Hi, my name is Paul, and I'm a nerd. Maybe you can tell, from the title of the talk.
I'm an open source programmer and a technology watcher and definitely a big big nerd.
So, yes, the title of this talk…

…**is** taken from Star Trek 6, the Undiscovered Country
(highly recommended, by the way).
And in particular, taken from the scene, early in the movie,

where the Klingon Chancellor, Gorkon,
proposes a toast to “the undiscovered country”,
and everyone looks at him blankly, so he clarifies,
<x> “to the future”.
So, yes, I am a nerd, so much so that I’m going to be coming **back** to this scene…
but first, more science fiction…

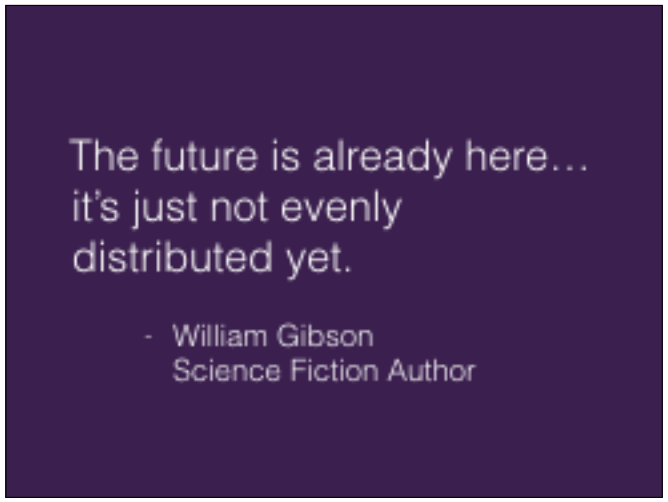

> "The future is already here... it's just not evenly distributed yet." William Gibson
And because the leading edge of technology appears in different fields,
at different times, we can learn a lot about where we are headed,
as a profession, just by looking around.

Because, as much as any author, banging away at a word processor,
or office worker, shuffling records in a database,
or musician, manipulating tracks on a laptop,
we **cartographers**, we **geographers**, we experts in the **WHERE** of things,
now work in a technologically mediated vocation.

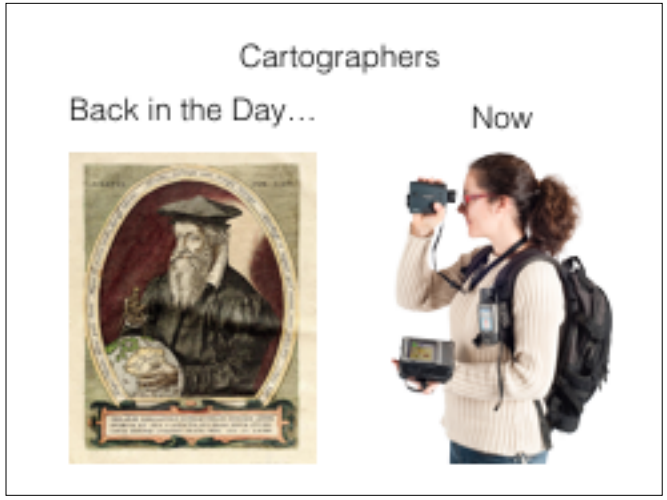

The older of us have experienced the transition,
from physical to digital, and perhaps sense that,
for all the change we've seen, over the last 30 years,
in the art of knowing and showing where things are,
we're just getting started…

For the most part, we have not been the deliberate architects of our future,
we've been swept along on the same current as other professions.
But the big changes in our field have been driven by the **same trends** that have changed the rest of society:
the rise of general purpose computing,
the rise of the power of the machines available to us,
and then, more recently, the inter-connection of those machines, into a global network.

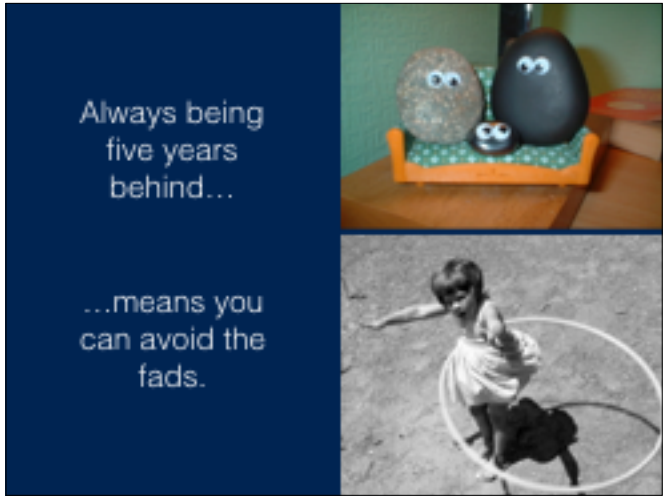

For better or worse,
the geographic or GIS world does not adopt new technology
as fast as the corporate world
or the world of consumer electronics.
On the downside, that means we're frequently 5 years behind.
On the upside that means it's not too hard to predict what will be changing our field in five years time.

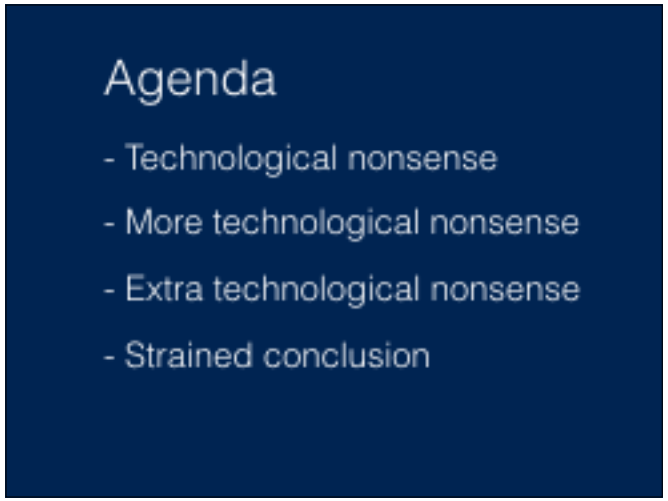

To a large extent what I have for you today is a laundry list of interesting technology trends that are going to effect our field over the next five years.

But if the past 30 years of computer-driven upheaval have taught us anything, it's that technological change doesn't occur in a vacuum. Each change has knock-on societal effects, and the effects are radically different as the technology moves down the availability curve from niche luxury to ubiquitous commodity.

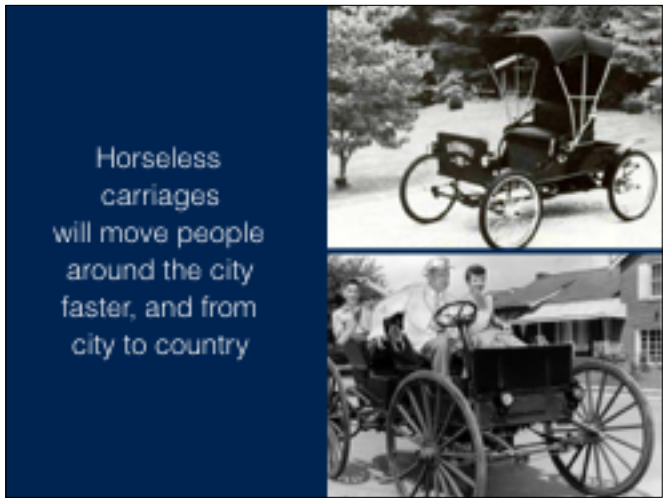

Looking at the expensive and unreliable horseless carriages in the late 1800s, futurists would have had to stretch to imagine

… highway interchanges, or strip malls,
developments that only make sense in a world of ubiquitous auto ownership.

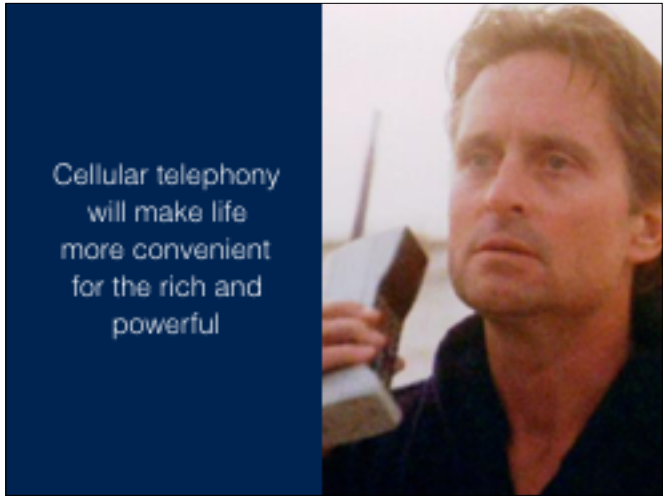

Looking at the cell-phone that was used as a status symbol for Gordon Gecko in the 1980's movie "Wall Street", it would have been hard to predict the social effects of cellular technology

when placed in the hands of third world farmers or slum dwellers…

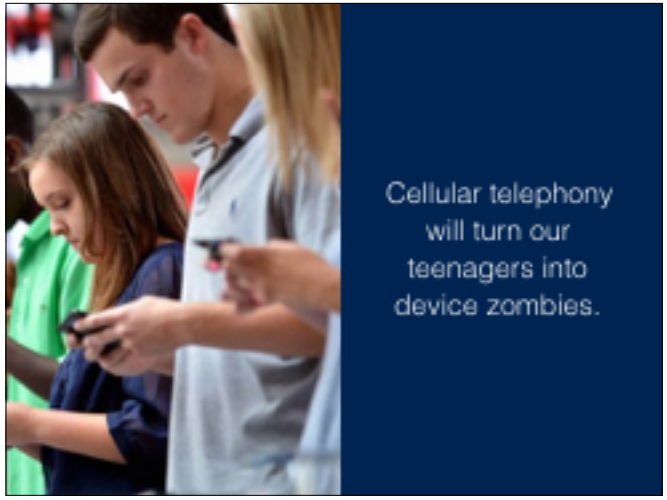

or first world children.
When **access** to technology moves from elites
to the mainstream, the effects are unpredictable.

Now, Twitter is generally a cesspool of nonsense,
but the 140 character limit does have the effect of generating nice terse quotes,
like this one from Hillary Mason:

> “First it becomes possible, then it becomes so cheap that anyone can do it... and that's where it gets interesting.” Hillary Mason https://twitter.com/hmason/status/772075074065166345

"Interesting". It's not **quite** a value neutral word. It's a positive thing. That's interesting! You really have to work to bend it to the negative. That's... interesting.
The direction from which you approach a subject can have a big effect on how you perceive it. I'm a techno-**optimist**, because I love technology, even though it scares me a little, around the edges.

So, here are the broad trend lines:
* technology continues to improve, but **better technology alone** doesn't change the world
* the really **interesting** things happen when technology **cost** falls far enough to enable adoption at scale
That is where it gets "interesting"".

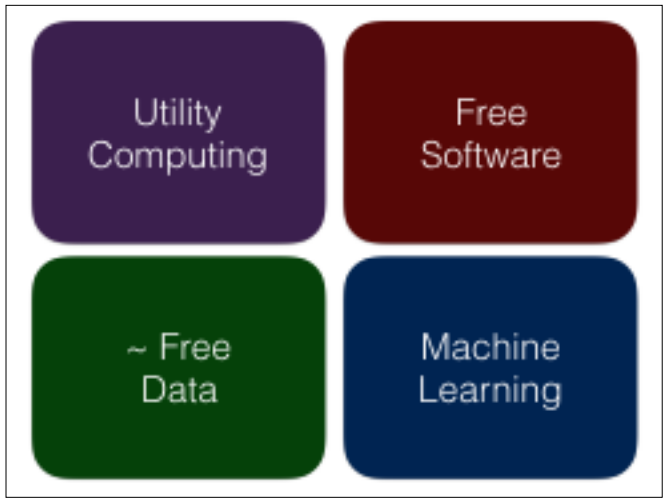

I'm going to talk about four trends, that are not only "interesting" from a capability point of view, but also from a transformational point of view.
- <x> Computing is a utility
- <x> Software is free
- <x> Data is almost free
- <x> Machine learning is available to all

None of these trends are about changes in **technology,** they've already happened, they are all about changes in **availability**.
That means available to **you**, which means transformation in your workplaces and your organizations.

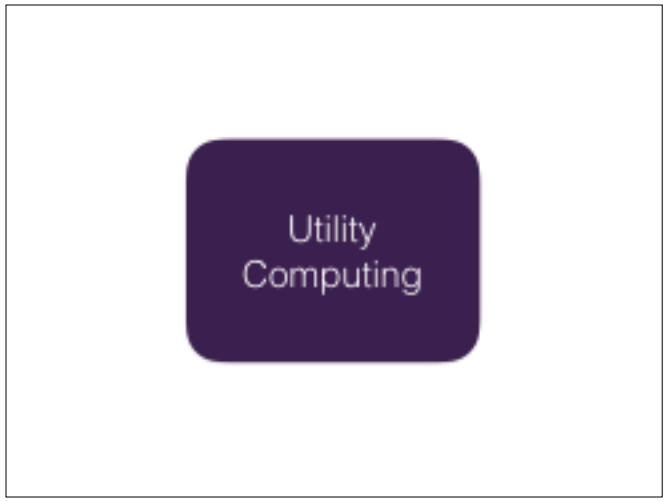

So first, utility computing.
As a technology story, utility computing or "cloud" computing is very much old news.
And yet, I'll take a risk here, raise your hand if your organization **still runs its own servers**.
This is our industry, always a little behind the crest of the wave.

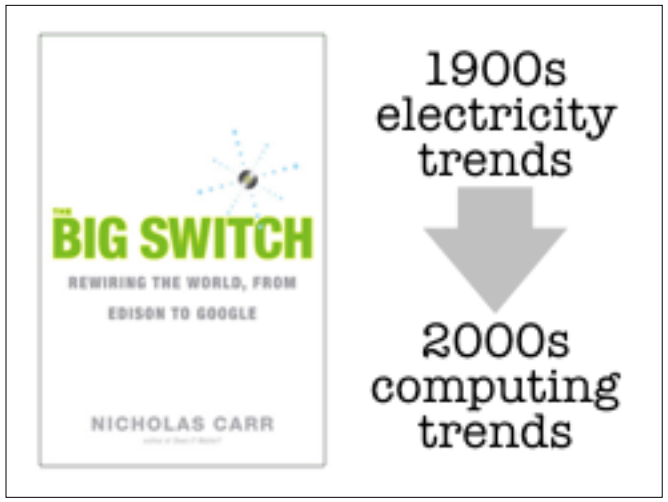

In 2013, Nicolas Carr wrote a book, "The Big Switch",
that compared the move to cloud computing architectures
to the rise of utility grade electricity.

In the late 1800s, electric lights were cool, but the big win for electricity was in changing the nature of factories.
Before electricity, machines (like this belt-driven lathe) had to be harnessed to central drive shafts for power, with complex systems of gears and belts.
The belt-and-drive system meant that factories had to be located very close a natural source of power, like a water wheel, or had to run their own steam engine.
The belts and gears were dangerous, people lost limbs.
Changing the configuration of an existing factory was nearly impossible, because you didn't just have to rework the location of machines, you had to rework the whole mechanical power transfer system.

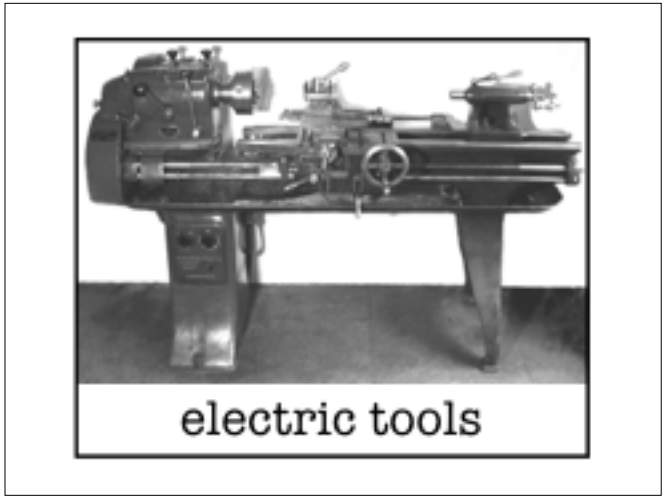

Electric powered tools, like this lathe, had their own **small electric engine** to spin the material.
No more belts and gears, no more belt-and-gear accidents.
Any tool can be located anywhere, just run a new wire.
Also, the factory can be located anywhere you can run an electric steam generator.

electric transport

Electricity also let new transportation companies run electric streetcars: faster, cheaper and less smelly than the horse-drawn omni-busses they replaced.

---

pearl street generator

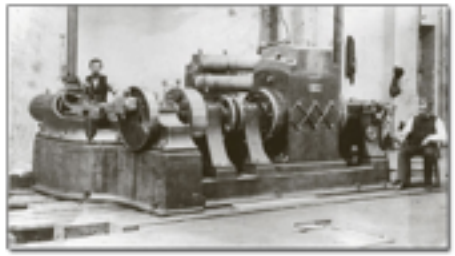

1882

Thomas Edison set up his first large scale demonstration generating station on Pearl Street in New York in 1882.
General Electric began selling generating machinery to industry shortly thereafter.

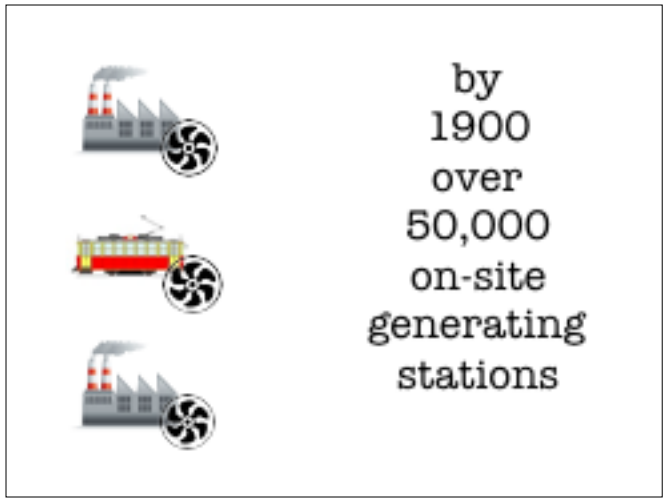

And by 1900, that's just 18 years later, the US Census recorded over **50,000** site-specific generating stations in the USA.
An explosion of new electrified technology, into every workplace in the country.
Does anything sound familiar?
Think: word-processing. Think, PCs: 1985 to 1995.

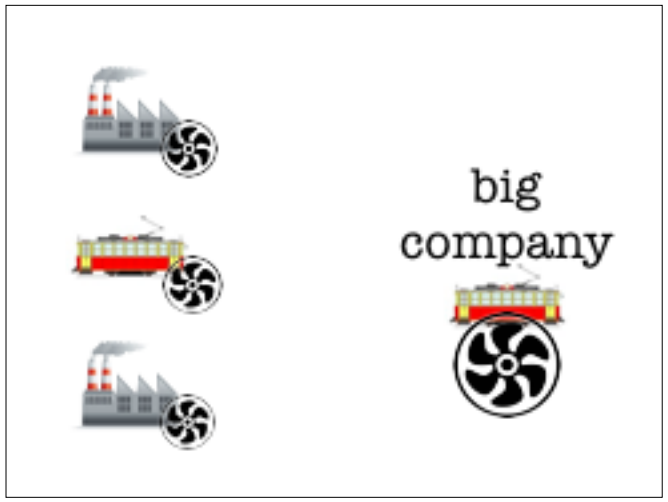

The folks with the biggest generating stations (often streetcar companies) and the best operating records began selling excess power to other users on the side.

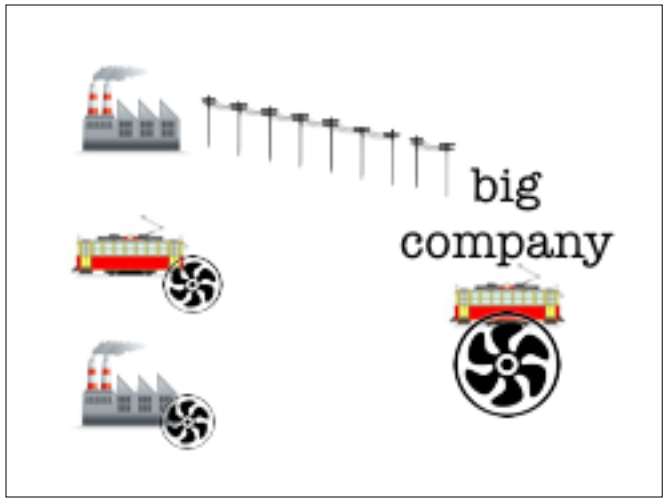

These users got the service reliability of a large company without having to make a huge infrastructure investment.
More sophisticated grid technology meant utilities could provide even greater reliability to their customers using multiple generating stations.

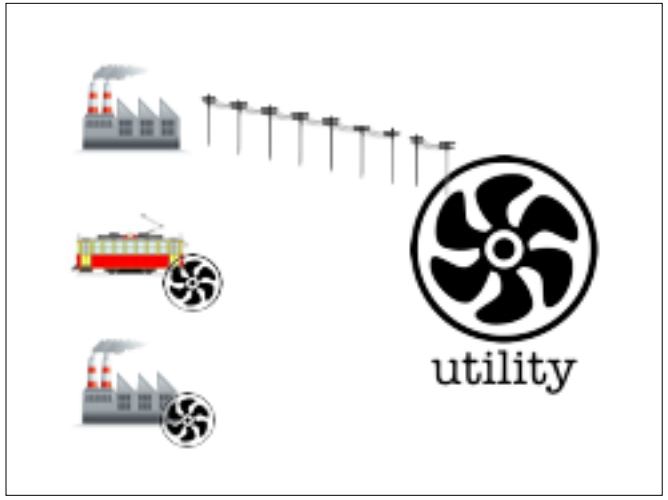

The former streetcar companies moved directly into generating as their main business and became utility power companies.

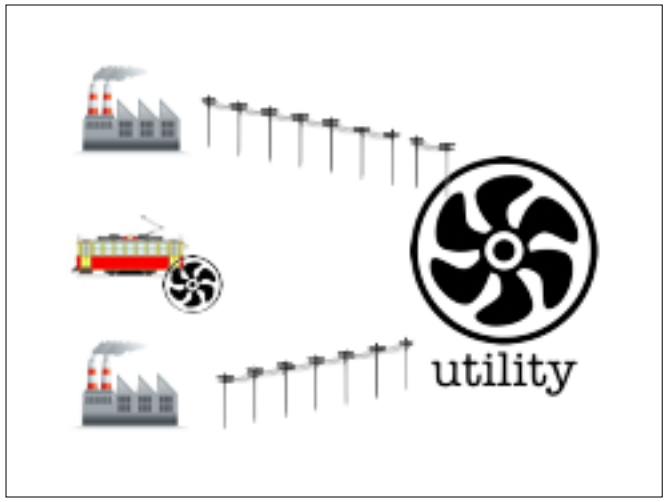

And the factories began shutting down their old generators, taking power directly from the grid.
Though some stubborn ones held out longer, preferring to control their own generators,…

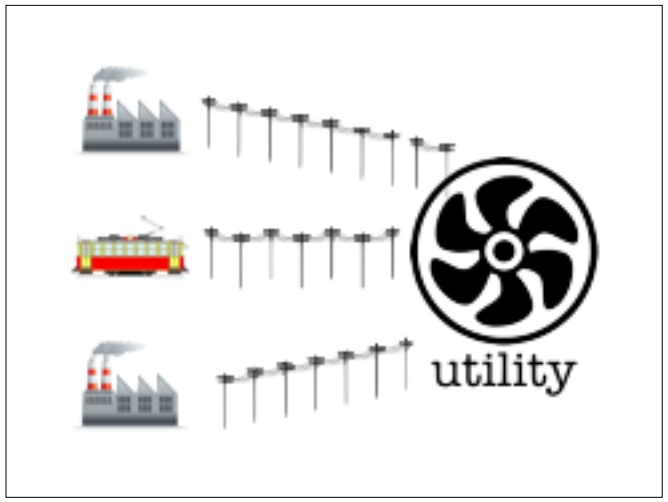

…they all, eventually, gave in.
We're seeing the same transition in IT.

Microsoft was a business software company, now they sell utility computing too.

Google was a search/advertising company, now they sell utility computing too.

Amazon was a web retailer, now over **50% of their profit** is coming from selling utility computing.
And those are just the companies that are **succeeding** at utility computing, there's dozens more like IBM and Oracle that are desperately trying to get into the business.
As a consumer, transitioning to utility computing is not as abrupt
as flipping the switch on grid electricity,
because of the nuances of information technology,
but the transition is just as certain,
because it is being driven by powerful economic forces.

Can you provide a data centre as cheap and reliable and redundant as the Amazon AWS cloud?
<X> No, Of course you cannot. So eventually, you'll switch.

Can you provide e-mail services as cheap and reliable and redundant as Google or Microsoft can?
<X> No, Of course you cannot. So eventually, you'll switch.
This is a guarantee.
All computing, eventually, will be utility computing.

There's no doubt that established organizations don't change very fast.
So don't judge trends or change by looking at **OLD** organizations.
Look at **NEW** organizations.

what would a **brand new company** do?

Does a NEW company set up their own data center?
Does a NEW company set up their own email server?
Do they run their own accounting software, timesheet tracking, calendar system, CRM database, hiring pipeline?

OldCo	NEW Co
in house data center	Amazon Web Services
in house e-mail	Google Apps
accounting system	QuickBooks Online
custom timesheet	Harvest
in house calendar	Google Apps
in house CRM	Salesforce.com
recruiting system	Resumator

I note all those functional categories,
because back when I set up my first business, in 1997,
I installed and managed *all* those information management systems on-site.
At the two relatively brand new companies I have worked for since 2009 (Boundless and Carto),
all those functions and more are consumed as services over the web.

this is about
access

Utility computing is transformative because it changes the terms of **ACCESS** to IT.

People have been using super-computers for years,
both expensive ones like the lovely old Cray computers and cheap ones,

like the Linux beowulf clusters of the late 90s.
But cheap or expensive, they had one thing in common:
you had to buy them, before you could use them,
and you could only use what you had bought.

One of the surprising things you can do with utility computing is build an on-demand supercomputer.
Requisition 100 servers, install your computation on them, run it, and then turn them off and hand them back.

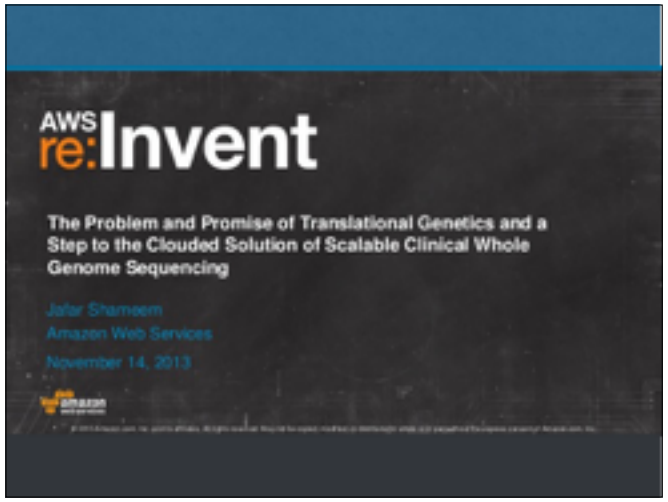

Geneticists are using the Amazon cloud in this way to run computations against huge genomics databases.
Rather than running 10 nodes of computation for a 100 hours in their lab, they run 1000 nodes of computation for 1 hour in the Amazon cloud.

Similarly, the Obama campaign in 2012 spun up nightly analysis jobs to crunch their huge collection of voter data and spit out new estimates of likely support for each voter.
These kind of workloads would never have been **considered** before utility computing,
because of the cost of **buying** and **deploying** so many machines,
Now they are routine.

It's not the capability.
There have **always** been supercomputers.
It's the price: "when it becomes so cheap that **anyone** can do it, **that's** where it gets interesting".

So, what should you do,
geographers and cartographers and GIS'ers,
to understand and ride this transformation?

Our industry is already a good way into the transition.
Esri recognizes the change to utility computing,
which is why they are promoting ArcGIS Online so heavily.

Start-ups like Mapbox and Carto are also offering new looks at the old business of spatial. MapBox from a location services point of view: maps, geocodes, routes.
Carto from a location intelligence point of view: analysis, data augmentation, visualization.

But don't just look within our industry. That's why we end up 5 years behind the curve.
Also look directly at what Amazon and Google and Microsoft have to offer.
Look at what services you can move off-site.
There is no big capital investment,
so there is no excuse not to experiment with **every** option, at least a little.

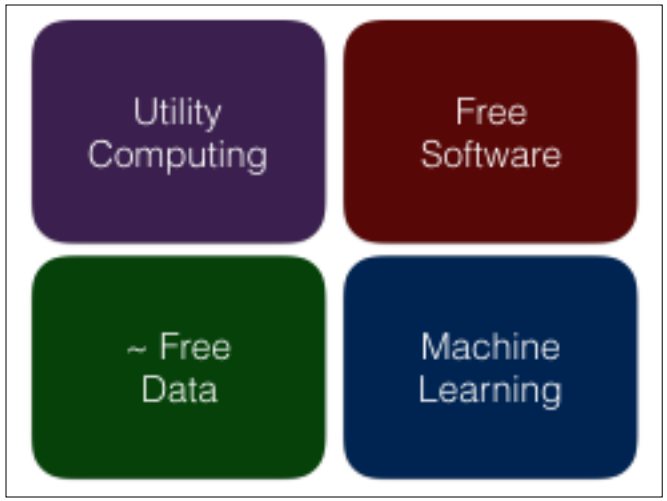

So,
- Computing is a utility
- Software is free
- Data is almost free
- Machine learning is available to all

Let’s look at free software, open source software.
Like utility computing, free software is very much an **old** story.
It's been transforming the technology world for a generation,
and has hopefully been on your radar as spatial technology experts for at least the last 5 or 10 years.
Before getting to my larger point, I'd like to point out a few pieces of free and open source software, for illustrative purposes:

QGIS is a free and open source desktop GIS software.
It's functionally similar to ArcGIS, and largely comparable in terms of features offered:
viewing, editing, vector data, raster data, analysis, and cartography are all supported.

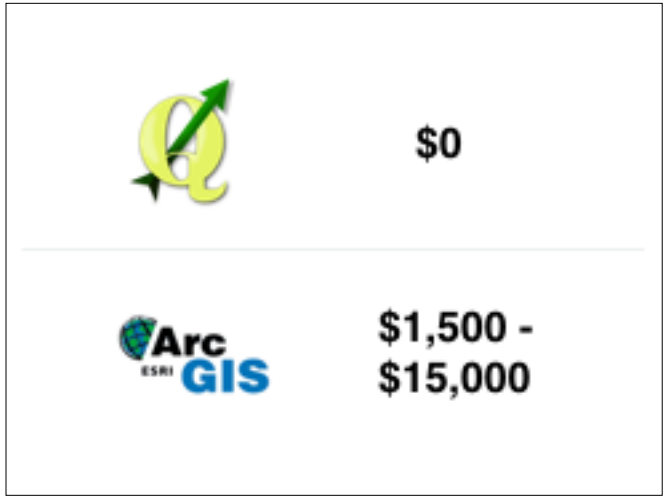

The price of ArcGIS ranges from
<x> $1500 to $15000, depending on what you need it to do.
The price of QGIS is
<x> $0, no matter what you want it to do.

this is about
access

This isn't a moral argument,
there's nothing **wrong** with Esri charging for their software,
but price limits utility. This is about **access**.

Where can you put a $0 fully featured GIS that you cannot put a $1500 one?
With an NGO in Subsaharan Africa,

or in the hands of 1000 school children,
or on the desktop of every member of your organization
just in case they need to deal with a shape file.

---

this is about
access

These are options that transform who does geographical analysis,
and the purposes for which it can be used,
and they are options that only open up when the price point falls to $0.

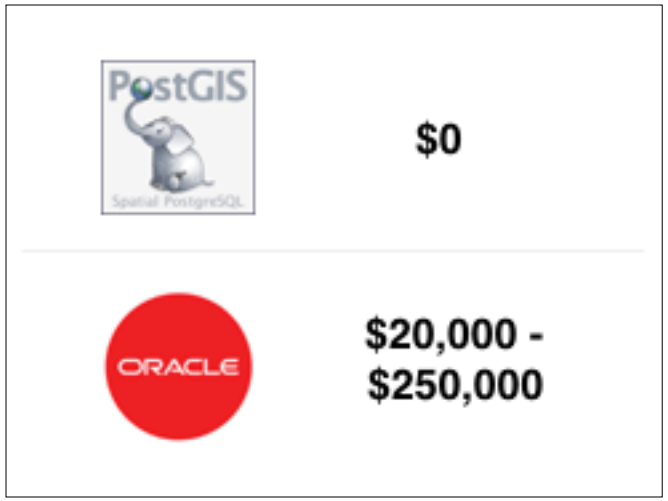

PostGIS is a spatial database engine.
It's functionally similar to Oracle Spatial or Microsoft SQL Server.
The price range of Oracle ranges from
<x> $20,000 to $250,000 depending on how many people you expect to serve.
The price of PostGIS is <x> $0, no matter how many people you expect to serve.

Mapnik and MapServer and GeoServer are all
map rendering engines and data transformation engines,
functionally similar to ArcGIS for Server.
The price of ArcGIS for Server ranges from
<x> $5000 to $20000.
The price of the free engines is... <x> well, free.

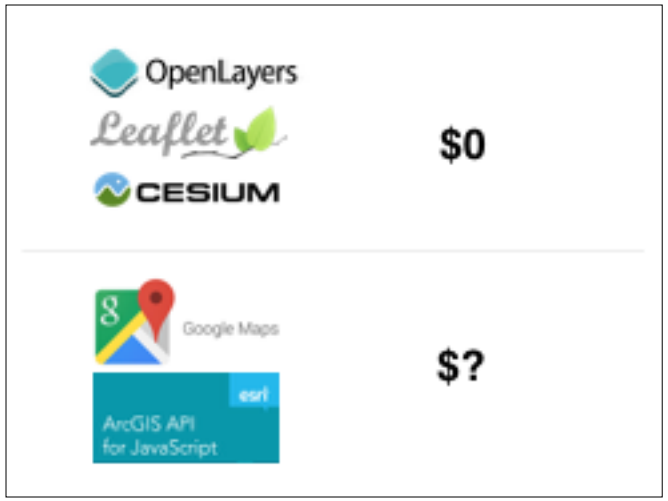

Leaflet and OpenLayers are Javascript web map APIs.
They are functionally similar to Google Maps or Esri's Javascript APIs.
Cesium is a 3D web map API.
It's functionally similar to Google Earth, but it runs entirely in a browser using WebGL.
The pricing of Google Maps or Esri Javascript APIs <x> varies depending on what you want to do. The price of Leaflet/OpenLayers/Cesium is <x> always $0.
There's a lot of free software, I could do this all day, and I still have some more software to tell you about. But I'd like to take a break and loop back to utility computing.

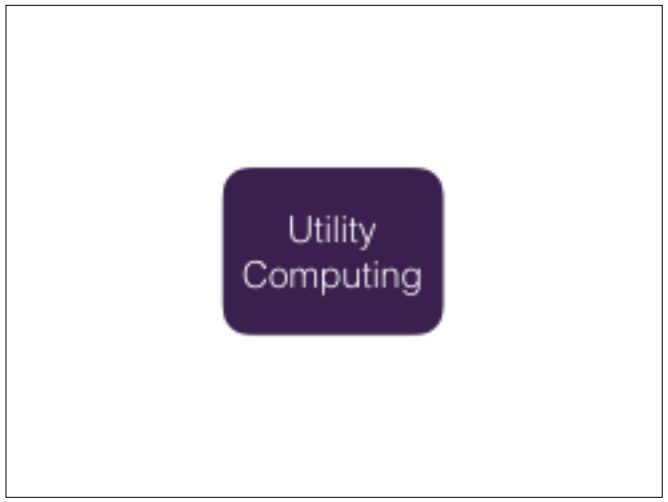

What do you get when you join utility computing with free software?

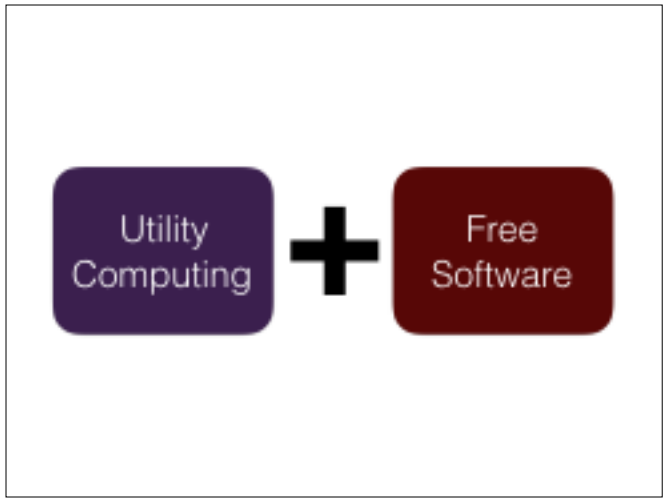

You get completely new businesses that can operate at scale.

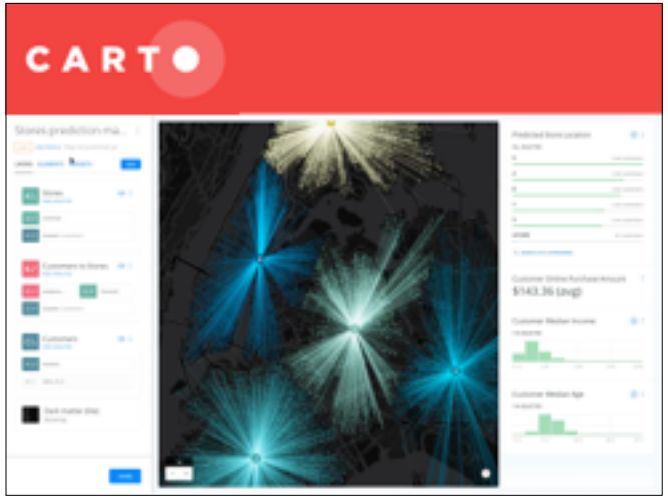

I work for Carto, so excuse me if I use them as an example, but here's our analytical builder in action:
It's showing a powerful widget-based UI for filtering and visualizing a spatial data set.
It's web based, so it has to handle as many users as are interested in it at any one time, which can be anywhere from zero to millions depending on the map and the moment.
It is only possible to **deliver** this capability because we can combine all these **free and low cost** things:

<x> utility computing so we can scale up and down depending on the current load (we use both Amazon and Google)
<x> all those utility computers run the free Linux operating system.
<x> all the data is stored in PostGIS, the free spatial database, which provides both storage and the analytical code for spatial intelligence
<x> the middle tier is built with the free Ruby-on-Rails framework
<x> the maps are drawn using a free Mapnik renderer
<x> the web front end map is a Leaflet map
That's a lot of free parts to start with! Carto has written a great deal of our own code, sure, but without the foundation of free building blocks, it would never have gotten started in the first place.

Carto provides spatial analysis and data access APIs.
It's a kind of geospatial information utility.
And it's built on top of a <x> commodity computing utility.
Which is built on top of a <x> commodity electrical utility.
And we know that our customers are
<x> building software-as-a-service offerings on top of our own.
Real estate information services, agricultural analysis services, satellite image catalogue services.

Each layer of cheap service and free software is
opening the door to a further layer.
Not because the technology is ground breaking,
but because it's free or cheap:
"when it become **so cheap** that anyone can it, **that's** where it gets interesting"

So, maybe you're not building a software-as-a-service business,
is this relevant to you?
Sure! Know you have yet another option for solving problems!

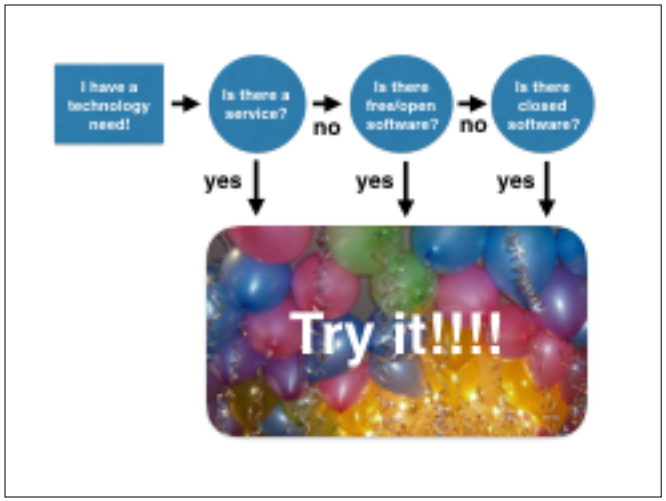

If you have a technology need, your first trip should **no longer** be to the software sales guy:
<x> is there a **service** that fulfills the need? is it economic for you?
<x> is there **free software** that fulfills the need? can you make it work for you?
<x> finally, is there **paid software** that fulfill the need? can you make it work for you?
Because there's **no capital cost** to experimenting,
with low cost services or free software, there's no excuse not to experiment!

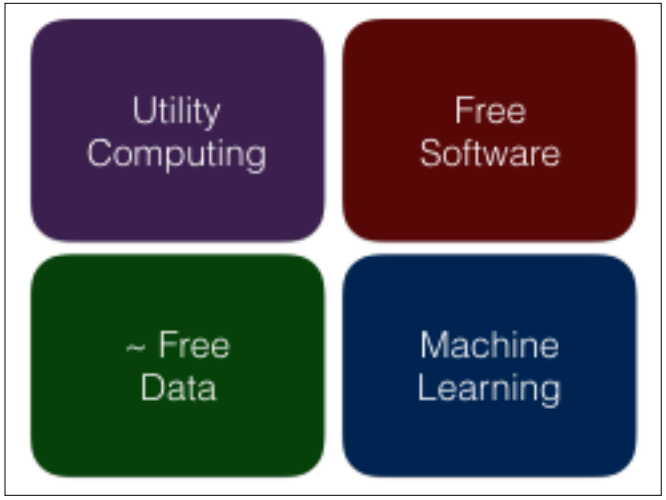

So,
- Computing is a utility
- Software is free
- Data is almost free
- Machine learning is available to all

Why is data almost free… what's going on?

This is a Texas Instruments TI4100 NAVSTAR.
It was the first commercial GPS receiver, and it came out in 1981.
It had an accuracy of around 10 metres.
It was 8 by 8 by 15 inches and weighed only 53 pounds,
not including the pair of car batteries you'd need to run it if you were in the field.
<x> It cost $120,000.

What could you do, with a TI4100 NAVSTAR, back in 1981?
What most mapping agencies did was
drag them to the tops of mountains to confirm or deny the location of their key survey monuments while establishing NAD83.
I suppose you could also mount one in your ship, if you had $120,000 to spare.

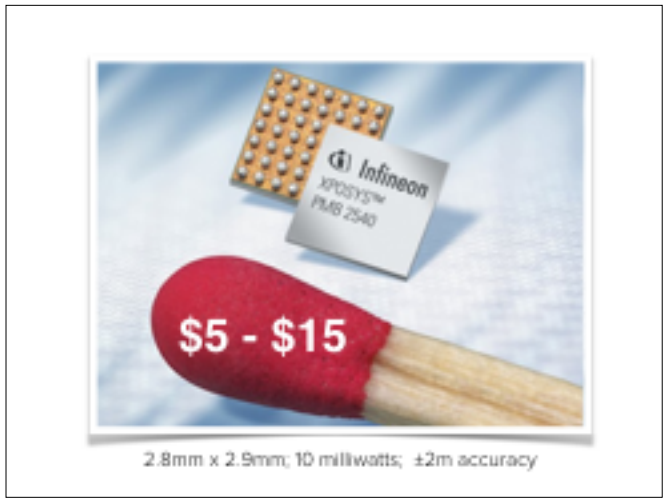

This is an Infineon XPOSYS, which went on the market in 2013.
It's 3 by 3 millimeters and draws 10 milliwatts.
It's accurate to within 2 meters.
<x> Depending on how many you buy, it'll cost you between $5 and $15.

What can you **NOT** do, with a GPS receiver the size of a match head?
"once it gets so cheap anyone can do it, then it gets interesting"

The first smart phone with a GPS chip came out in 2008.
So, a sub-$1000 internet connected mobile computer that always knows where it is.
Lots of folks saw the potential in that…

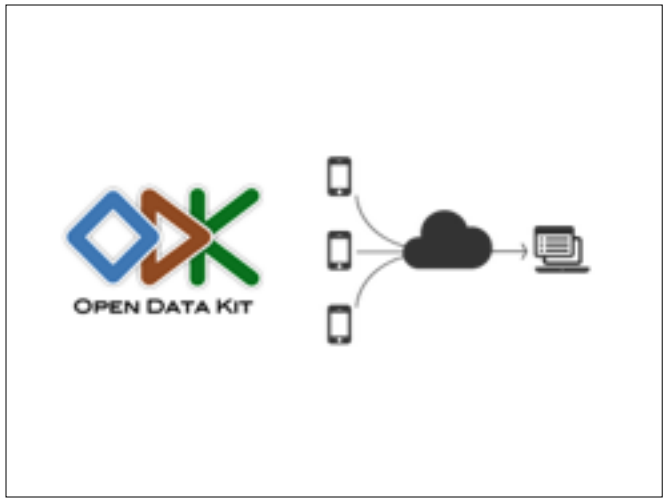

There's an open source project called "Open Data Kit"
that turns an Android device into a mobile data collection unit.

Walk around, type data into a form, it ends up in the cloud, in real time.

There's a software-as-a-service company called Fulcrum
that does pretty much exactly the same thing, but on a fee for service basis.
When location is **that cheap**, people start collecting it everywhere.

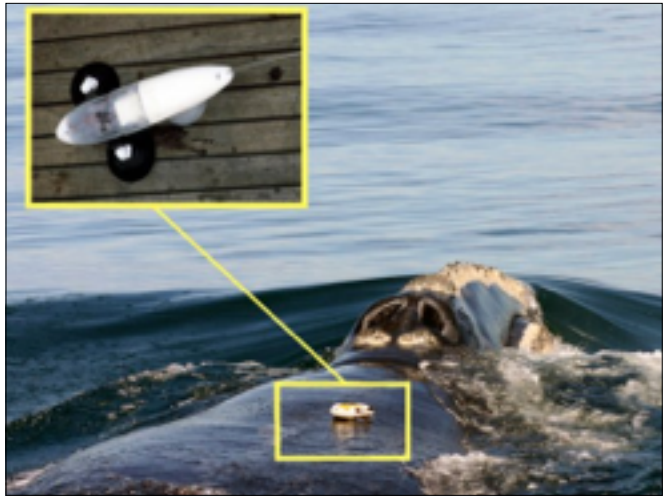

There are GPS chips on whales. http://ocean.si.edu/ocean-photos/whale-tracking-device

GPS chips on wolves. https://www.google.ca/search?q=gps+on+wolves

---

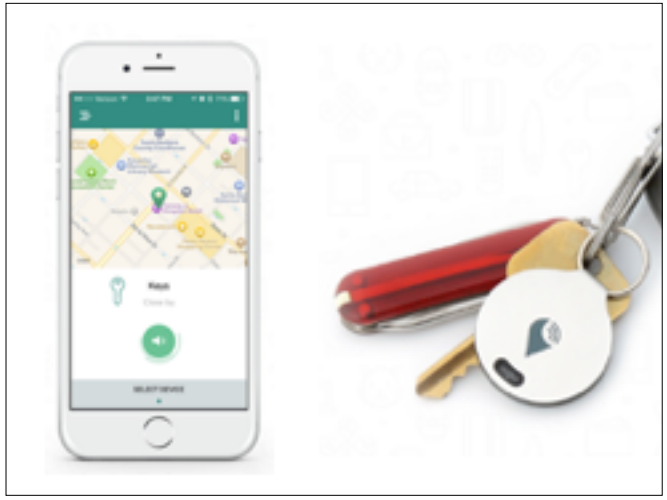

There are GPS chips on keychains, in case you drop them. https://www.thetrackr.com/

There are GPS chips on children, in case you lose them (or they lose themselves).
(I love that this model is for “kids **and** pets”, because really…)
Knowing **where things are** at all times is now practically free.
Attaching a **computer** to something is also now practically free.

Last year a UK computer magazine included a "Raspberry Pi Zero" **pasted into the cover** of every copy they sold.
Here are the specs of the Pi Zero,

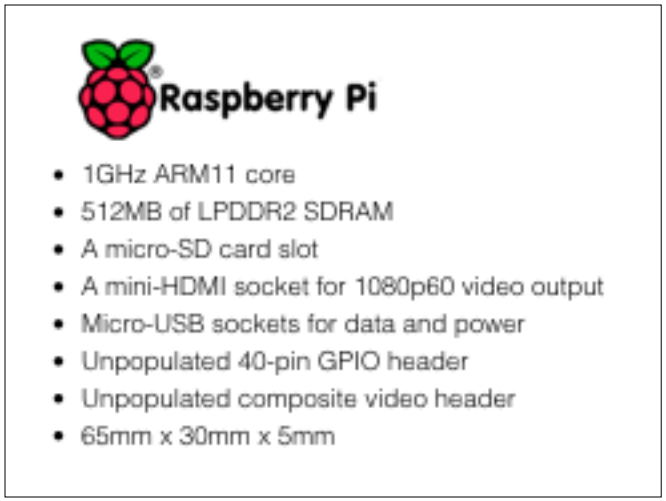

it's a full Linux computer, a lot better than the first computer I owned, actually better than the computer I owned in 2000.
Computers are now so cheap they are being handed out like America Online CDROMs.

Collecting imagery is now also plummeting from an expensive annual operation to a nearly free continuous operation.
What do you get when you combine a <x> GPS chip, a computer, and a model airplane? All the parts are incredibly cheap.

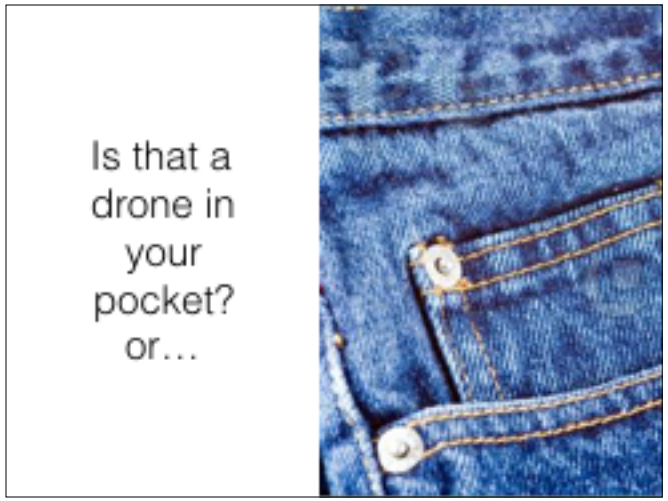

Just as consumer demand for cellular phones drove down the cost of CPUs, sensors, and wireless chips,

---

the consumer demand for drones is
driving down the cost of micro-aerial platform components.
And now, the market is lousy with drone start-ups,
all with a different take on bringing low altitude observations to market.

CyPhy, Ghost, 3D Robotics, and Skycatch
are all marketing drones specifically for data capture.

Many of them can operate autonomously, so regular automatic overflights are completely possible, generating a dynamic data stream of **changes**.
Industries with obvious need for regular monitoring, like construction and agriculture, are being targeted first.
Acquiring agricultural imagery annually or biannually via the NAIP program used to be a **big deal**.
Now acquiring it **weekly** is going to rapidly become the norm.

Because micro-drone technology is also being used in consumer oriented applications,
like GoPro cameras and toys for children,
I expect prices to continue to fall,
which could lead to even more surprising applications.

Cheap aerial platforms
mean that things previously **UN**mapped can be mapped.
This is the point at which things get **INTERESTING**.

The Hakai Institute on Calvert Island on the north-west BC coast used a drone to carry out a low-altitude map of bull-kelp in their study area, and observed a remarkable ecological story as a result.

Sea otters like to eat sea urchins, and urchins like to eat bull kelp. So otters suppress urchins, which allows kelp beds to flourish.

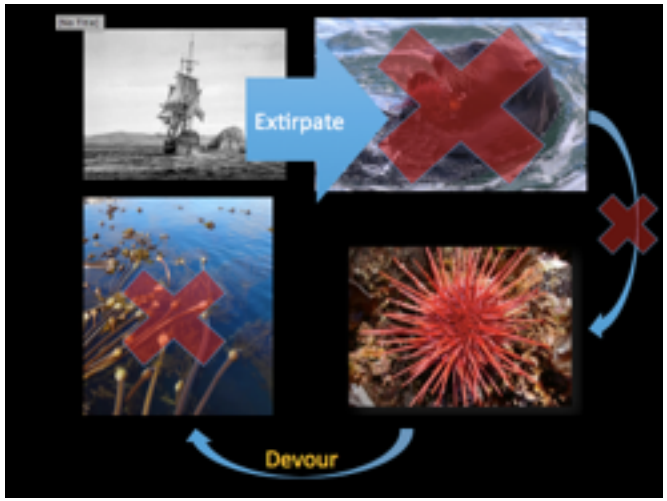

In the 19th and 20th centuries, sea otters were hunted nearly to extinction. So without sea otters, the urchin population exploded, they ate all the bull kelp, and kelp coverage got very low, and then stayed low for generations.

In the last decade, as sea otter populations have come back, and returned to old habits, the return of the old balance was visible over the course of just years of observation.
The otters came back, they ate the sea urchins, and the kelp beds grew back. The institute was only able to map this because they had access to cheap drone technology, that allowed them to map at high enough resolution to perceive the kelp, at just the right time and weather to avoid glare off the water, when visibility was high.

There are **whole categories** of things we've never mapped before that are going to start being mapped routinely.
"this is where it gets **interesting**"
So cheap sensors are increasing the **tempo** at which we collect information about the built and natural environment.

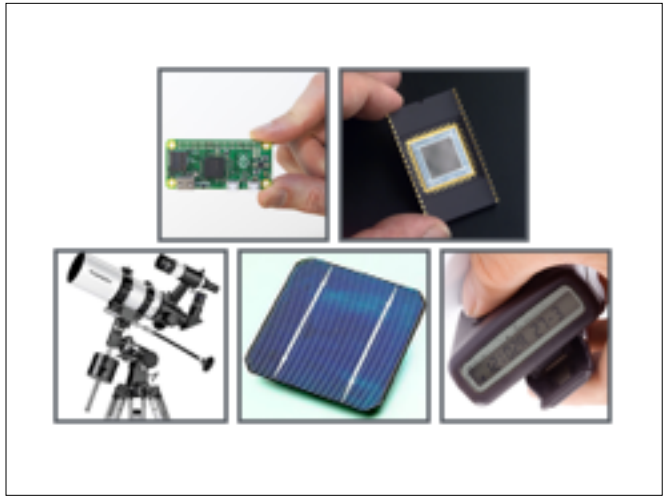

What do you get when you combine a cheap microprocessor with a commodity camera sensor, a hobbyist telescope, some solar panels, and a pager radio?
And then you loft the whole mess into space?

You get an orbital imagery platform!
These satellites from Planet Labs measure 10 by 10 by 30 centimetres,
so about the size of your forearm.
The measurements aren't random,
the 10 centimetre dimension allows the satellite to be added to standardized “micro-sat" payloads tucked in the corners of other launches.

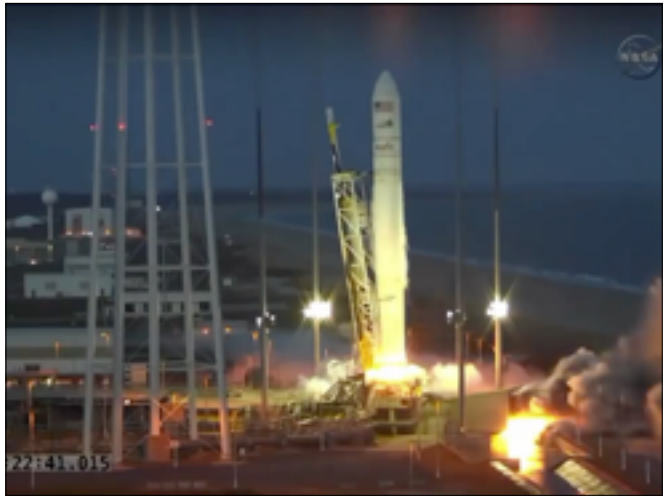

By hitchhiking on other payload launches and using cheap hardware,
Planet Labs has been able to launch dozens of earth observation satellites in just a few years.
By riding the downward cost curve of technology and now launch costs, this Silicon Valley start-up now has **more orbital sensors** than any organization **including** the US Government.
They make their satellites small and cheap so they can launch them in bulk, and so that the occasional failure isn't a catastrophic setback.
They lost 26 satellites in this launch, but they were back launching new satellites within a few months.

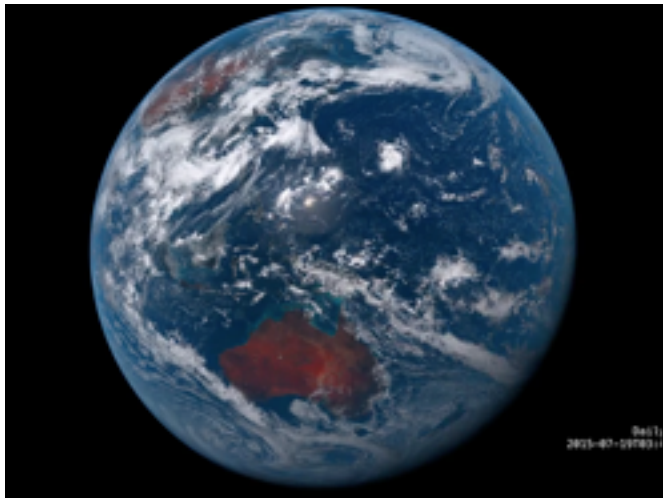

Planet Labs will image the whole earth daily when they have their full constellation in orbit.
For us as geo data consumers,
it means we are sitting on the threshold of an age of imagery that is not only cheap,
but also incredibly dynamic.

(Himawari-8)

That's a sea change in how we think about earth observation.
Earth observation is not an episodic activity anymore, it's continuous.
Previously we've been restricted to imaging areas immediately **after** an important event. Here's the tailing pond at Mount Polley in British Columbia shortly **after** the containment dam collapsed.

Before: July 29
After: Aug 5

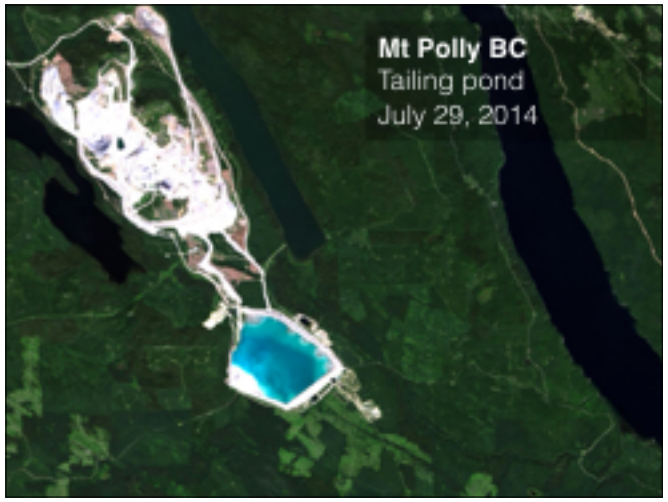

It would be nice to see how it looked immediately **before** the breach, but the closest landsat scene is several **weeks** earlier.

Before: July 29
After: Aug 5

The second most astonishing thing about Planet Labs,
after the fact that they are actually successfully building, launching and managing a constellation of dozens of micro-satellites,
is that they aren't even the only company in their category!
Skybox, UrtheCast, PlanetIq, and GeoOptics are all putting non-traditional sensors in orbit
and planning to radically reduce the cost of earth observation data.

As the **costs come down**, "interesting" things will come out.

I talked early about free software, focussing on the cost, but I'd like to revisit it briefly:
free software is free because it's built collaboratively using a license that promotes sharing of intellectual property.
Each collaborator has her own motivations for working on the software, and fixing or building the particular things they work on,
but in the end, they all share in the final product.
It's like a barn raising,

but where everyone gets to take home a copy of the barn, and so do all the spectators.

but where everyone gets to take home a copy of the barn, and so do all the spectators.

There is a free and open **data** analogue to free and open software, **collaborative knowledge building**.
The most widely known example is Wikipedia,
but there's an example in our field called OpenStreetMap.

The project started out as a reaction to the closed nature of UK Ordnance Survey mapping.
They maintained a monopoly on physical mapping in the UK, and leveraged as much revenue as possible from that monopoly, charging what municipalities and regions and corporations were able to pay for the data.
They thought they were doing OK, because their "clients" were happy.
The only people unhappy weren't clients, they were whiney academics and civil society types who couldn't afford the data, or wanted to process it in ways the Ordnance Survey license wouldn't allow.

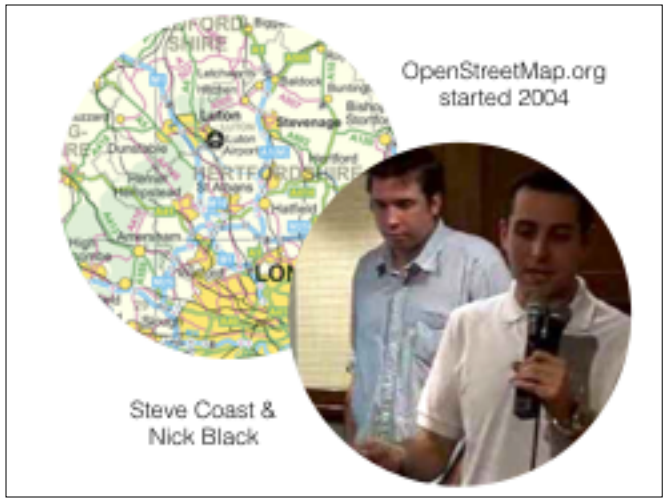

And eventually, not because they necessarily thought they would succeed, but more to give the finger to the Ordnance Survey,
a couple students spent a month building a basic infrastructure for crowd sourcing maps,
and then started telling people about how to do it.
Mapping their local neighbourhoods, building an unencumbered version of the data
that anyone could use, any way they wanted to.

Since then, the growth in coverage and accuracy has been astounding.
As a complete, fairly consistent map of the world that doesn't stop at national or provincial borders, OpenStreetMap is an easy resource to pick up and build on.
Academics test algorithms on it, companies use it as a default base map, new developers skip government sources of data and go straight to OpenStreetMap.
The OSM base is now in front of 100s of millions of users daily.

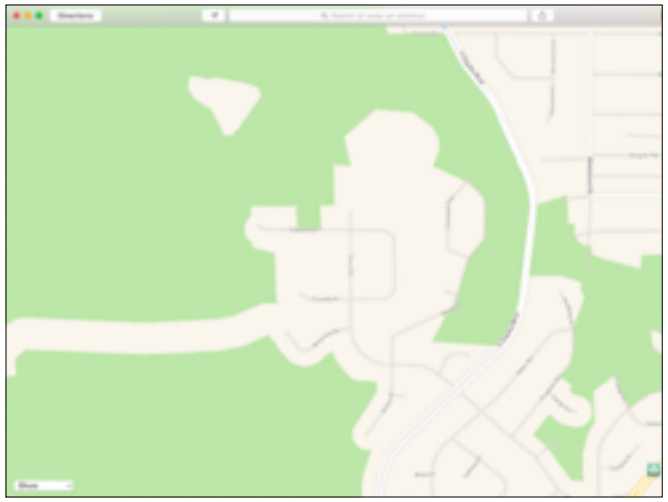

- Apple maps makes partial use of openstreetmap data

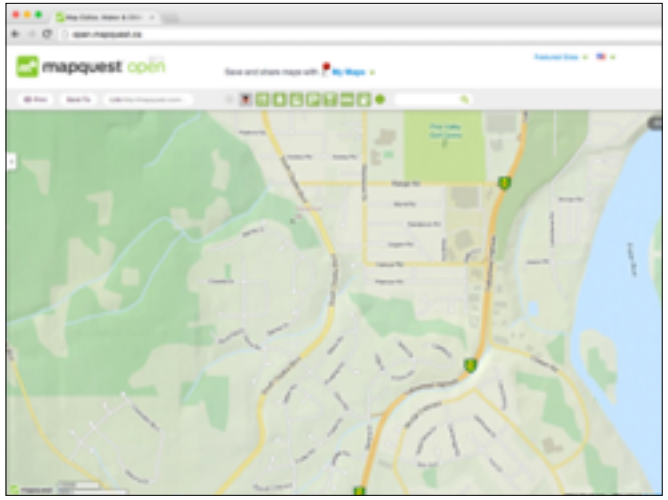

- MapQuest makes extensive use of openstreetmap
- Tableau provides OpenStreetMap map tiles they render themselves

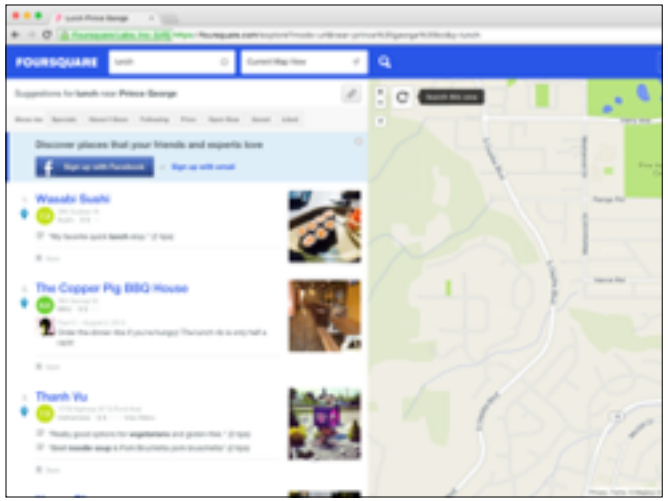

- Foursquare and Pinterest and National Geographic,
- Road Tripper, the Financial Times, gitHub, Etsy, The Guardian, USA Today,

---

the Wall Street Journal, the FCC and the Washington Post
are only *some* of the institutional and corporate customers
who use OpenStreet Map on their public maps, via map tiles produced by MapBox

- Everyone of our 200,000 users at CartoDB, including City of New York, NewsWeek, the LA Times, UN Environmental Program, and the US National Parks Service, makes use of OpenStreetMap in the base-maps we provide for free with our data visualization and analysis services

- So, OSM provides a base map, but what about addressing?

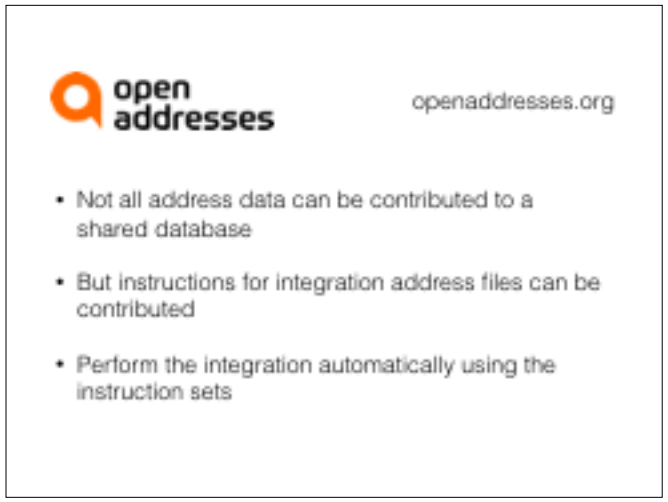

There's a project for addresses too, the imaginatively named "OpenAddresses".
This project recognized that the source of accurate and up-to-date address information was local governments that **might not be able** to upload their data to a collaboratively licensed central database.
So instead of building one big database of addresses,
they have built one big database of **recipes** for gathering and **standardizing** public address data.
For each source of address data, OpenAddresses holds a recipe card for where to download it from,
and how to transform it to fit it into the OpenAddresses schema.

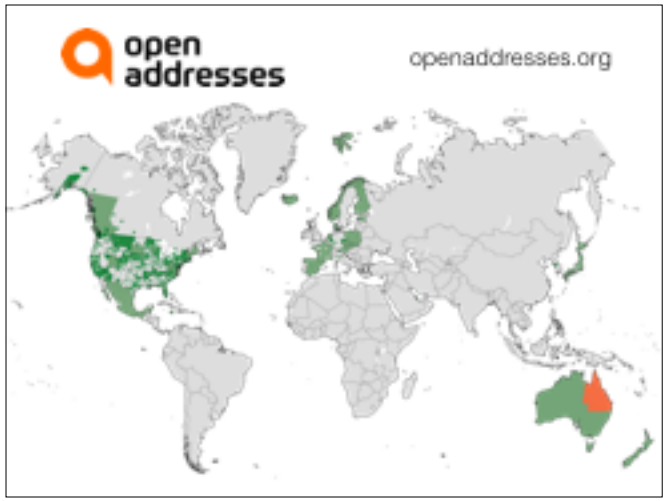

They now have pretty good coverage in the USA and are growing internationally.
If you want to share your address data widely,
writing an OpenAddresses recipe will be far more effective than standing up an open data portal.

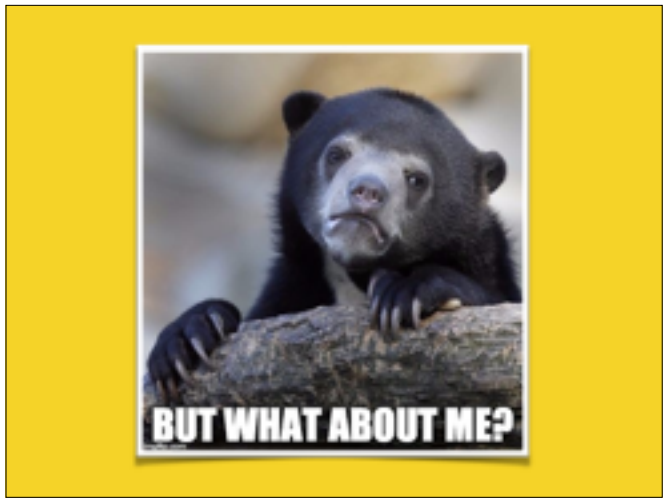

So data is being created more and more cheaply,
through ubiquitous GPS sensors,
through cheap drone sensors,
through cheap satellite sensors,
and through crowdsourcing efforts like OpenStreetMap and OpenAddresses and others.
So what?

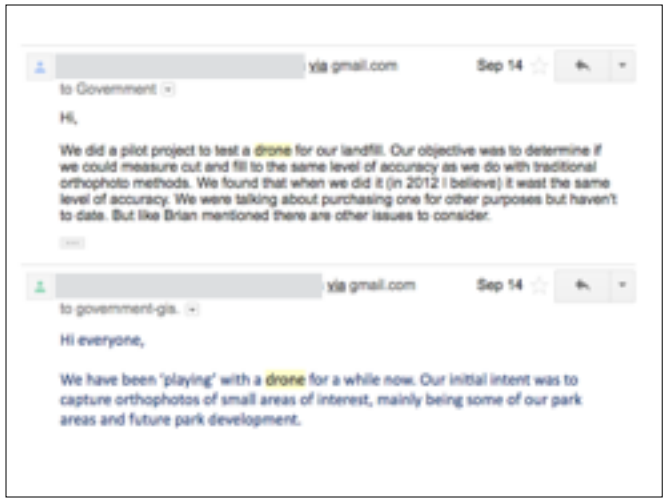

via gmail.com Sep 14

to Government

Hi,

We did a pilot project to test a drone for our landfill. Our objective was to determine if we could measure cut and fill to the same level of accuracy as we do with traditional orthophoto methods. We found that when we did it (in 2012 I believe) it wasn't the same level of accuracy. We were talking about purchasing one for other purposes but haven't to date. But like Brian mentioned there are other issues to consider.

via gmail.com Sep 14

to government-gis.

Hi everyone,

We have been 'playing' with a drone for a while now. Our initial intent was to capture orthophotos of small areas of interest, mainly being some of our park areas and future park development.

I am surprised and gratified at the amount of experimentation that governments are already doing with drone technology,
and there seems to be a rich local ecosystem of vendors willing to fly drone missions or help organizations get experience.

Satellite data is only getting cheaper, check out the web sites of some of the new providers,

and also check out what is available for free.

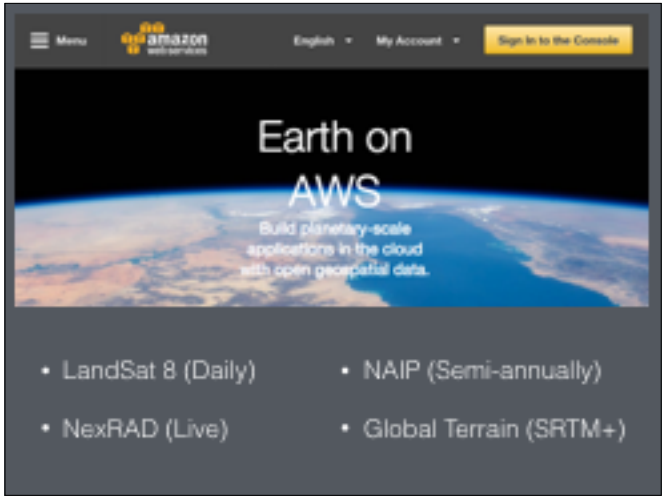

The "Earth on AWS" initiative has brought huge volumes of public domain remote sensing data into the AWS cloud, free for use.
Landsat 8 daily, live NextRad radar, NAIP imagery, and global terrain data.
The benefit to you is access to a huge corpus of data;
the benefit to Amazon is that, while you don't pay for **data**, you may pay for **CPU time**, which presumably you will use when you analyze the data.

And look what happens when data are made freely available:
this smartphone app lets you browse the historical stack of LandSat imagery in AWS, and view any of them.
The imagery is **free to access**, and that is "when things get **interesting**"

As far as what to do about collaborative data,
I'd like to relate the experience of the Portland Trimet transit agency.

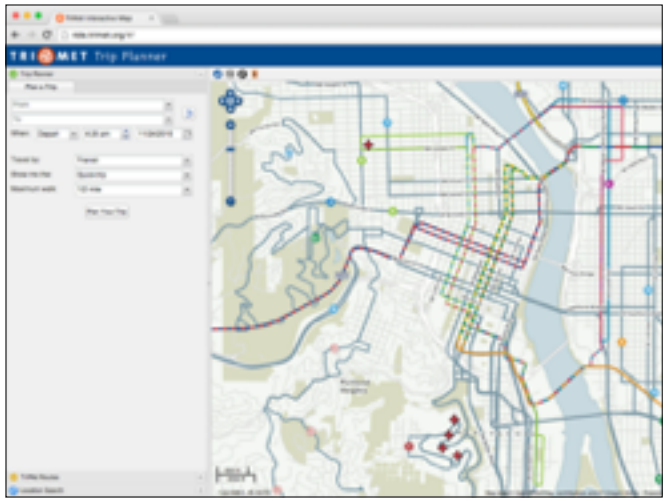

They run transit for the metro area of Portland, Oregon.
As a transit agency, they run a trip planner, and they need to keep an up-to-date street network to feed into that, and into their online maps.
For many years, they purchased data from NavTech,
but when they made corrections, they had to maintain a local copy and then reintegrate changes from NavTech when new data was shipped.
It was expensive to license, yet **still** involved a lot of staff work.

About 5 years ago, they looked at OpenStreetMap.
It wasn't quite up to the quality they needed, but they didn't say "oh well, guess we can't use it then".
Instead they estimated what it would cost to upgrade OSM to the quality they required, and compared that with the cost of their NavTech license.
Upgrading OSM was less than a year of license cost, and thereafter wouldn't cost them anything.
So they hired some interns who spent several month upgrading OSM, and then they **cancelled** their NavTech license.

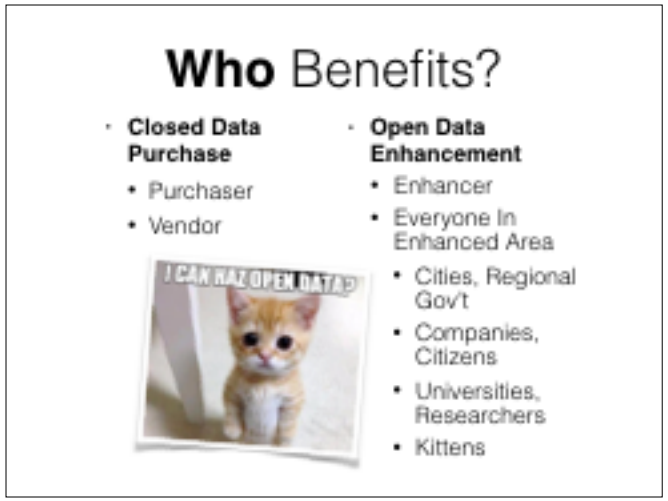

The best part is that not only have they benefitted, but **everyone in Portland** is now benefitting from the upgrades they made and continue to make, in the Portland area of OSM.
The data is up to date, and is being kept to a high quality by an agency that has operational reasons to ensure it remains accurate.
By treating the OSM data as if it were "theirs", they've gotten a better result than if they owned and kept a separate copy to themselves.
So,
* try out new sources of data, they could surprise you and
* consider engaging with **collaborative data communities**, you may find a new way to manage data of interest in your community.

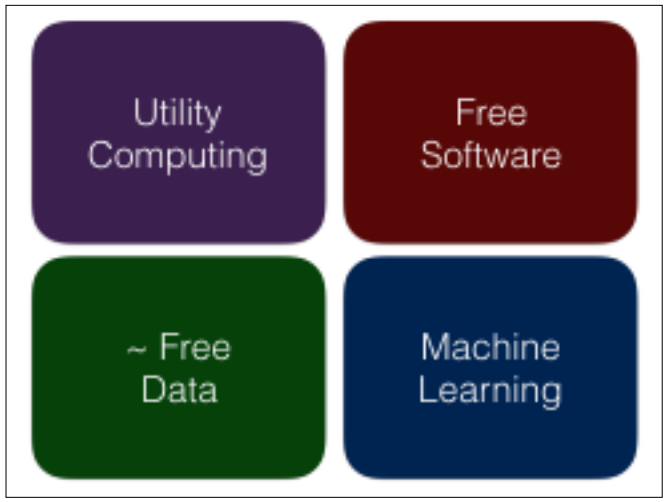

So,
- Computing is a utility
- Software is free
- Data is almost free
- Machine learning is available to all

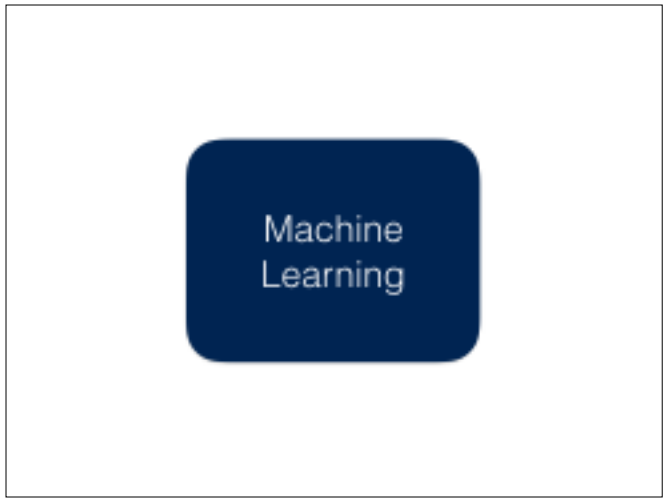

The last area, where cost is going way down, and accessibility is going way up, is machine learning.
We've heard a lot about self-driving cars, we're all hopefully very impressed with how well GMail filters spam, and we may have even asked Siri or "OK Google" a question and got a reasonably good answer,

---

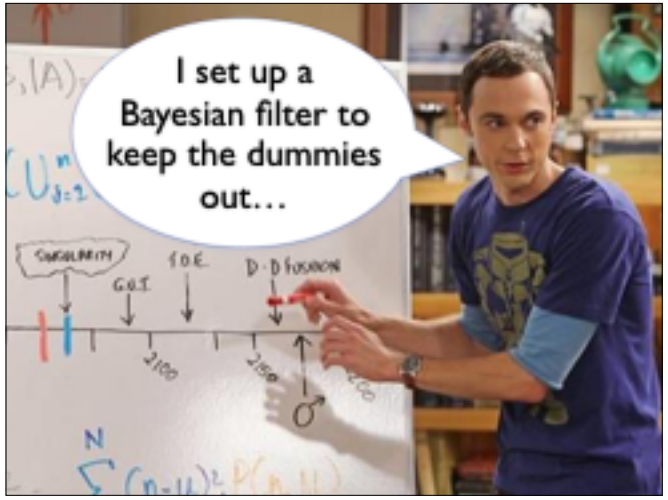

but machine learning probably **still** feels like something restricted to big Silicon Valley companies who have a stable of PhDs.

Not any more. Machine learning is easier. Well, easier.
Two things are going on:

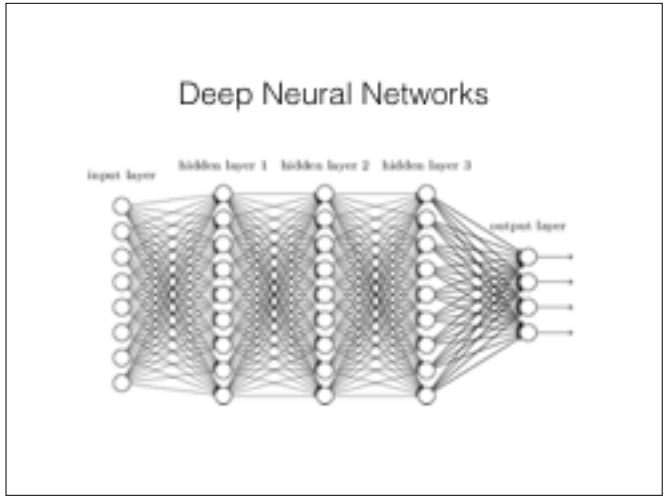

First, machine learning is just getting a lot better.
So called "deep" neural networks are able to create astounding results.
They work by stacking together simple neural networks.
The layers in the "middle" are called "hidden" because their inputs and outputs are not created by people, they are taking inputs from neural nets and feeding the output to further nets.
Until recently, the kind of computation needed for this trick wasn't economical.

My favourite example of a deep neural net is not spatial, it's the "deep art" project,
which incidentally runs entirely on the Amazon cloud.
Here's a picture of my son, playing baseball.
<x> And here's a painting by Kandinsky.
<x> And here's a picture of my son playing baseball, in the style of Kandinsky.

Or as a Roman mosaic.

Or in the style of Emily Carr.
All produced automatically with a deep neural network.

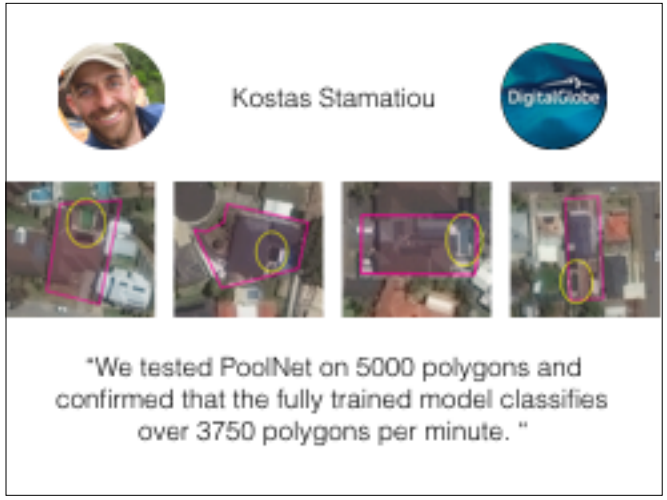

More close to home for us GIS folks, here's what a deep neural network can do as a feature classifier.
For tax purposes, the City of Adelaide wanted to know which properties had pools.
There are 700,000 properties. That's a lot of work!
So researchers started crowd sourcing to create a training set,
and then used that set to train a neural network, which automatically located all the pools.

In some cases, where the pool was partially occluded by trees or coverings, the neural network was actually better at classifying pools than the human crowd!

Or, returning to PlanetLabs for a moment.
One of the problems with launching cheap satellites and taking pictures is that cheap satellites have relatively poor attitude control,
so based on attitude and orbit information,
the images can only be placed on the ground with 100km accuracy.
Which is really, almost no accuracy at all, for a scene that is only 10km across.

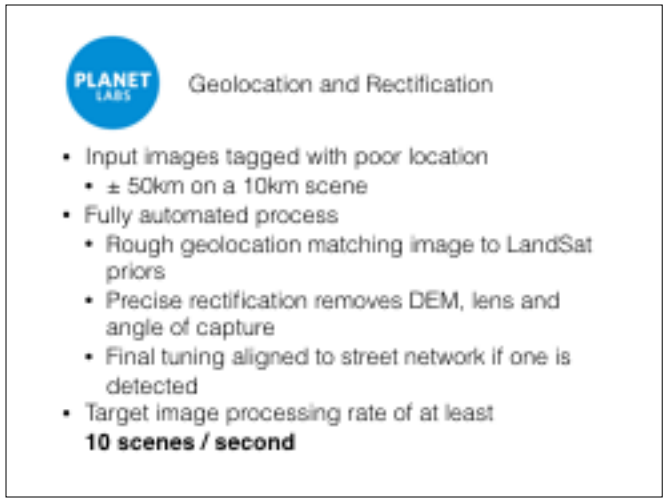

Without proper geo-referencing, the images are basically junk, so what does Planet do?
They figure out the correct referencing in post-processing, using machine learning.
There's a wonderful talk about this from a conference last year, called "Rectifying the Planet”,
but to get a sense of the difficulty of the problem,
the process has to be fully automatic, since 10 fresh scenes arrive for processing every **second**.

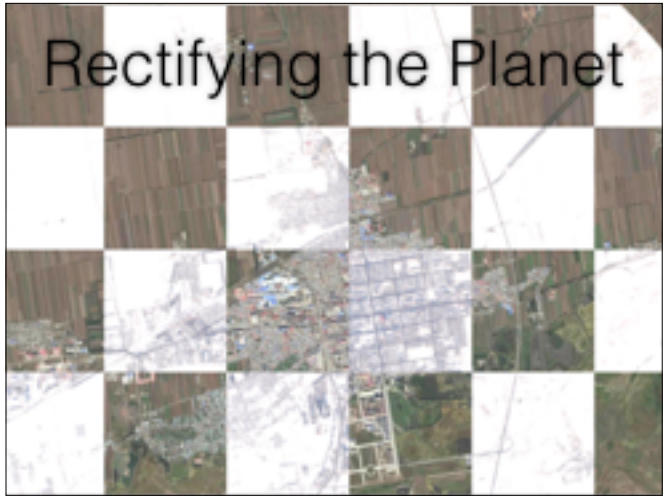

This is a pair of scenes from two seasons, after being run through the rectification process, and the registration is impressively accurate.
That they can achieve the result is incredible, and this profound accomplishment is only possible because of:

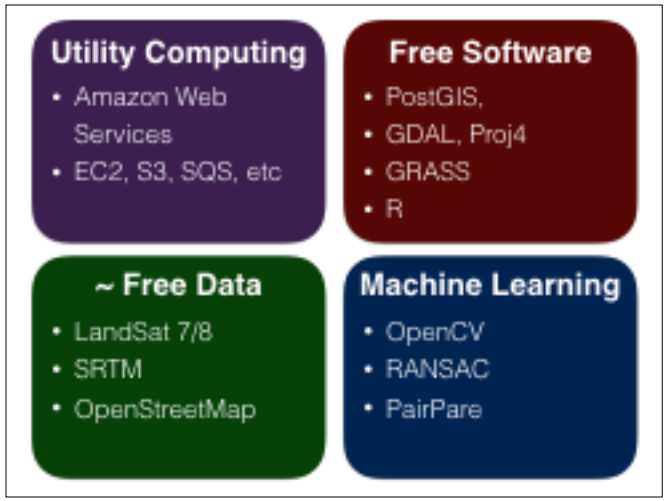

* free data (landsat reference imagery, and terrain data) and
* free software (GDAL, PostGIS, OpenCV, R, GRASS) and
* utility computing to handle the load, which
* allowed this amazing piece of **machine learning** to be built

---

Because of **access**, because the pieces got **cheap** enough,
Planet can do something very "**interesting**":
cheaply capture and process complete imagery of the earth.

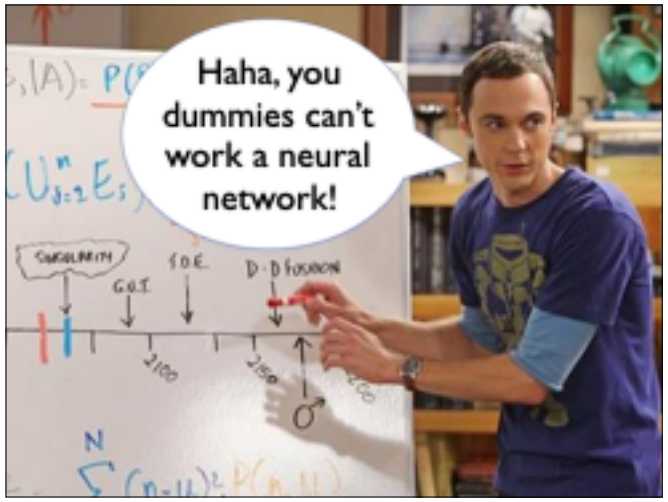

So far, while this is all very cool,
I have to admit that it still seems like it requires a PhD to do it.
And actually, I do think that it was mostly PhDs doing the work in those example.

But just as big utility computing companies
first perfected cloud technology for their internal use,
Amazon for e-commerce, Google for search,
and then started selling it to the general public.

In the **same** way, these **same** companies are
taking what they've learned **about AI** and making it available to the great unwashed.
Making it available to us non-PhDs.

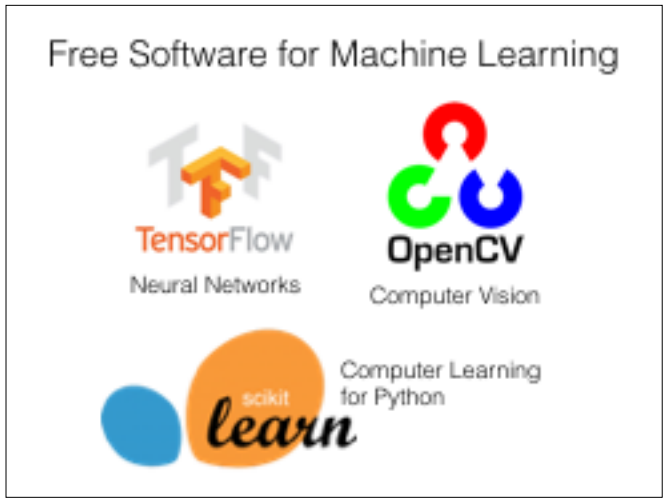

On the free software side, you can use the same library as Google uses for machine learning and neural networks,
<x> “TensorFlow”, which was released as open source last year.
<x> You can also use libraries like OpenCV for computer vision
<x> and SciKit-Learn for machine learning, both libraries are used in the academic community for research, but are now available for all.
Even better, Google and Microsoft have both started to release access to machine learning through utility computing **APIs**.
These allow you to access the power of AI algorithms with no server or software setup.

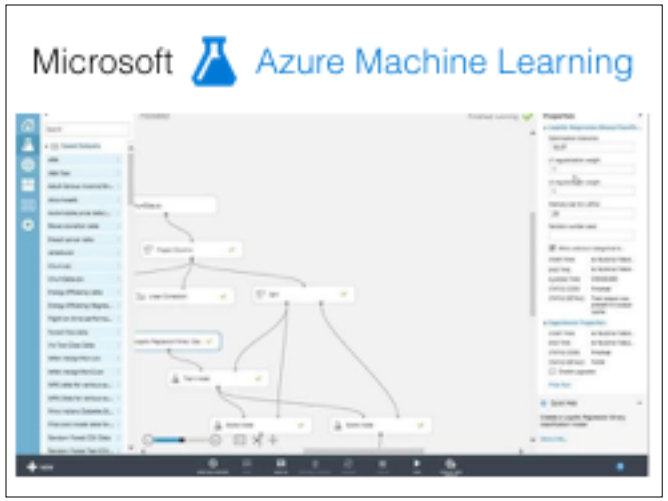

On the Microsoft Azure cloud, you can build a predictive learning model just by dragging together some boxes, then running your data through it to train it.

On the Google Cloud, some very impressive basic AI algorithms are now available with simple API calls.
* Natural language understanding
* Computer vision and object recognition
* Speech recognition
* Language translation, and
* General purpose machine learning,
so you can run models you build with the TensorFlow library using the huge Google neural network processing cloud
At this point, high quality machine learning is available to **anyone with basic programming skills**, on a simple service-based pricing basis.

So, for example, in Japan,
A young man returned home to help his aging parents run the family cucumber farm.
He was dismayed at how many hours his mother had to spend hand-grading cucumbers during the harvest season.
And so he decided to train a machine learning algorithm to classify cucumbers.

Then he built a little machine, to apply the classification and sort the cucumbers into bins.
Note the key role of extremely cheap hardware, as well as machine learning.
He did all the image classification himself with TensorFlow.
The only limiting step was the one that required human labor: generating a large enough collection of training images to feed the neural network.

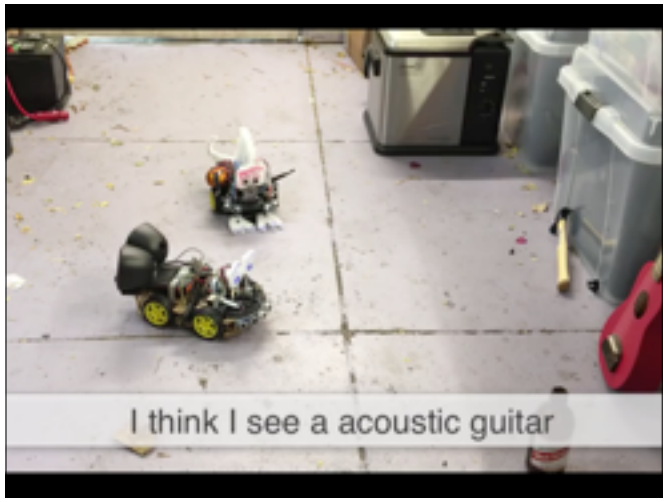

I came across this example only last month.
It's a hobby project, combining some simple wheels with a micro-controller, a camera, and Google's TensorFlow library.
The final product is a small robot that drives around and says the name of every object it recognizes.
The whole thing, the platform, the controller, the camera, the software, the whole thing cost less than $100.
https://www.oreilly.com/learning/how-to-build-a-robot-that-sees-with-100-and-tensorflow

---

So machine learning is here, and it’s available to anyone with even moderate technical still.

Even if you don't have a stable of PhDs, you might have tasks that machine learning can help with.
If you have repetitive jobs that involve classifying things into groups, there's probably a machine learning solution waiting for you out there.

Before I wrap up, I want to revisit my title for this talk, the “Undiscovered Country”, from Star Trek VI.
The screenwriter, Nicolas Meyer, had a good time with the scene, which has a bunch of quotations from Shakespeare, some in the original Klingon.
And he sells Gorkon’s toast as a quote from Shakespeare, which it **kind** of is, but it’s not about the future.
The phrase, “**the undiscovered country**”, comes from the soliloquy in **Hamlet**, which runs like this.

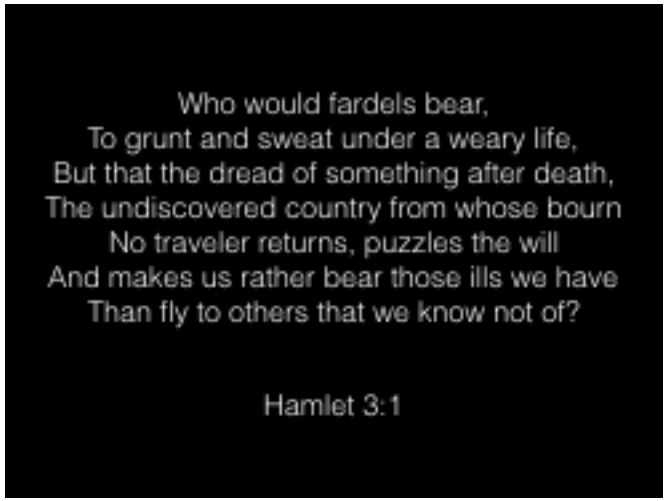

Who would fardels bear,
To grunt and sweat under a weary life,
But that the dread of something after death,
The undiscovered country from whose bourn
No traveler returns, puzzles the will
And makes us rather bear those ills we have
Than fly to others that we know not of?

Perhaps the undiscovered country is not the “future”,
but instead some place just as **interesting**,
yet far more dreadful.

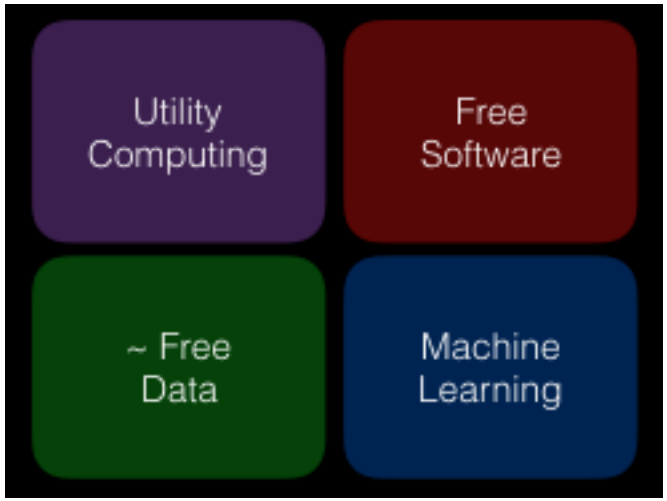

* cheap sensors, gathering imagery and location information, continuously in a stream of almost free data
* free software filtering and storing that data
* machine learning algorithms built with free software, extracting the features each camera sees, classifying each object, each person, each movement
* all running on utility computing infrastructure at scale.

Prices keep falling. For sensors, for platforms, for computers. And the capabilities of free software keep on growing. This is where we're headed, it's just a question of when.

This is a technology collection capable of telling you what roads need paving, where to trim the trees by the power lines, where the racoons are building a den by the school, where the fire hydrant has a leak that needs repairing. It's capable of doing it all in real time. This isn’t just “location intelligence”, it’s “location omniscience”.

In some ways this is the ultimate GIS nirvana: the "smart city" that IBM has been going on about in commercials lately.

Arthur C Clarke wrote a science fiction story, "The City and the Stars", about a city like this, and as **he** described it, it was a Utopia.
Everyone had what they needed, the city maintained itself through the aeons, the only downside was getting bored after centuries of good living.
Utopian idealism. Sounds nice! But it can go funny places...

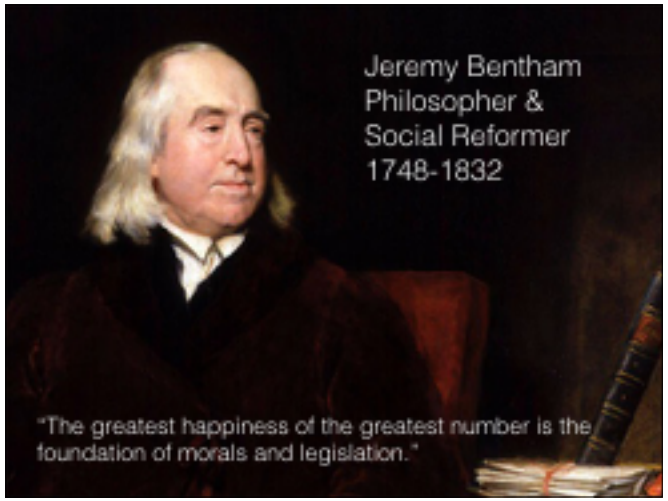

Jeremy Bentham was an English philosopher who lived in the early 1800s.
He advocated, for the time, pretty modern ideas about individual liberty,
and generally thought that more individual freedom was better.
He thought that people mostly have a good sense of right and wrong,
even if we don't always exercise it.

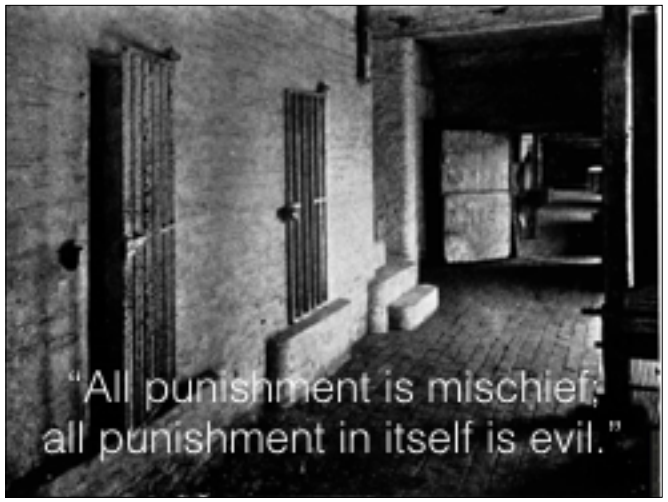

In that spirit, Bentham proposed a humane improvement
on the crowded, understaffed prisons of the era.
The goal was to provide a safer, quieter prison experience
where the prisoner could contemplate their crimes
and improve themselves before release, rather than fight it out in gang cells.

Bentham designed a new kind of prison, the **Panopticon**.
Each prisoner would have an individual cell, so no more fighting and mayhem.
And the whole thing could be run with minimal staff,
because the prisoners would **self-regulate**.
Deviations from the rules would result in harsh punishments.
And prisoners would know that **any deviation would be punished**, without fail.

How could they know that?
Because they would be under continuous surveillance.
Bentham reorganized the cell blocks around a central viewing area.
Every cell was completely visible from the central watch tower.
The watch tower was arranged so that from the outside,
it was not possible to know which way the guard was looking.
<x> Every prisoner would have to assume they were **always being watched**,
<x> Every prisoner would have **regulate their behaviour** accordingly.

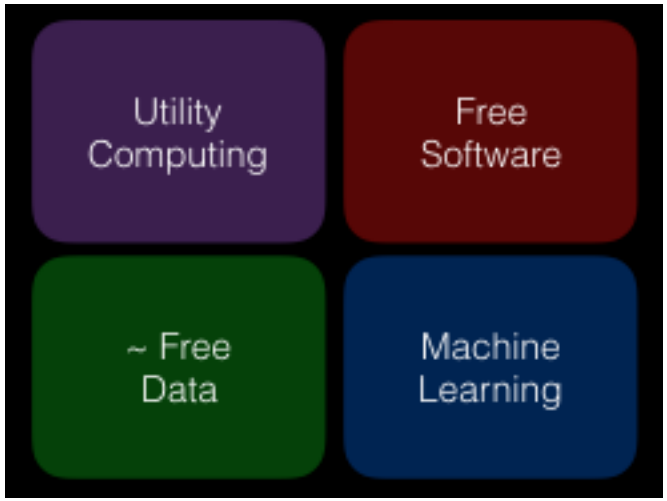

Return to the present, to the near future.
* data, from cheap sensors, everywhere
* free software filtering and storing
* machine learning algorithms, always running
* utility computing, so it never turns off

You don't need to carry a phone, for your location to be always known.
"Where's Paul? Let's look at the video feeds, ah,

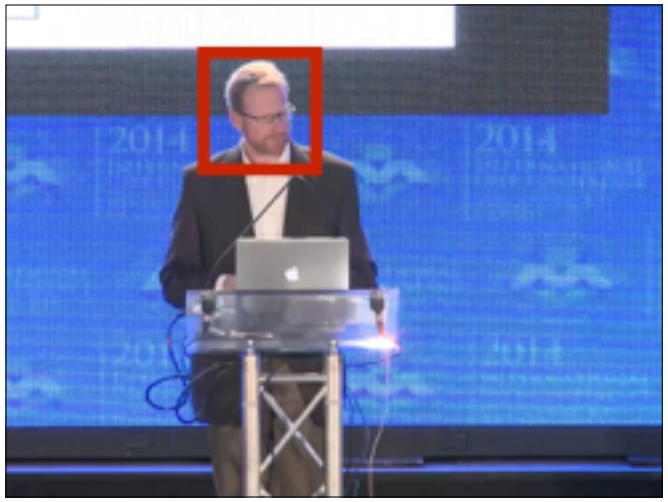

I see the system has found his face, he's giving a subversive talk right now in Montreal.
Let's log who is in the audience for more detailed profiling."

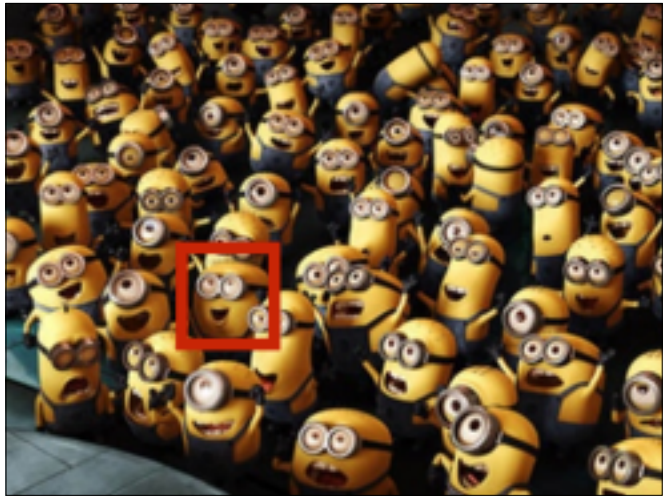

“Bring that one in for questioning”

The top two cities for closed circuit camera coverage right now are London, with 420K, and Beijing, with 470K.
It doesn't take a lot of imagination to figure out why those particular cities are leading the pack.

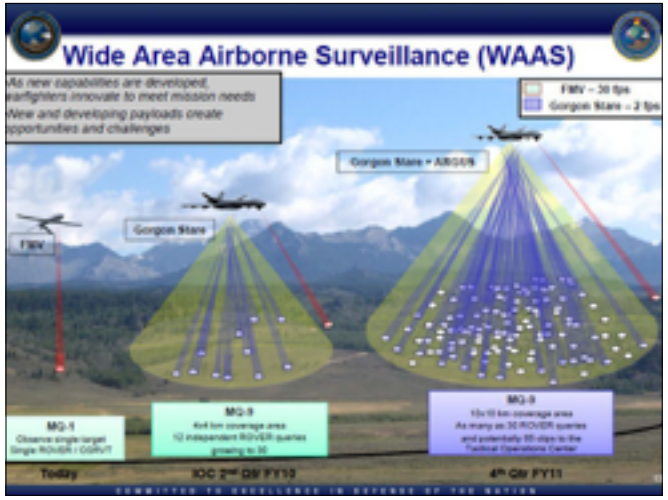

For the past 10 years, the US Airforce has been flying surveillance drones over Iraq and Afghanistan. That might sound like old news, but these drones are a lot more than a video camera on a model plane.

The system has the warm fuzzy name “Gorgon Stare”.
They aim for total awareness.
They have 365 cameras, and capture video at 12 frames per second over the entire field of view, in visible and infra-red.
At their standard operating altitude, that means then can watch an area of 100 km2 continuously.
These planes have been circling and circling the cities of the mid-east for the past decade.

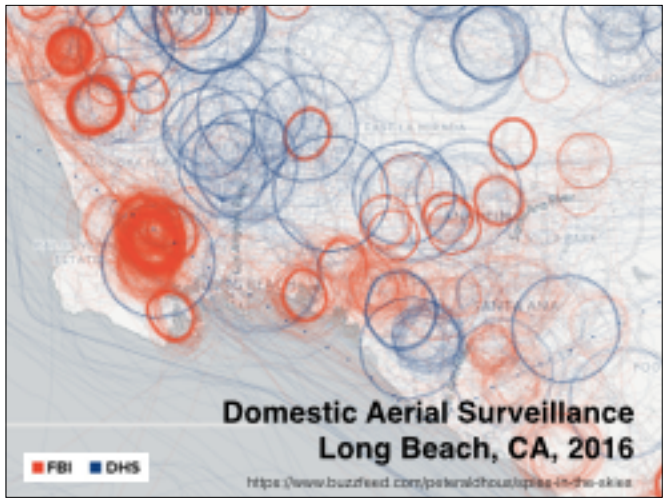

And who knows what is circling the cities back home?
The FBI and Department of Homeland Security wouldn’t comment on this Buzzfeed report,
that found federally owned planes in constant circling patterns over many major American cities.
Cheap sensors, everywhere, all the time.
This is just the first taste of life in the Panopticon.
And it’s ominous! Surely people will not go along with this!

And yet, social acceptance of life in the Panopticon is already growing, we're getting **used to it** very fast.
After the Vancouver hockey riots, we pored over photos, we turned in the culprits.

And in Montreal too.
We turned them in and we muttered to ourselves,
"Those idiots, Why didn't they realize they were **being recorded**, and **regulate their behaviour** accordingly."

The Panopticon doesn't **only** work on the side of authority, it can also be mobilized against it.
We are demanding that police carry cameras, that their cars carry cameras, in the hopes that they will regulate **their** behaviour accordingly.
That the consequences of bad cops being caught doing terrible things, as we have seen in the USA this past year, will encourage **self-regulation** by other officers.
Maybe this all sounds very **good** to you.
Rioters and violent police officers **SHOULD** regulate their behaviour, they **SHOULD** conform to social norms.

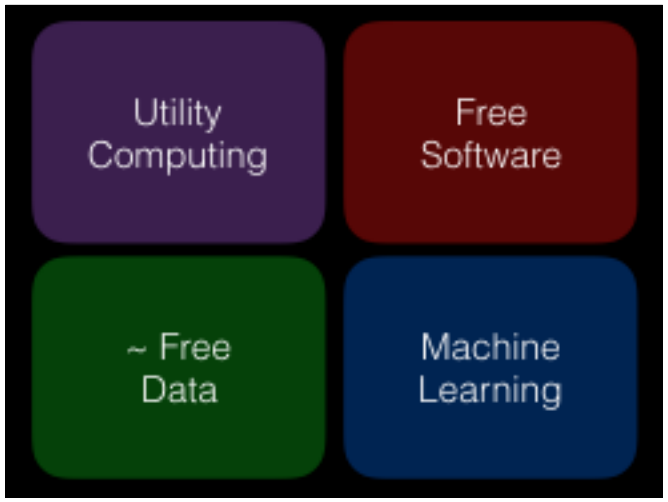

We're only getting started though.
* data, from cheap sensors everywhere
* free software filtering and storing
* machine learning algorithms, always running
* utility computing, so it never turns off

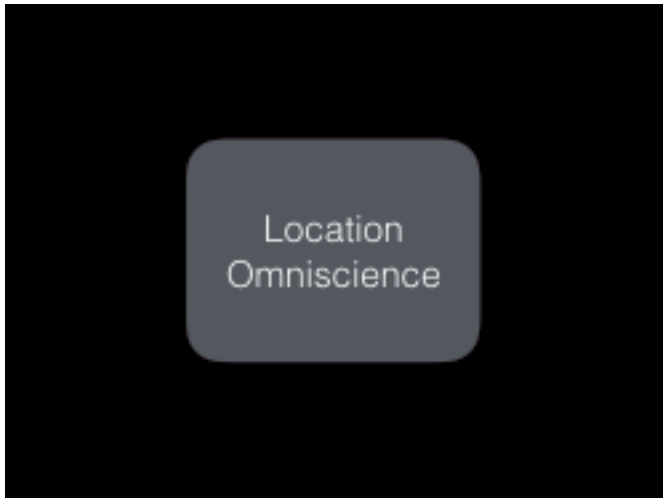

The beauty of the Panopticon, our digital overseer, is that we don't need to get 100% coverage

---

to get the **EFFECT** of 100% coverage. You don't need to be under surveillance to **FEEL** like you should act, as if you are.

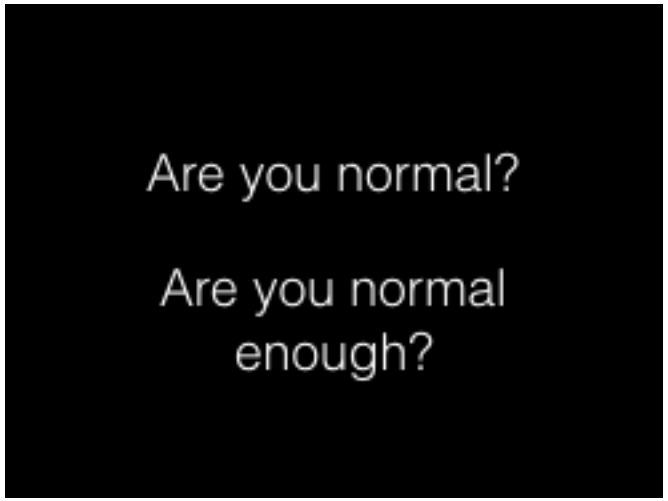

Conforming to social norms only sounds good when they're **your** social norms
and when you don't have to do it
24 hours a day, 7 days a week.

Our new, digital Panopticon promises to erase the last of the private spaces,
which should terrifying to all of us,
because very few of us are perfectly normal,
very few of us are perfectly average in every way.

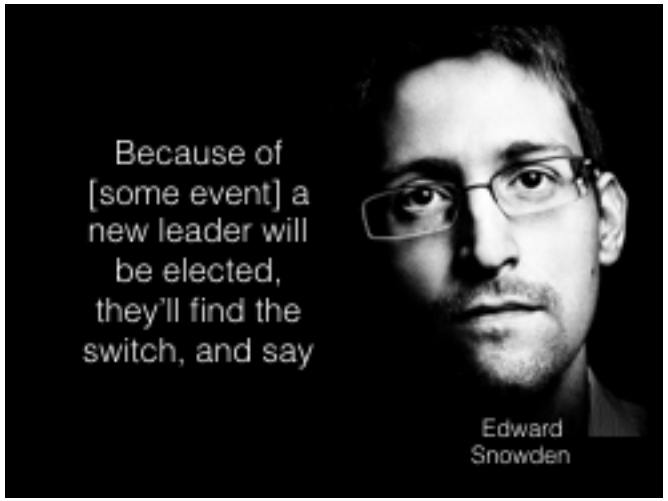

Edward Snowden, has described the risks in a terrifying way:
<x>
Because of [some event] a new leader will be elected,
they'll find the switch, and say

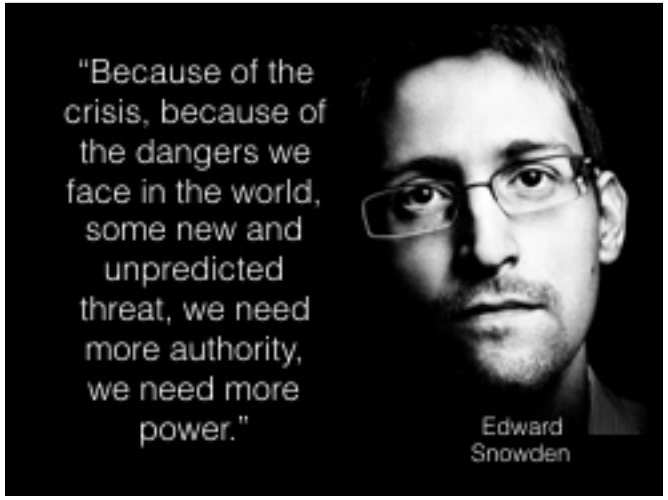

'Because of the crisis, because of the dangers we face in the world, some new and unpredicted threat, we need more authority, we need more power.'

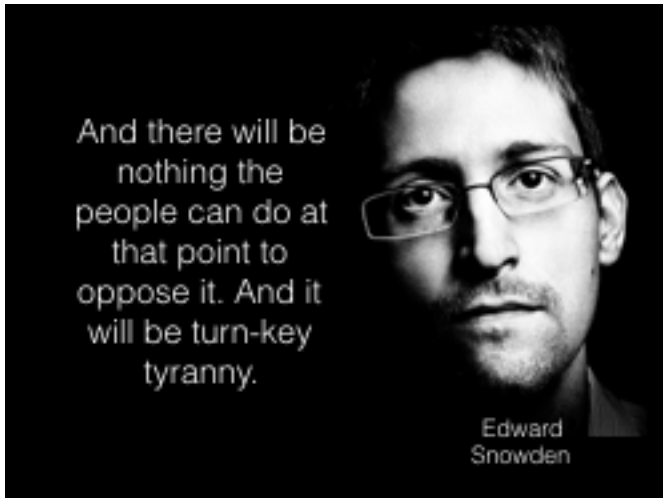

And there will be nothing the people can do at that point to oppose it. And it will be turnkey tyranny."

We have been given the gift of living in interesting times.
As interesting times go, they are pretty good:
we aren't living through the Chinese Cultural Revolution,
or the Irish Potato Famine,
this is pretty tame stuff, so far.

We're living in a time when you get buy a smartwatch for the price of a Starbucks panini.
We're living in a time when the number of tools at hand has never been higher, and the cost of using them has never been lower.
We're living in a time when new interesting things are being invented every day.

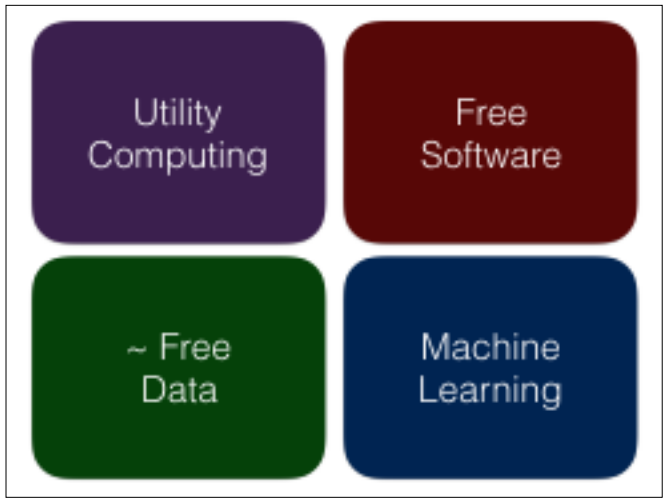

We have
* cheap sensors, so we can measure and optimize our communities
* free software, always getting more powerful
* machine learning algorithms, to make sense of the world
* utility computing, so no problem is too large

And all those cheap components
are going to unlock new surprises we haven't imagined yet.
After this conference, let’s all go home, and do something interesting.
Thank you very much.