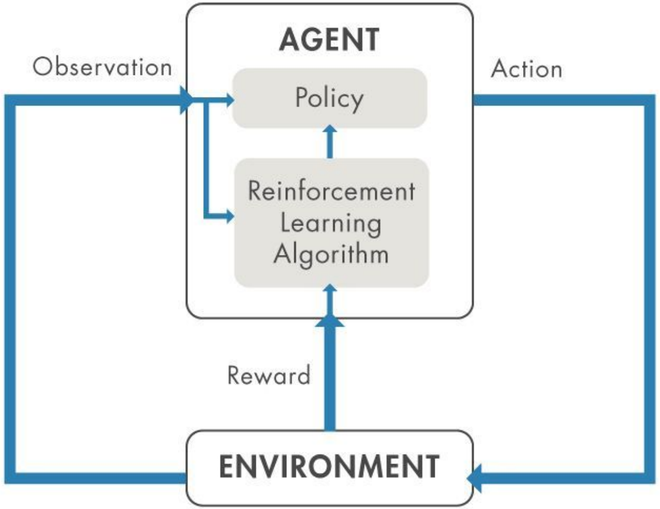

RL concepts

There are many algorithms in RL for building a good policy

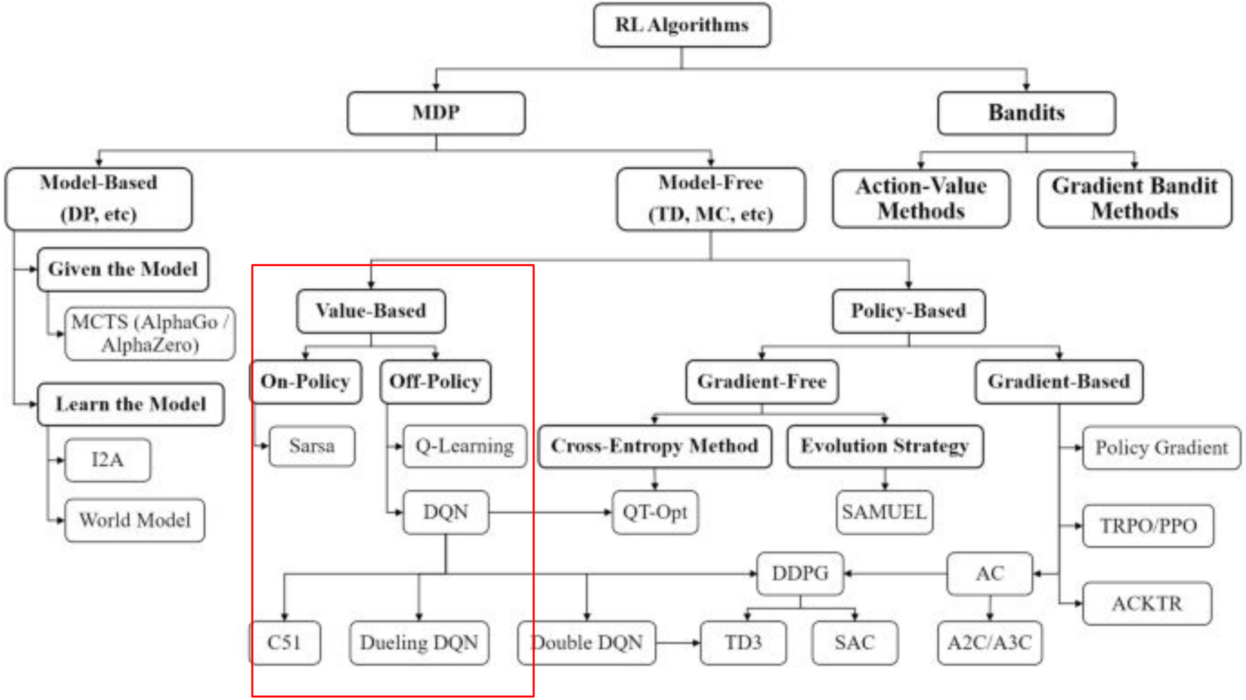

Let's look at Q-learning and DQN

Distribution shifts: a problem in RL (and all machine learning)

Whenever your model is run on data that deviates from what it was trained on (a “distribution shift”), its performance may suffer.

This is why we use various forms of validation to test for generalization performance.

Distribution shifts: a problem in RL (and all machine learning)

In RL, many different aspects of the data and its relationships can change, including:

-the statistics of the environment

-how actions impact the environment

-how the environment is observed

-what a reliable predictor of reward is

-reward magnitudes

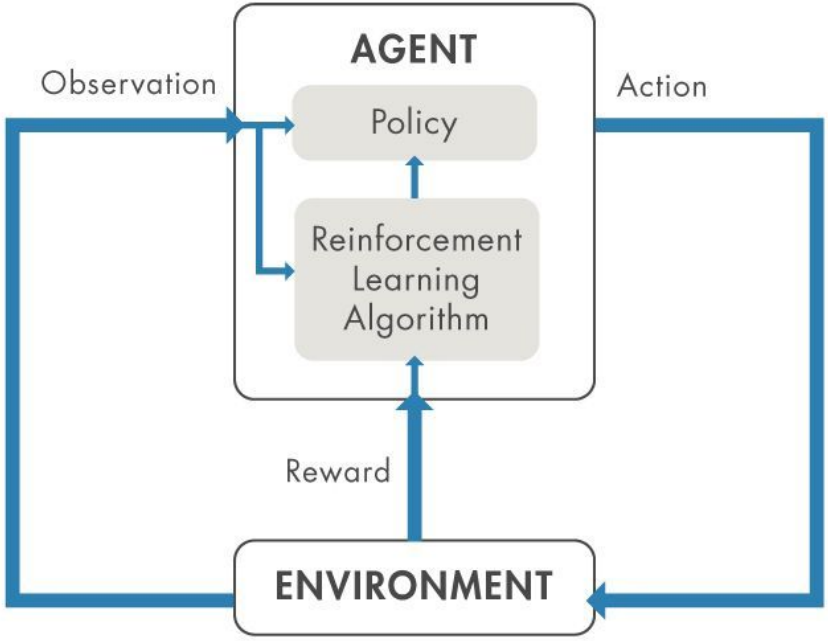

Benchmarks

Much research in machine learning is centered around benchmarks: well-defined tasks and datasets that anyone can train a model on. People aim to be the best on the "scoreboard". This spurs research and offers direct comparisons across approaches.

Well-known example: ImageNet

ImageNet Challenge

IMAGENET

- 1,000 object classes (categories).
- Images:
  - 1.2 M train
  - 100k test.

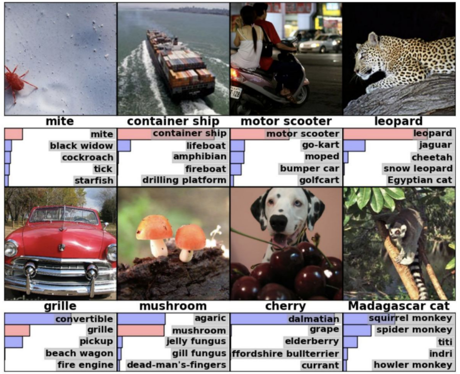

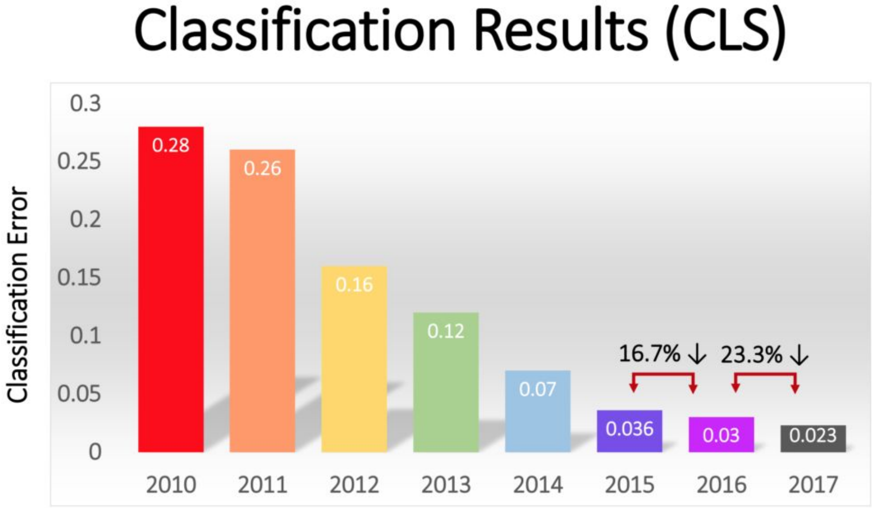

For you reading

The paper introduces 2 benchmarks: one on EV charging and one on the electricity markets. You only need to focus on the EV charging work. (Note: a more recent version of the paper has more RL tasks, if you want to consider it for your project)

A “pilot signal” sets the amount of electric current going into the car.

Exam recap

Grade Distribution