Lecture 5.5

Algorithms in the Real World

Dan Webber, PhD

Wait, who's *this* guy?

Hi! I’m Dan Webber (he/him)

- BA in Computer Science, Amherst College
  - Afterwards: a couple years as a software developer in fintech and e-commerce
- PhD in Philosophy, University of Pittsburgh
  - Moral theory
    - Basically: thinking really hard about *value*
- Now: Postdoc, EIS and HAI at Stanford
  - Embedding ethics into CS courses like this one!

Stanford University Human-Centered Artificial Intelligence

Stanford | McCoy Family Center for Ethics in Society

Lecture 5.5

Algorithms in the Real World

California High-Speed Rail

Goal: Renewable-powered, lightning-fast travel between SF and LA

But there are so many different routes we could take!

High-Speed Rail and Algorithms

- How could we use an algorithm to find the best route?
  - What an interesting question! Let’s come back to that in a few weeks…
- Okay… how could we use an algorithm to *select* the best route from a set of proposals?

Today

How can comparison-based algorithms help us with real world problems (and how can they *not* help us)?

1. Representing rail route selection as an algorithms problem

2. Comparison challenges: incommensurability

3. Measurement: problems with proxies, estimation

Suppose you have a big stack of route proposals…

Route Name	Time SF -> LA (min)
Coastal	180
I-5 Express	135
Eastern Valley	160
…	…
…	…

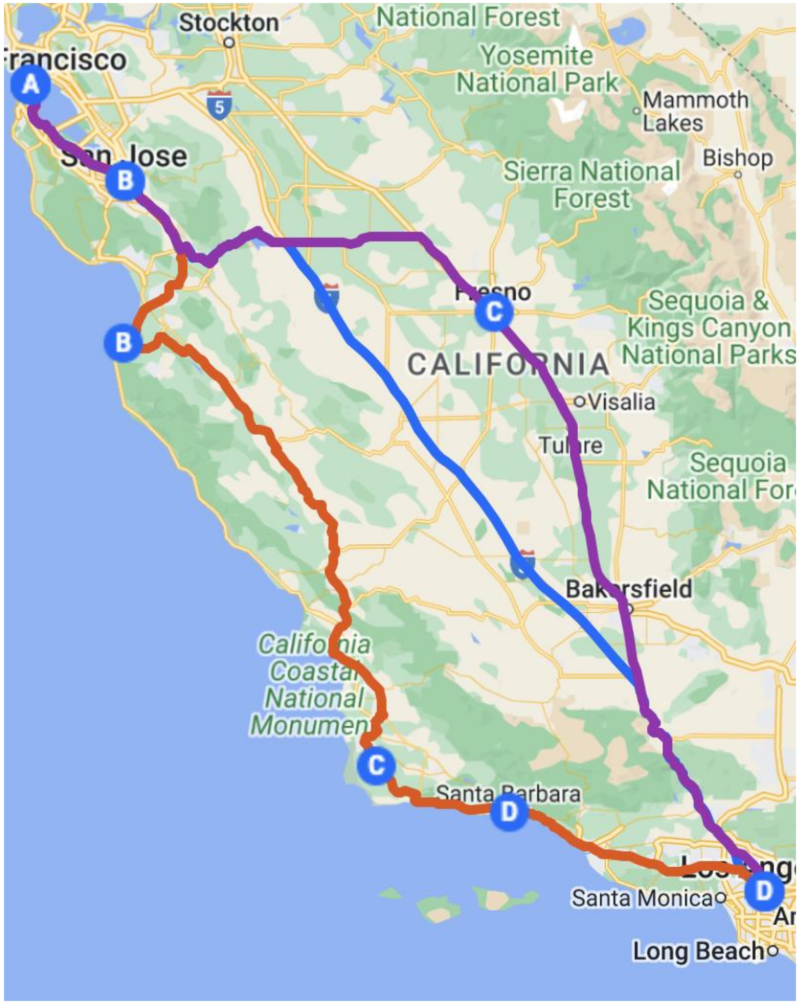

… and you want to consider them one by one, from most to least promising.

Route Name	Time SF -> LA (min)
Coastal	180
I-5 Express	135
Eastern Valley	160
…	…
…	…

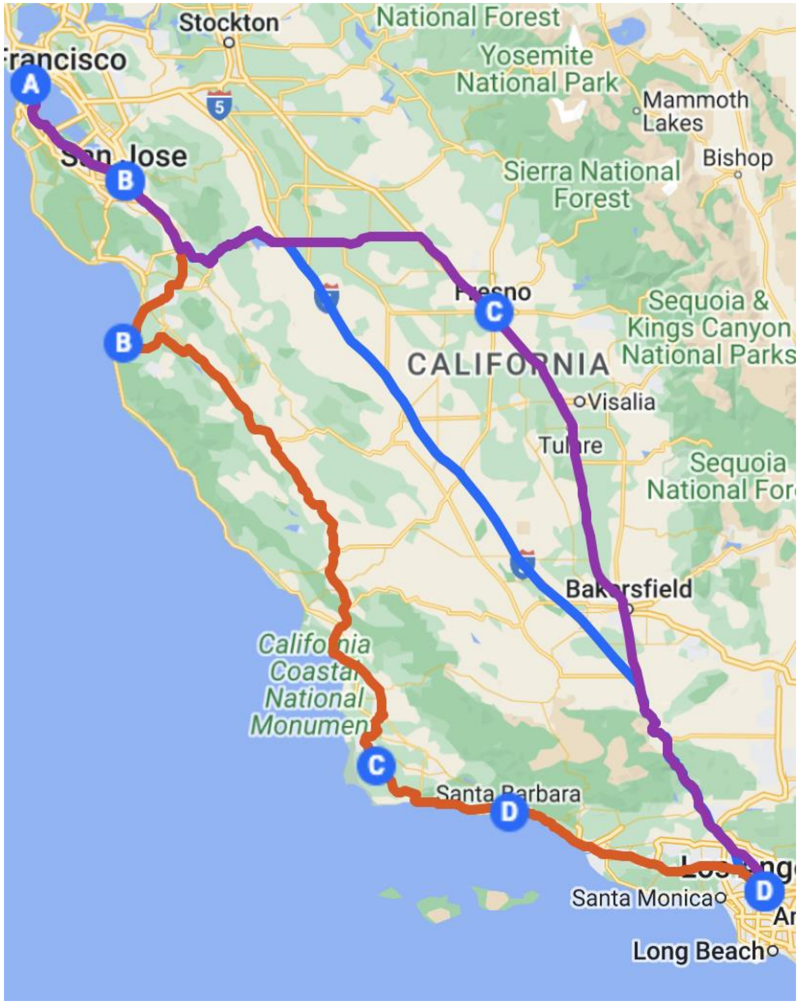

How could an algorithm help you here?

Route Name	**Time SF -> LA (min)**
Coastal	180
I-5 Express	135
Eastern Valley	160
…	…
…	…

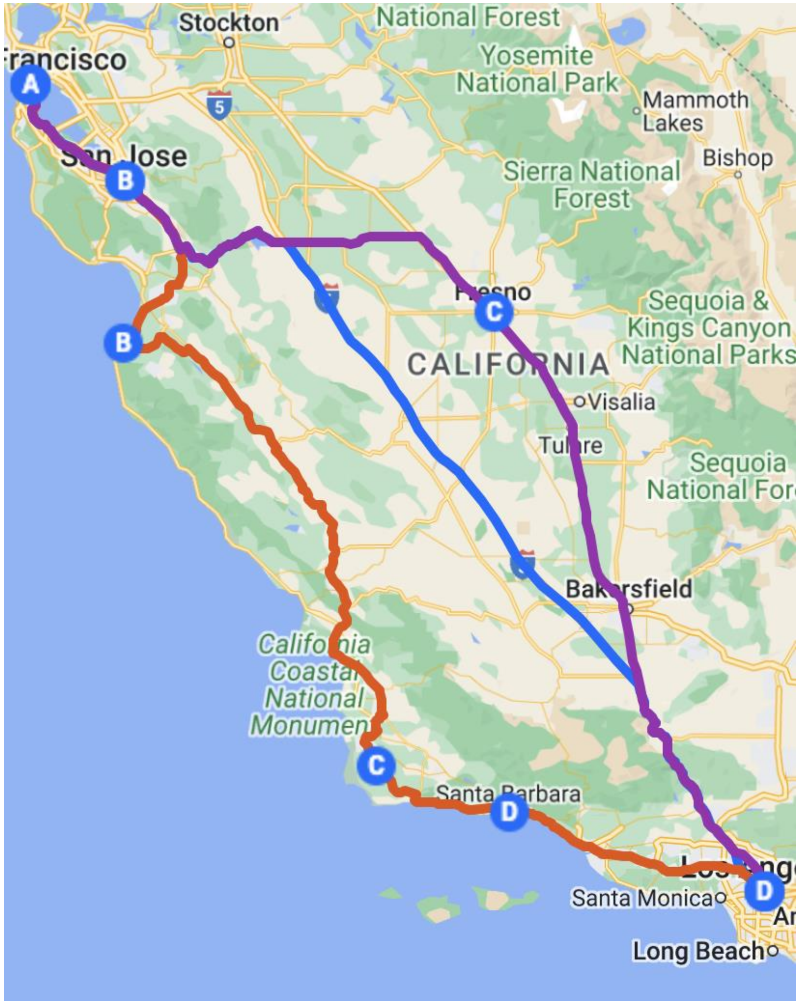

Sort the list (by time)!

Route Name	Time SF -> LA (min)
I-5 Express	135
…	…
…	…
Eastern Valley	160
…	…

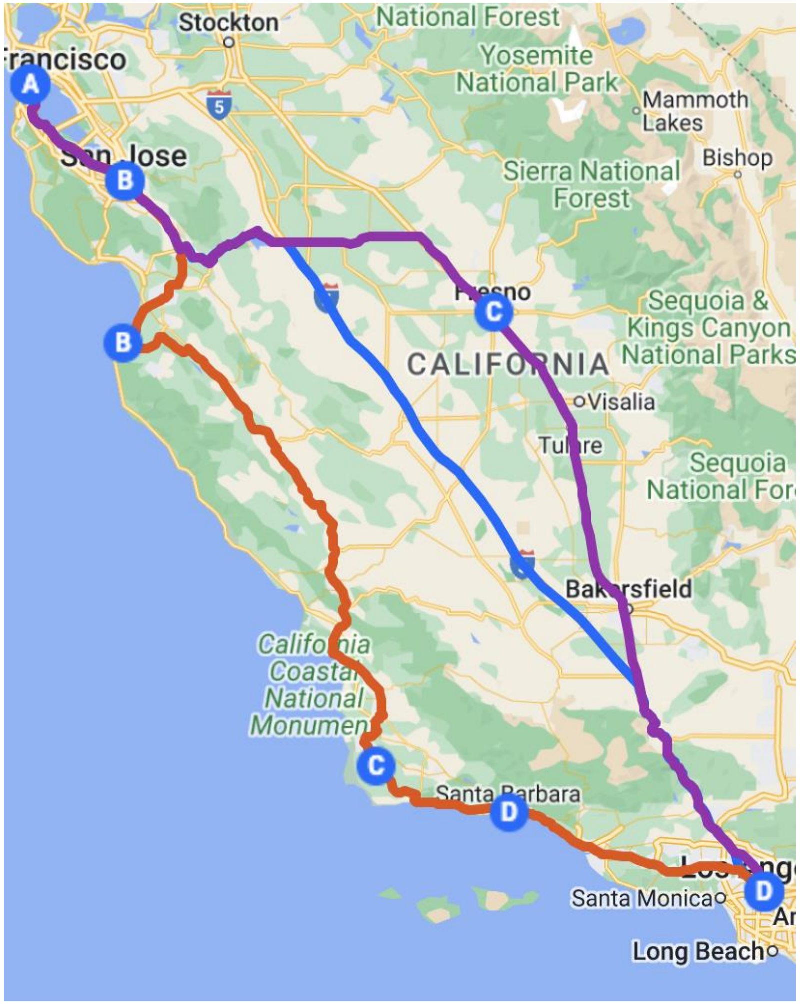

That's not the route they're actually planning to build, though…

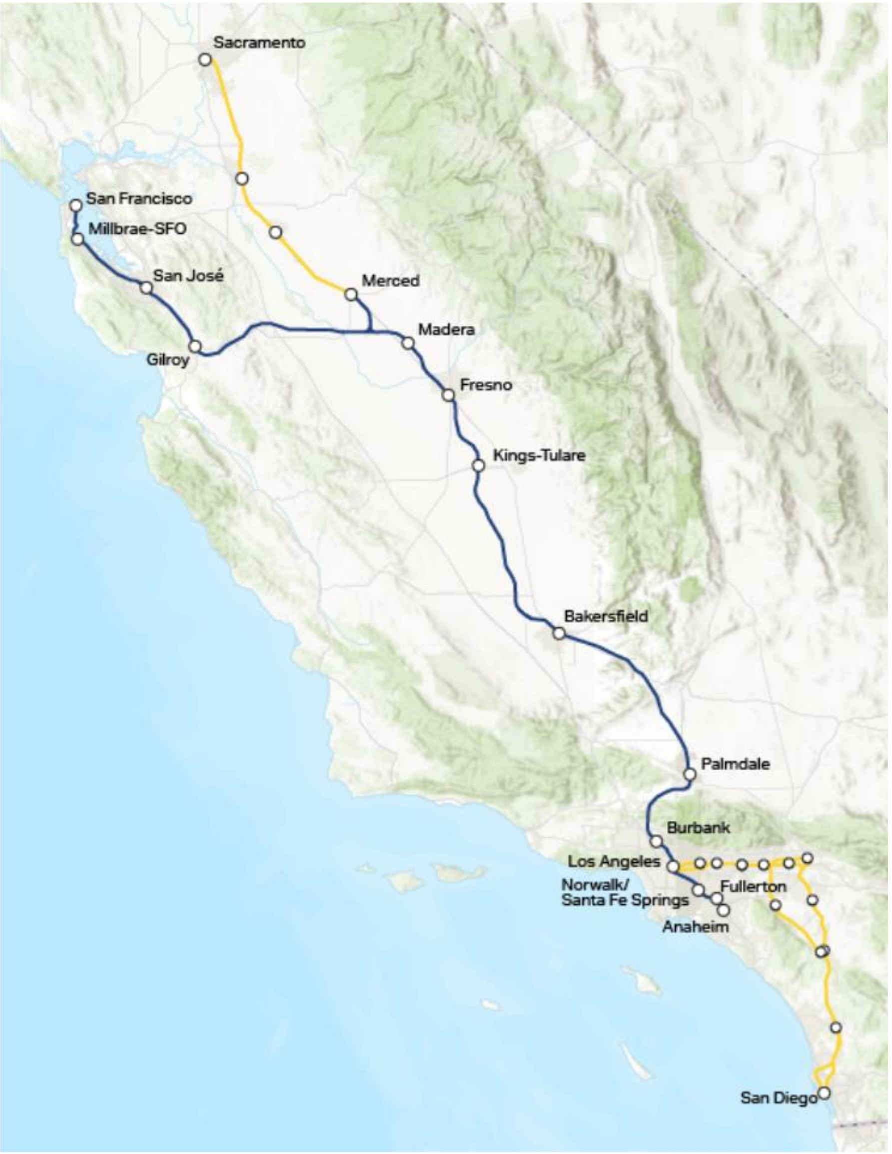

What else might matter for route selection besides speed?

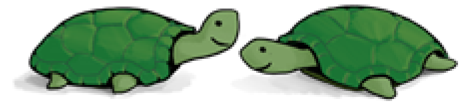

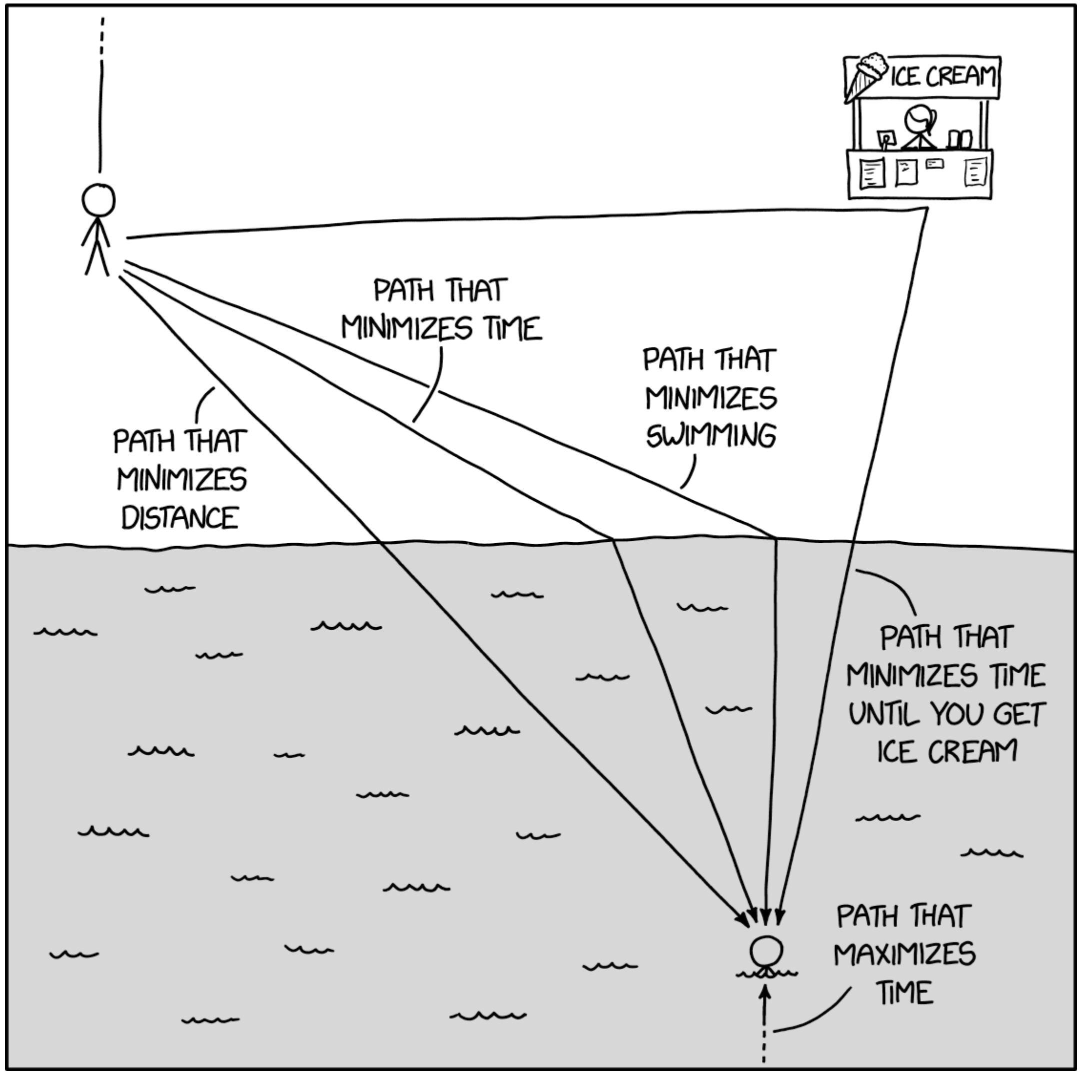

Example: emissions

- High-speed rail will be entirely powered by renewable-generated electricity
- It will reduce carbon emissions primarily by taking cars off the road, so emissions reduction depends on *ridership*
- The state estimates that the eastern route will reduce emissions by ~2 Mmt per year
- Would you expect the “direct” I-5 route to reduce emissions by more or less than this?

Taking emissions into account

- Suppose we also have emissions reduction estimates for each proposal. How should we sort the proposals now?

Route Name	**Time SF -> LA (min)**	**Emissions reduction (Mmt)**
Coastal	180	2
I-5 Express	135	1.5
Eastern Valley	160	2
…	…	…
…	…	…

One way: a weight function

- Rather than sorting by time, sort by a function of both time and emissions reduction
- We want to rank the proposals from best to worst in terms of some **f(t, e)**, which reduces time and emissions to a single value

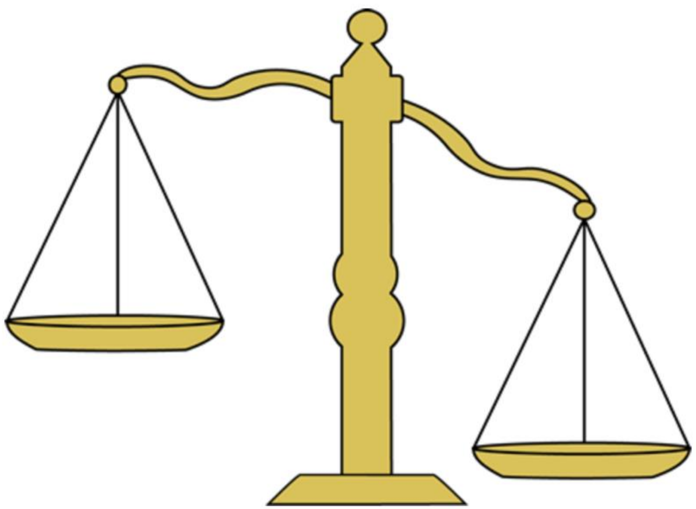

Sorting with a weight function

Route Name	**Time SF -> LA (min)**	**Emissions reduction (Mmt)**	**f(t, e)**
Coastal	180	2	
I-5 Express	135	1.5	
Eastern Valley	160	2	
…	…	…	
…	…	…	

Sorting with a weight function

- Problem: which function should we use for **f(t, e)**?

Route Name	**Time SF -> LA (min)**	**Emissions reduction (Mmt)**	**f(t, e)**
Coastal	180	2	???
I-5 Express	135	1.5	???
Eastern Valley	160	2	???
…	…	…	…
…	…	…	…

Comparison is key, but can be hard!

Is there a **common measure**?

numbers have them

- 1 < 2; 1==1; 1 > 0

some factors don’t

- ⏳ v.s. ⛽

Are our **proxy measures** appropriate?

What are the pitfalls of trying to determine one value by measuring a different one?

Incommensurability

The First Barrier to Comparison

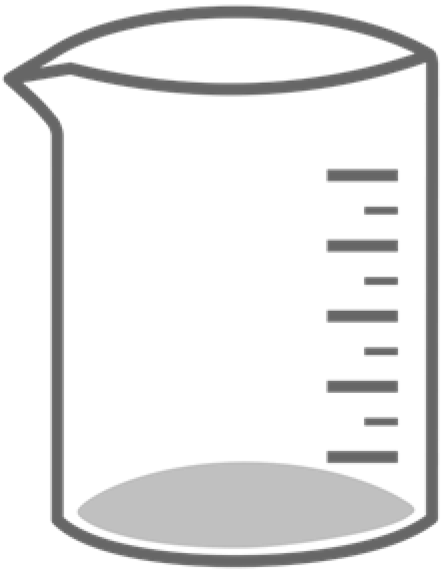

Incommensurability

Two things are *incommensurable* if they lack a *common measure* of value—that is, if there is no single measure that can be used to evaluate both.

Incommensurability makes it difficult to establish measure relationships such as “more than” or “less than, “better than” or “worse than” **by how much**.

For example, what is the relationship between milliliters (ml) and centimeters (cm)? How would you rank different amounts of water and physical lengths in the same list? You wouldn’t!

Some things are *commensurable*

- Apple stock and Microsoft stock are commensurable in terms of *dollars*: we can measure and compare them by dollar value

- Touchdowns, field goals, PATs, and safeties are commensurable in terms of *points*: we can measure and compare them by how many points they're worth

When two things are commensurable, they can easily be reduced to a single value

- f(AAPL, MSFT) = (223.66*AAPL) + (446.71*MSFT)
  - This function reduces Apple and Microsoft stock to a single dollar value

- f(TD, FG, S, EP, TPC) = (6*TD) + (3*FG) + (2*S) + (1*EP) + (2*TPC)
  - This function reduces touchdowns, field goals, safeties, extra points, and two-point conversions to a single point value

Not so with incommensurable values

What's the function for reducing career advancement and friendship to a single number?

Or oil and biodiversity?

Or… time and carbon emissions?

We can evaluate route proposals on each scale separately, but what if we need to consider both?

Incommensurable != Incomparable

- We might be able to tell that X hours saved is better or worse than Y tons of emissions, even if we can't say *how much* better or worse.
- (Are there *any* two things that are truly *incomparable*?)

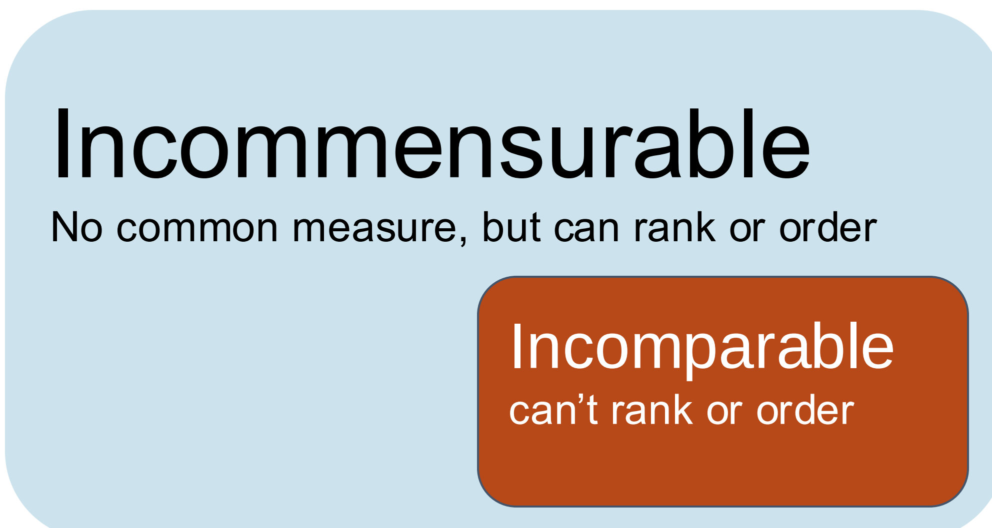

What are some other values that are commensurable? Incommensurable?

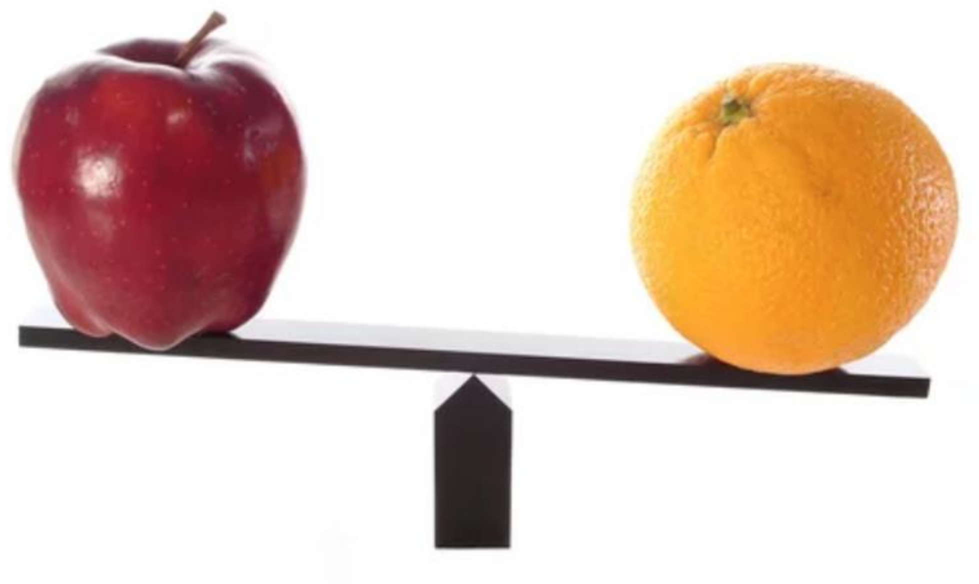

Upshot for Algorithms

Many algorithms work best on values that have a built-in standard of comparison, like numbers.

In real-world situations, some problems you might try to formalize and solve algorithmically will include incommensurable values with no obvious standard of comparison.

There might not be a single “optimal” choice, and results will vary depending on how you choose to weigh different values. Technical considerations alone won’t determine the correct weight function.

Back to the problem at hand…

Route Name	Time SF -> LA (min)	Emissions reduction (Mmt)	f(t, e)
Coastal	180	2	
I-5 Express	135	1.5	
Eastern Valley	160	2	
…	…	…	
…	…	…	

A silly example

Let **f(t, e)** = t – e

(Why minus?)

(Why do I call this example silly?)

Route Name	**Time SF -> LA (min)**	**Emissions reduction (Mmt)**	**f(t, e)**
Coastal	180	2	178
I-5 Express	135	1.5	133.5
Eastern Valley	160	2	158
…	…	…	…
…	…	…	…

A silly example: sorted

Let **f(t, e)** = t – e

- Essentially, this function says that 1 minute of time per rider is worth 1 Mmt of emissions

Route Name	**Time SF -> LA (min)**	**Emissions reduction (Mmt)**	**f(t, e)**
I-5 Express	135	1.5	133.5
…	…	…	…
Eastern Valley	160	2	158
…	…	…	…
…	…	…	…

What if we change the function? Or even just the units?

Let **f(t, e)** = t – e, but where t is measured in **hours**

- Without the unit change, **f(t, e)** = t/60 – e
- 1 *hour* per rider is worth 1 Mmt of emissions

Route Name	Time SF -> LA (hour)	Emissions reduction (Mmt)	f(t, e)
Coastal	3	2	1
I-5 Express	2.25	1.5	.75
Eastern Valley	2.66	2	.66
…	…	…	…
…	…	…	…

What if we change the function? Or even just the units?

Let **f(t, e)** = t – e, but where t is measured in **hours**

- Without the unit change, **f(t, e)** = t/60 – e
- 1 *hour* per rider is worth 1 Mmt of emissions

Route Name	Time SF -> LA (hour)	Emissions reduction (Mmt)	f(t, e)
Eastern Valley	2.66	2	.66
…	…	…	…
I-5 Express	2.25	1.5	.75
…	…	…	…
…	…	…	…

What about values that can't even be reduced to a number *on their own*?

The Amah Mutsun Tribal Band of Costanoan/Ohlone Indians filed a comment on California's high-speed rail proposal arguing that the proposed route would "significantly impact areas of profound cultural, spiritual, ceremonial, and historical significance to the Tribe," which would "contribute to the historic trauma the Amah Mutsun have suffered as a result of three consecutive periods of brutal colonization."

This is obviously an important consideration. But how could we factor it into our ranking algorithm?

Measurement

The Second Barrier to Comparison

Measurements vs. The World

Often the value that is feasible to measure is an imperfect proxy for the thing we would actually like to measure.

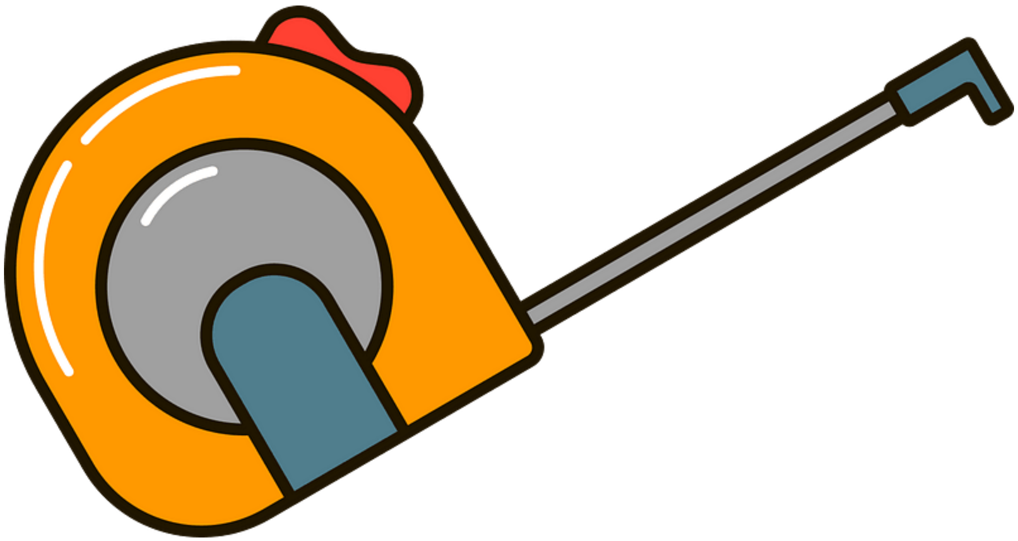

GDP (Gross Domestic Product)

- GDP is (supposed to be) the total value of all goods and services produced within a country.
- GDP is ”easy” (challenging but tractable) to measure.
- It is meant to measure economic health, but does it?
- As a formalization that allows measurement, does it capture the aspect of the world we want to formalize?

GDP as a Measure of Economic Health

As a measure, it has (at least) three kinds of problems:

- Idealization: Imperfect counting of the thing it measures (what production & work is it missing?)
- Abstraction: counting all 10 million dollar expenditures as equally contributing to economic health (what kinds of expenditures might be more or less valuable?)
- Imperfect tracking of the thing we really wanted to measure, which might have other values not captured

Measuring ridership (in advance)

- Remember: emissions reduction from high-speed rail depends on *ridership.*
- Ridership is easy to measure after the fact. But we need to know how many people would ride each route *before* we decide whether to build it.
- What *already existing* data could we use to estimate ridership? And how might that data be an imperfect proxy for rail ridership? 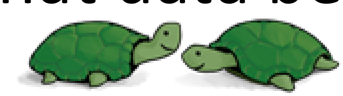

Inclusion & Exclusion from Data

Risks of Inclusion in Data

- Violation of privacy
- Data scraped without consent
- "Consent" not informed
- Risks to Fourth Amendment rights (probable cause)
- Risks to other civil liberties
- Misinterpretation of data

Exclusion from Algorithmic Data

How does the exclusion of some people’s life habits, patterns, preferences, beliefs, and desires from data affect our later use of that data in our algorithms?

Risks of *Exclusion From* Data

- Data doesn't describe you
- Your needs are not accounted for in political decision-making
- Exclusion from resources
- State might take action to make you more "legible"

Incommensurability, Measurement, and Inclusion

Think about the problems you have solved in this class or could imagine solving with your current skills.

- Can you think of any that involve incommensurable values?
- What questions might you ask about the measurements that went into producing the numbers on which you relied?
- What kinds of people might be included or excluded from the data because of how it was gathered?

Today

How can comparison-based algorithms help us with real world problems (and how can they *not* help us)?

1. Representing rail route selection as an algorithms problem ✔
2. Comparison challenges: incommensurability ✔
3. Measurement: problems with proxies, estimation ✔

Want to talk more about ethics?

(Not right now!)

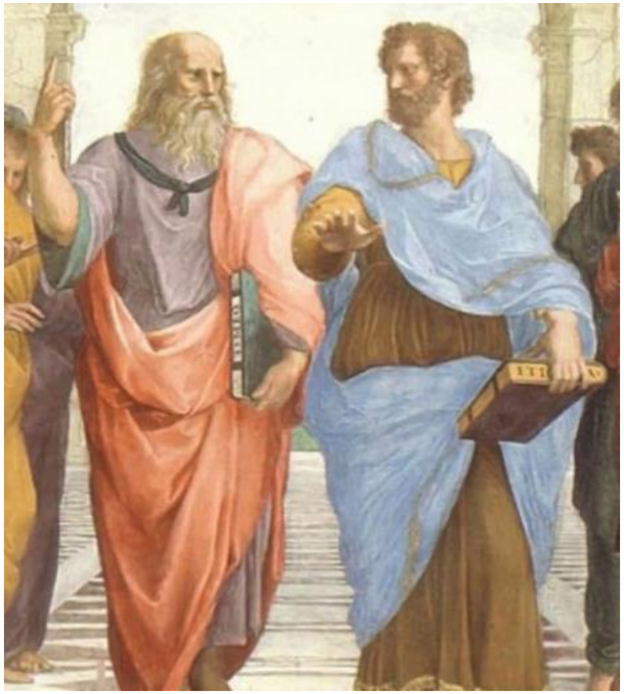

Dan Webber | webberdf@stanford.edu

Email to set up a meeting!