

# The next 10 years of quantitative biology

Michael Schatz

March 25, 2014

Keystone Meeting on Big Data in Biology

[@mike\\_schatz](#) / [#KSBigData](#)

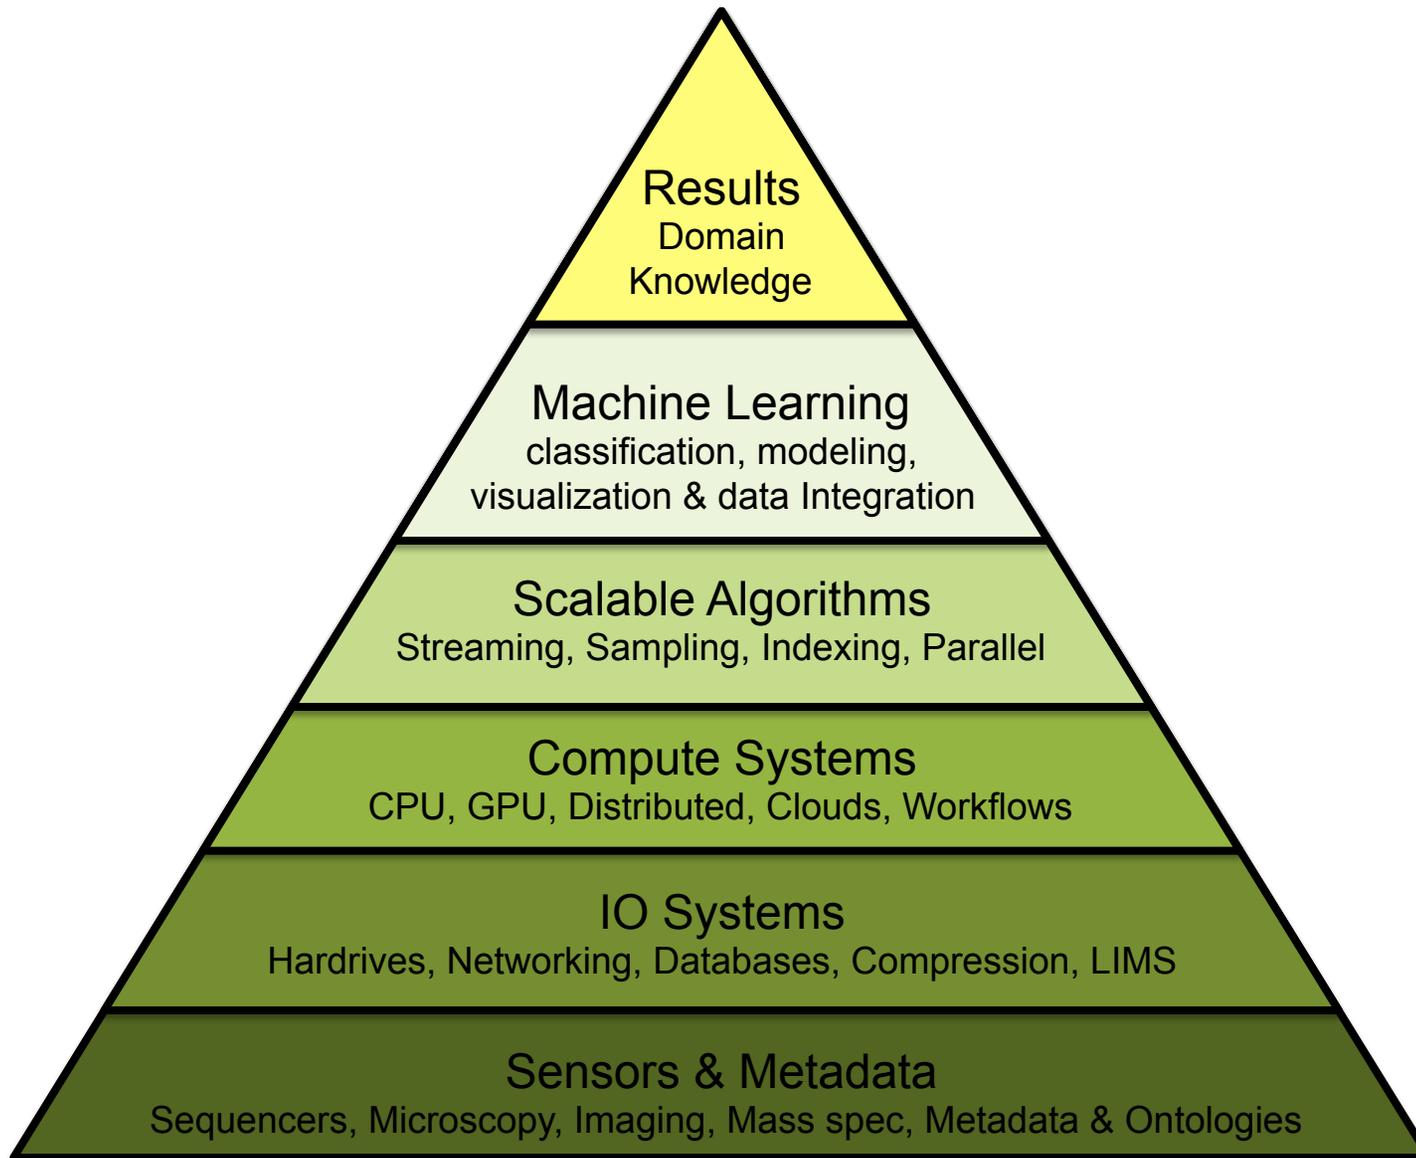


# Unsolved Questions in Biology

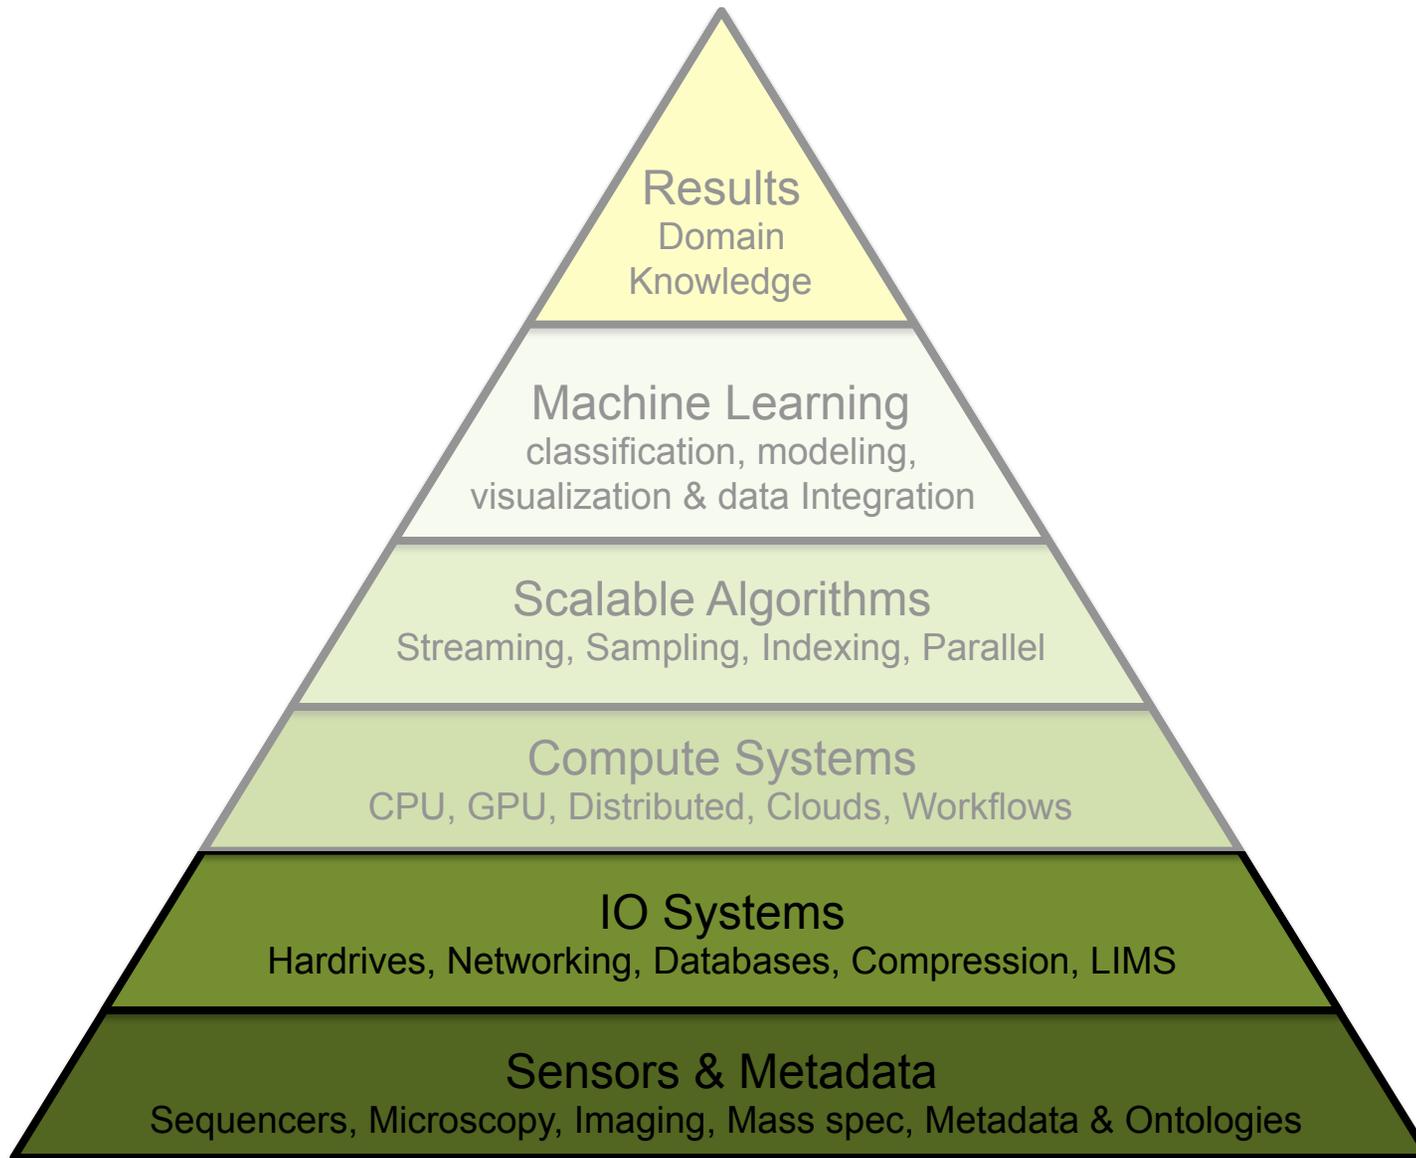
- What is your genome sequence?
- How does your genome compare to my genome?
- Where are the genes and how active are they?
- How does gene activity change during development?
- How does splicing change during development?
- How does methylation change during development?
- How does chromatin change during development?
- How is your genome folded in the cell?
- Where do proteins bind and regulate genes?
- What virus and microbes are living inside you?
- How do your mutations relate to disease?
- What drugs should we give you?
- Plus hundreds and hundreds more



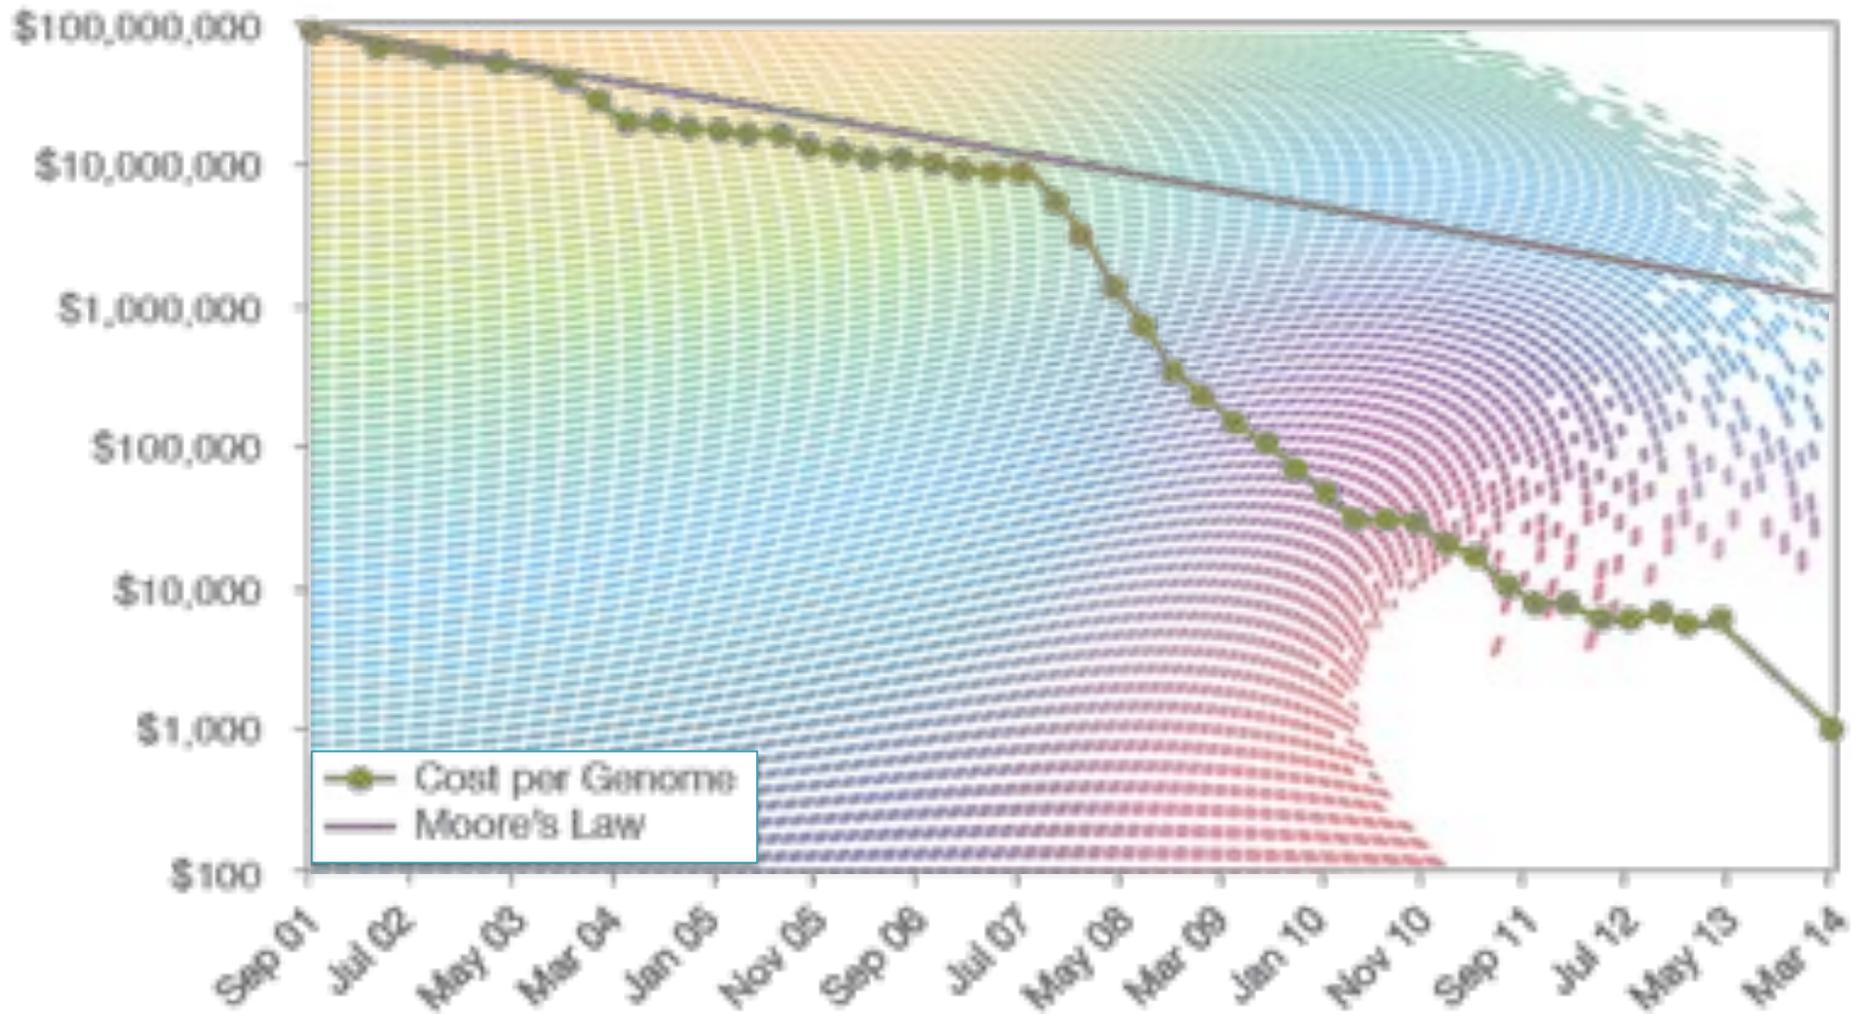
# Quantitative Biology Technologies



# Quantitative Biology Technologies



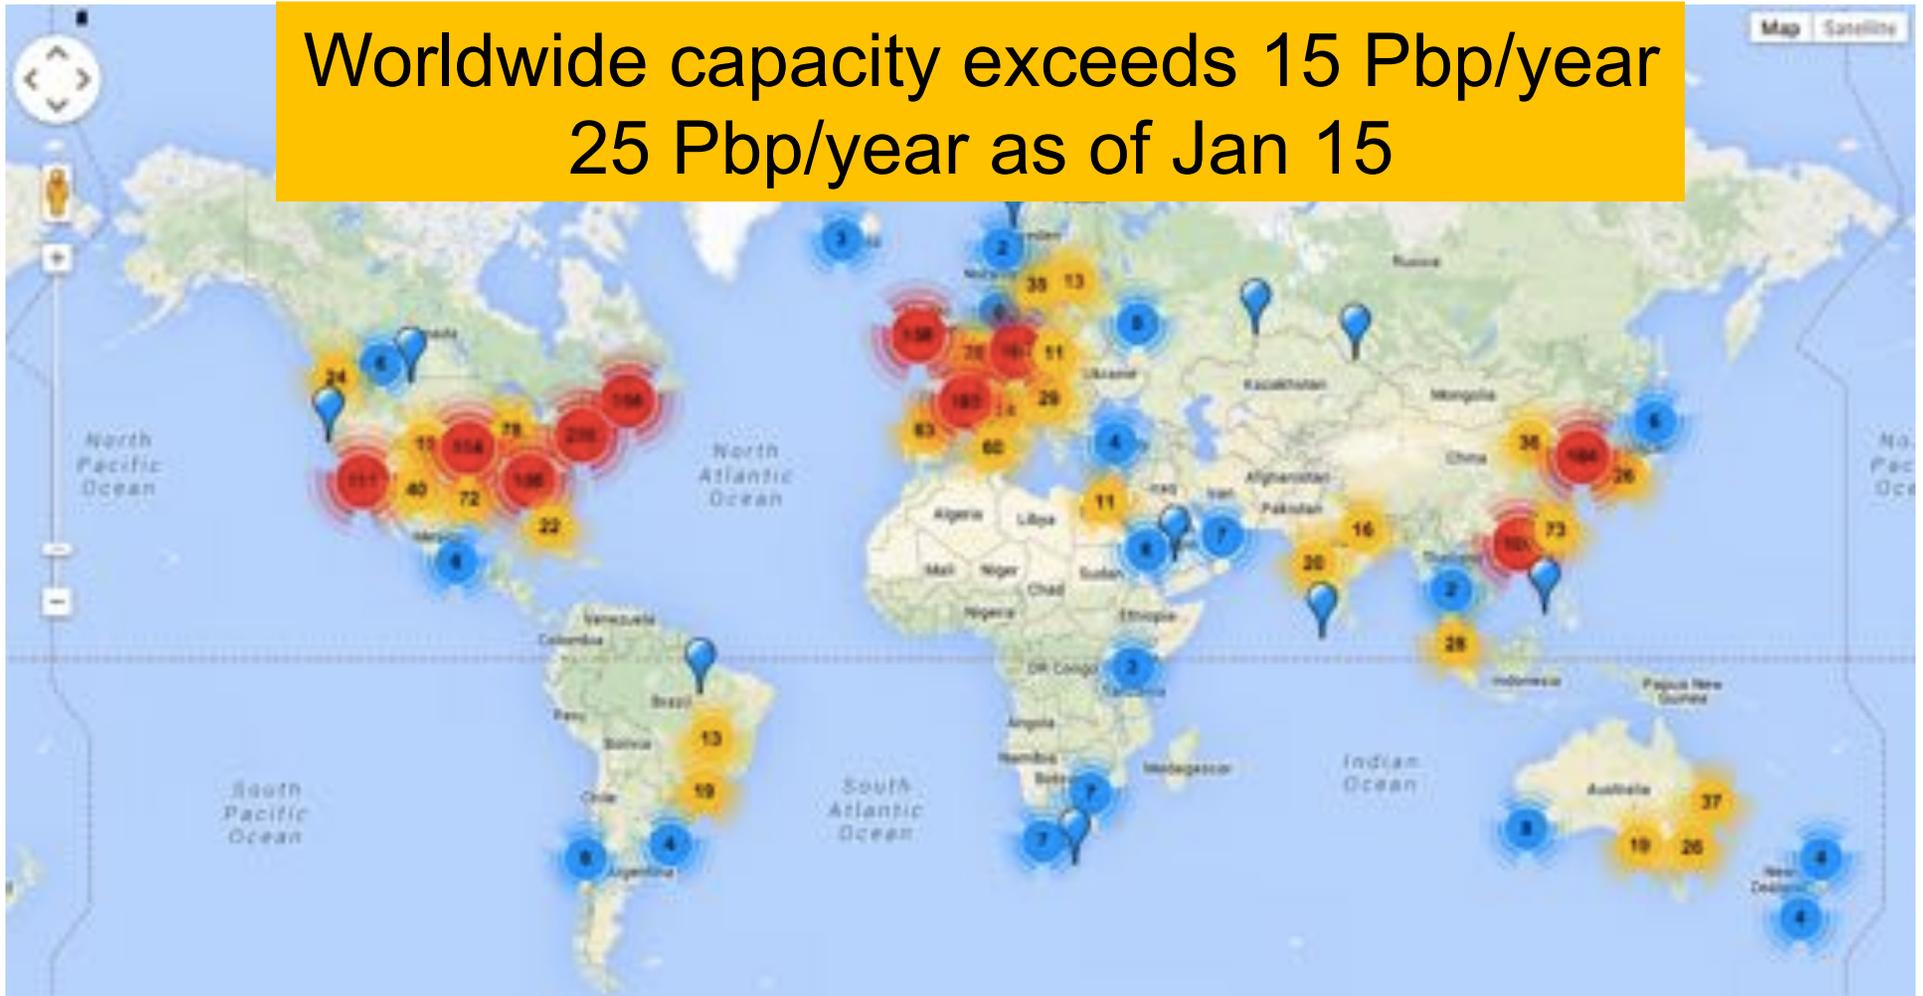
# Cost per Genome



<http://www.genome.gov/sequencingcosts/>

# Sequencing Centers

Worldwide capacity exceeds 15 Pbp/year  
25 Pbp/year as of Jan 15



**Next Generation Genomics: World Map of High-throughput Sequencers**  
<http://omicsmaps.com>

# How much is a petabyte?

Unit	Size
Byte	1
Kilobyte	1,000
Megabyte	1,000,000
Gigabyte	1,000,000,000
Terabyte	1,000,000,000,000
Petabyte	1,000,000,000,000,000

\*Technically a kilobyte is  $2^{10}$  and a petabyte is  $2^{50}$

# How much is a petabyte?



100 GB / Genome  
4.7GB / DVD  
~20 DVDs / Genome

X

10,000 Genomes

=

1PB Data  
200,000 DVDs



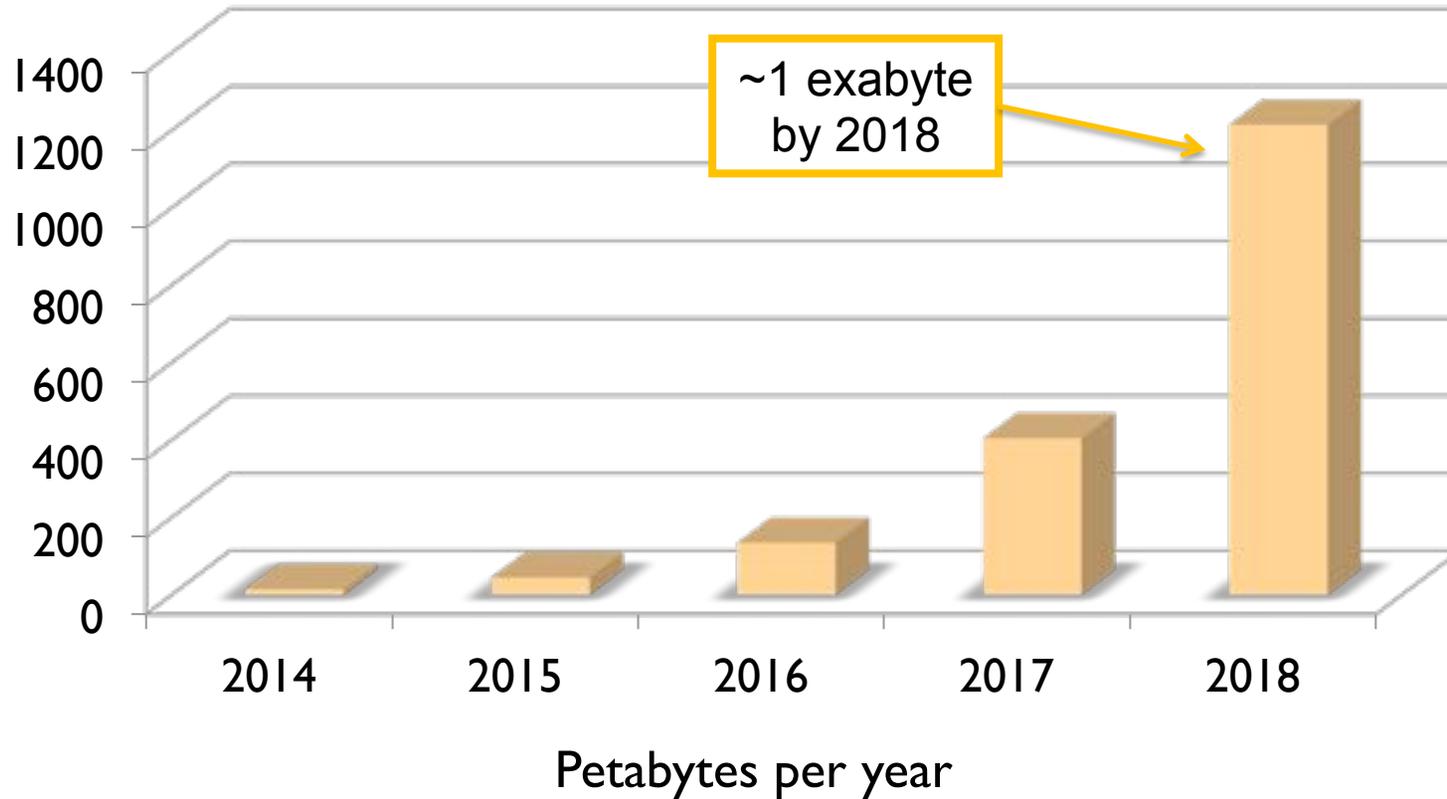
787 feet of DVDs  
~1/6 of a mile tall



500 2 TB drives  
\$500k

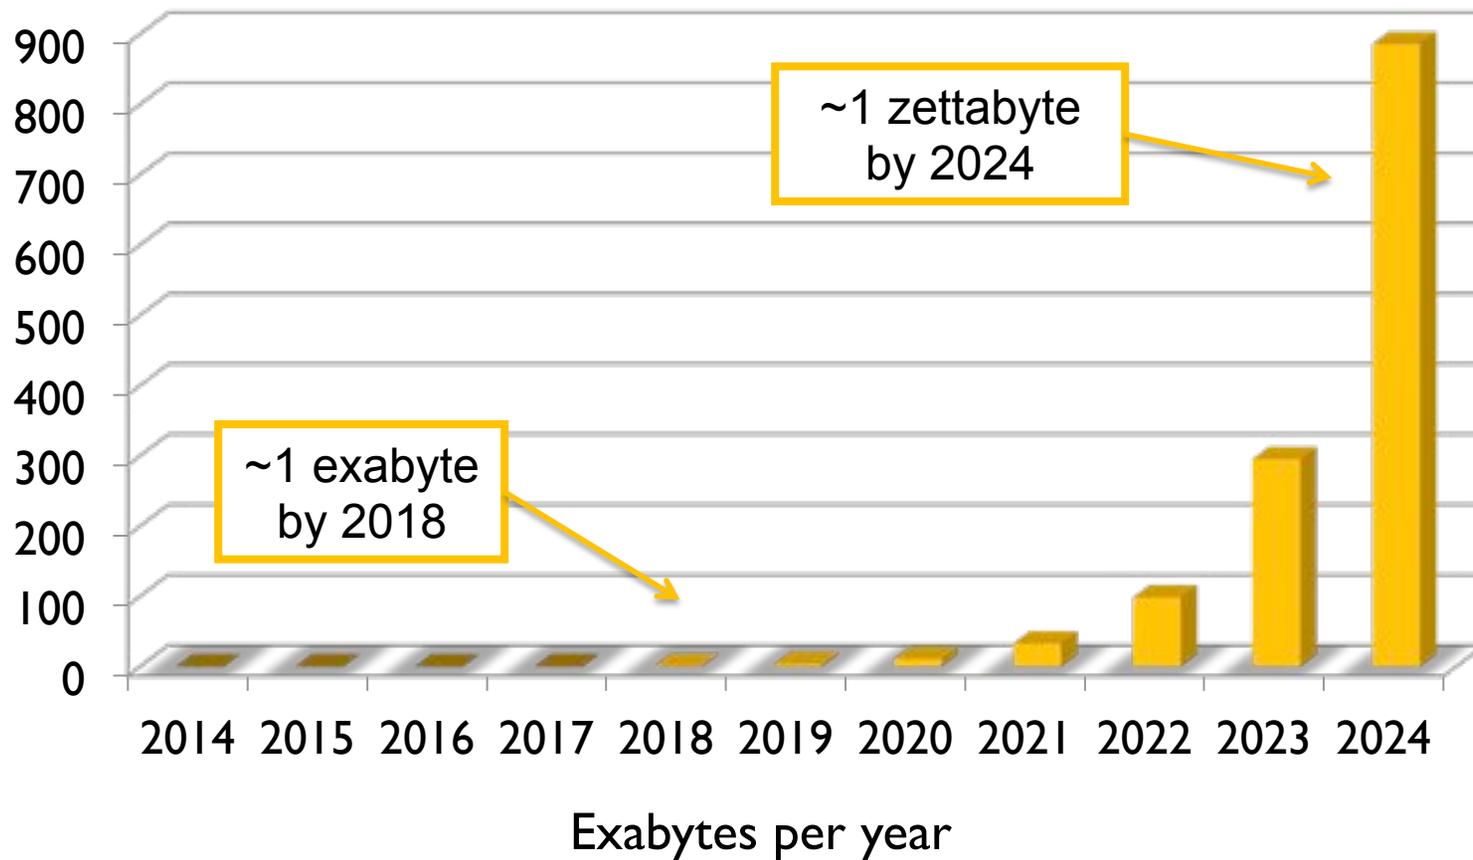
# DNA Data Tsunami

*Current world-wide sequencing capacity is growing at ~3x per year!*



# DNA Data Tsunami

*Current world-wide sequencing capacity is growing at ~3x per year!*



# How much is a zettabyte?

Unit	Size
Byte	1
Kilobyte	1,000
Megabyte	1,000,000
Gigabyte	1,000,000,000
Terabyte	1,000,000,000,000
Petabyte	1,000,000,000,000,000
Exabyte	1,000,000,000,000,000,000
Zettabyte	1,000,000,000,000,000,000,000

# How much is a zettabyte?



100 GB / Genome  
4.7GB / DVD  
~20 DVDs / Genome

X

10,000,000,000 Genomes

=

1ZB Data  
200,000,000,000 DVDs



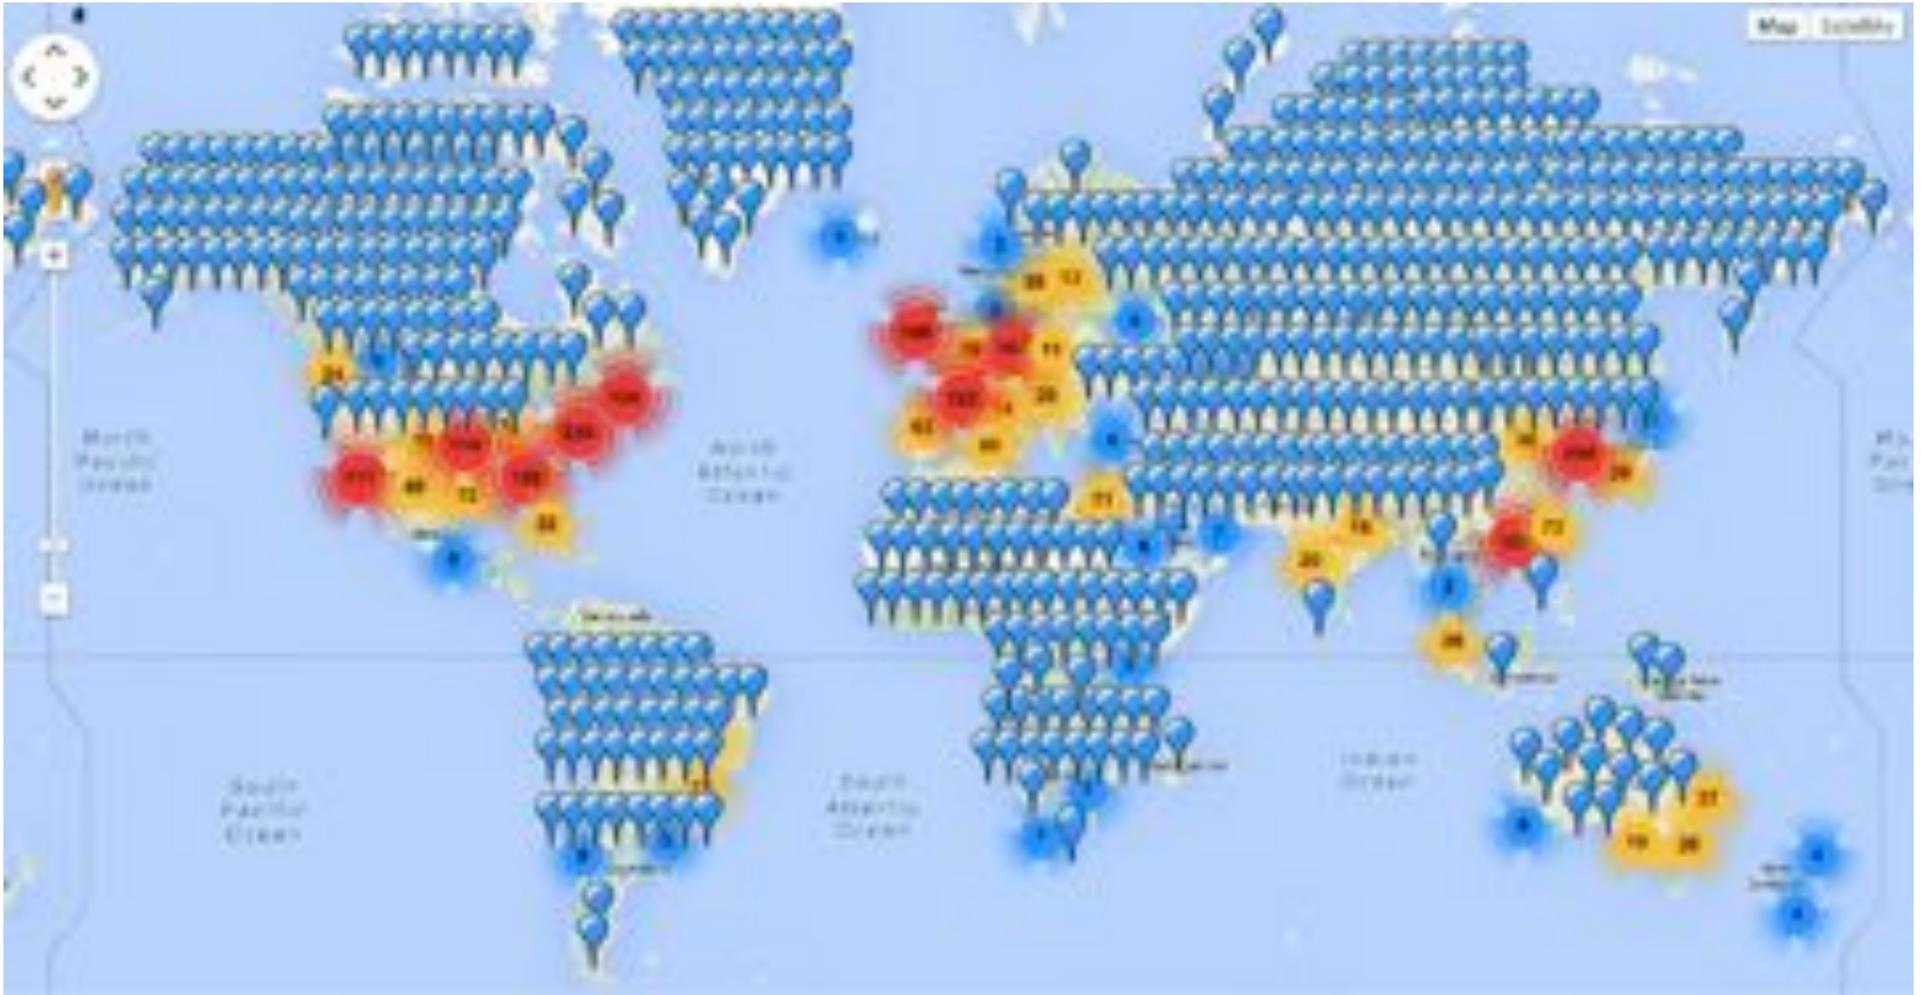
150,000 miles of DVDs  
~ 1/2 distance to moon



Both currently ~100Pb  
But growing exponentially



# Sequencing Centers



**Next Generation Genomics: World Map of High-throughput Sequencers**  
<http://omicsmaps.com>

# Biological Sensor Network



(@ewanbirney)



(@latimes)

***The rise of a digital immune system***

Schatz, MC, Phillippy, AM (2012) GigaScience 1:4

# Data Production & Collection

## Expect massive growth to sequencing and other biological sensor data over the next 10 years

- Exascale biology is certain, zettascale on the horizon
- Compression helps, but need to aggressively throw out data
- Requires careful consideration of the “preciousness” of the sample

## Major data producers concentrated in hospitals, universities, agricultural companies, research institutes

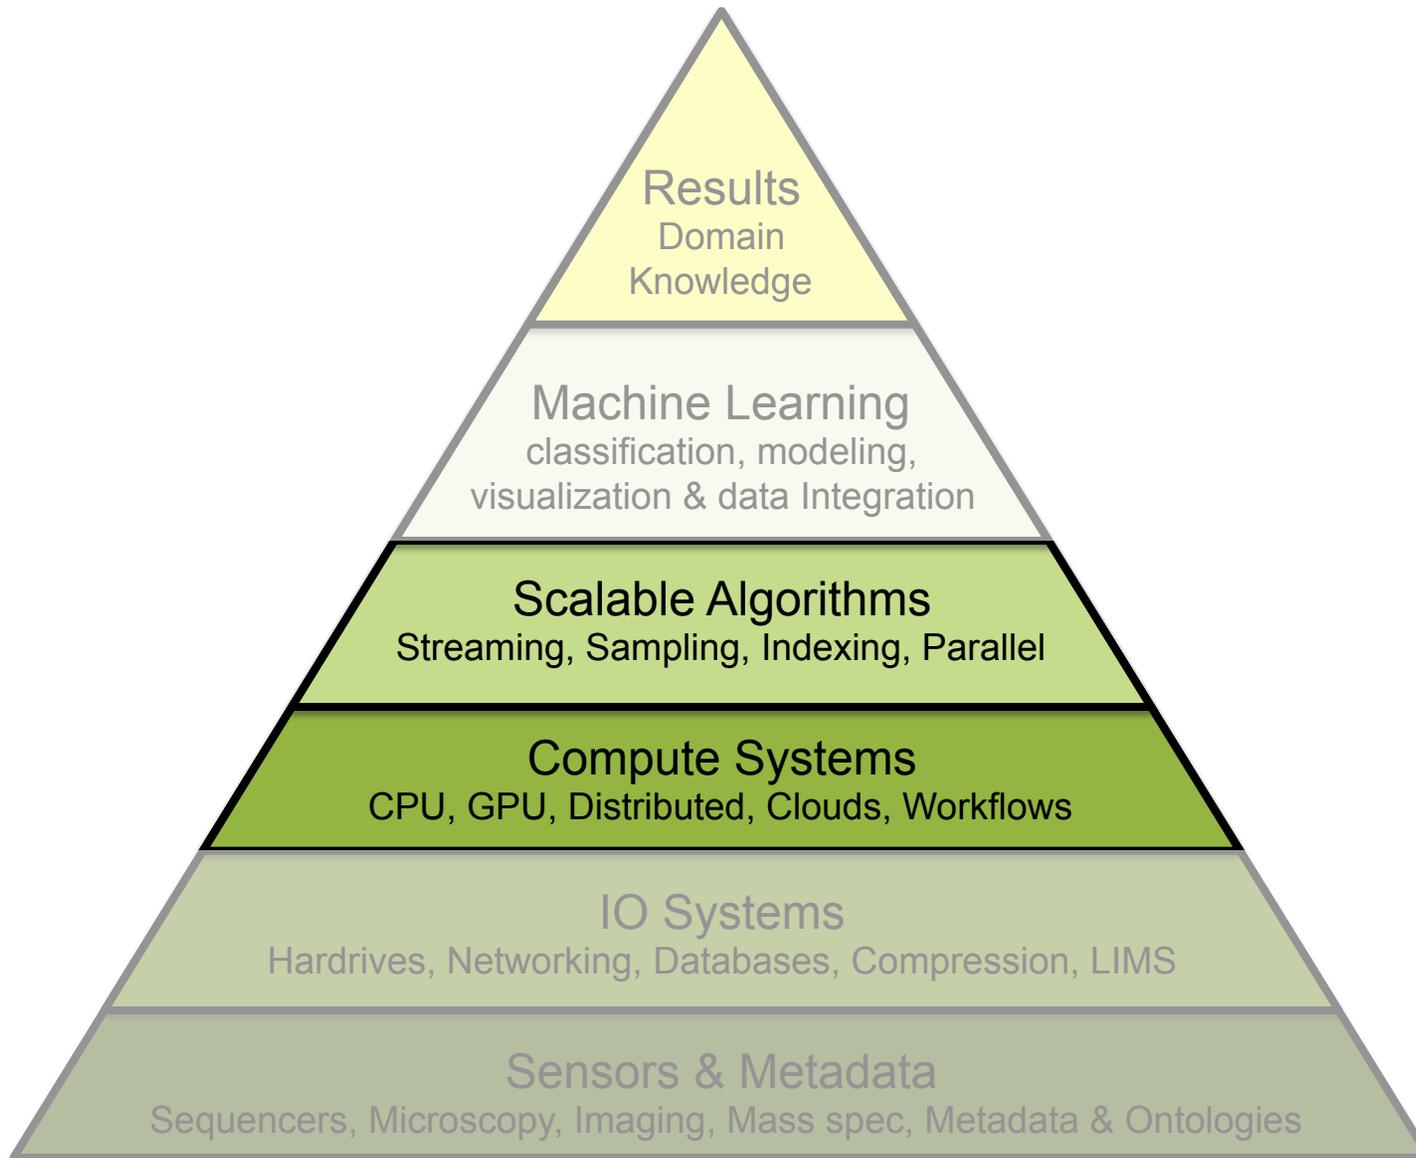
- Major efforts in human health and disease, agriculture, bioenergy

## But also widely distributed mobile sensors

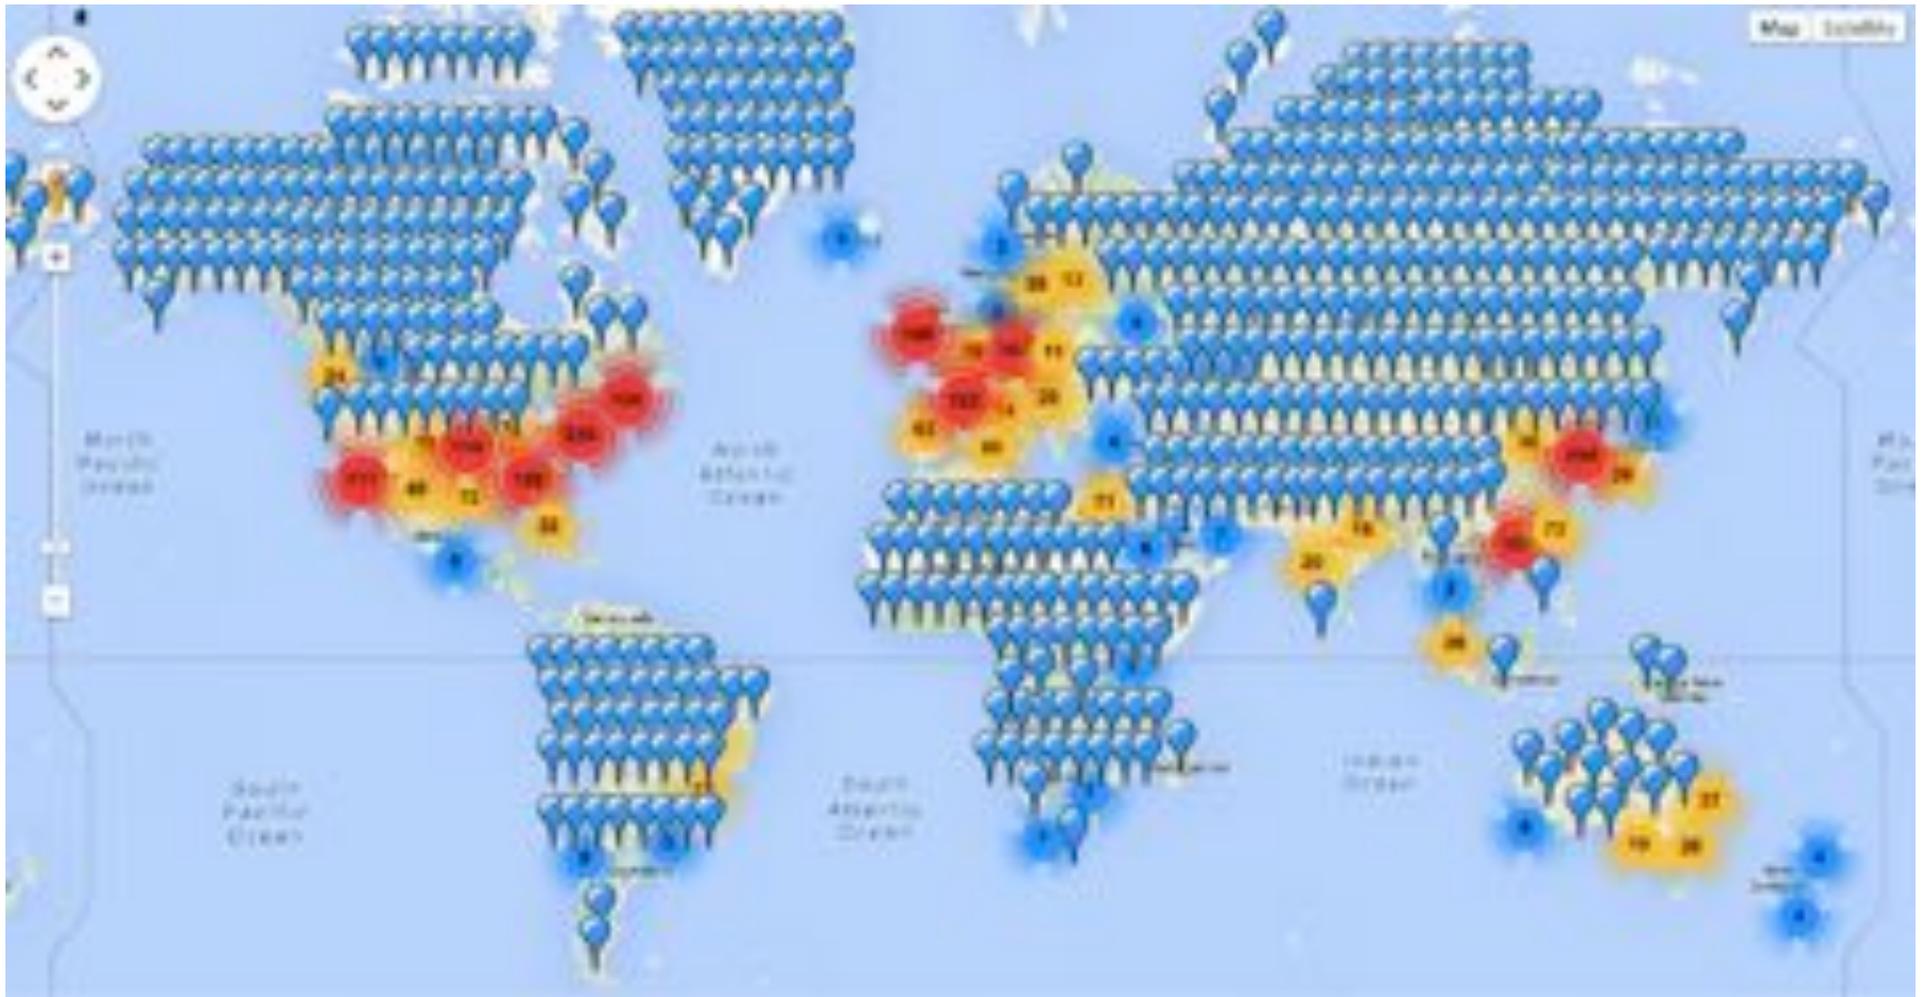
- Schools, offices, sports arenas, transportations centers, farms & food distribution centers
- Monitoring and surveillance, as ubiquitous as weather stations
- The rise of a digital immune system?



# Quantitative Biology Technologies



# Sequencing Centers



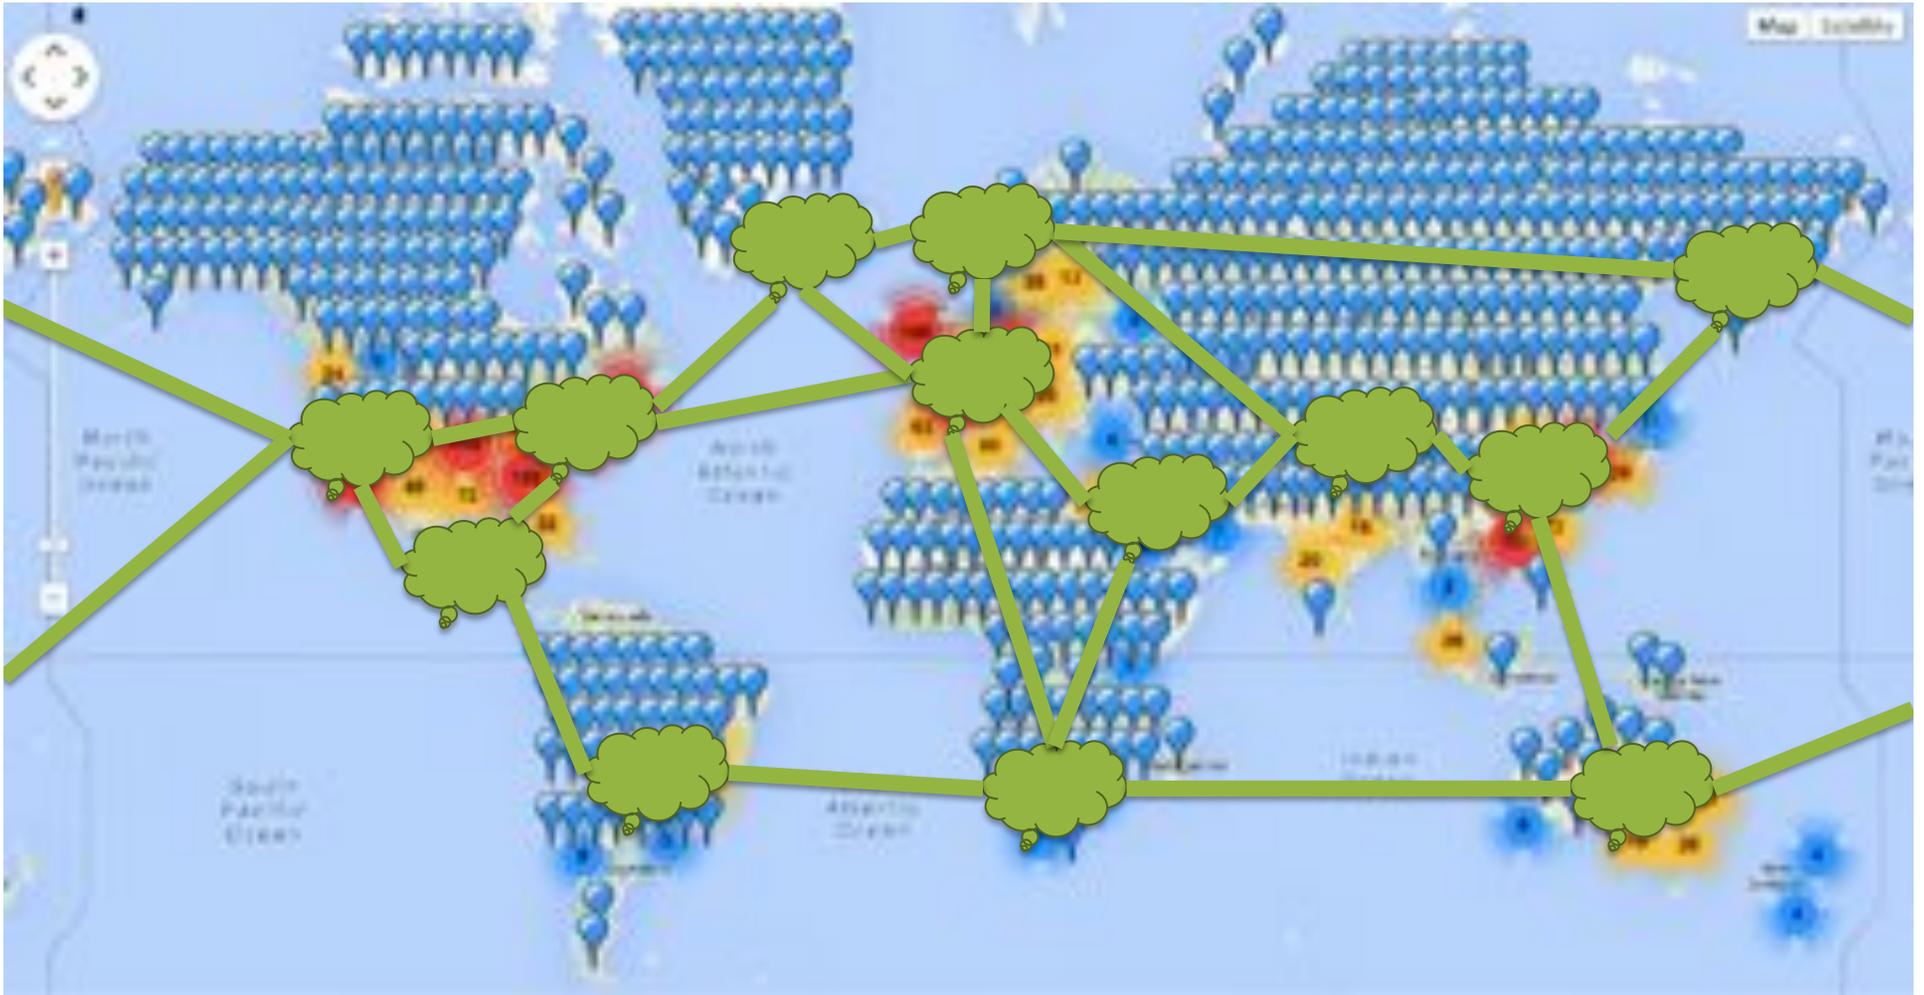
# Informatics Centers



## ***The DNA Data Deluge***

Schatz, MC and Langmead, B (2013) *IEEE Spectrum*. July, 2013

# Informatics Centers



## ***The DNA Data Deluge***

Schatz, MC and Langmead, B (2013) *IEEE Spectrum*. July, 2013

# Parallel Algorithm Spectrum

## Embarrassingly Parallel



Cluster Computing  
Each item is Independent

## Loosely Coupled



MapReduce  
Independent-Sync-Independent

## Tightly Coupled

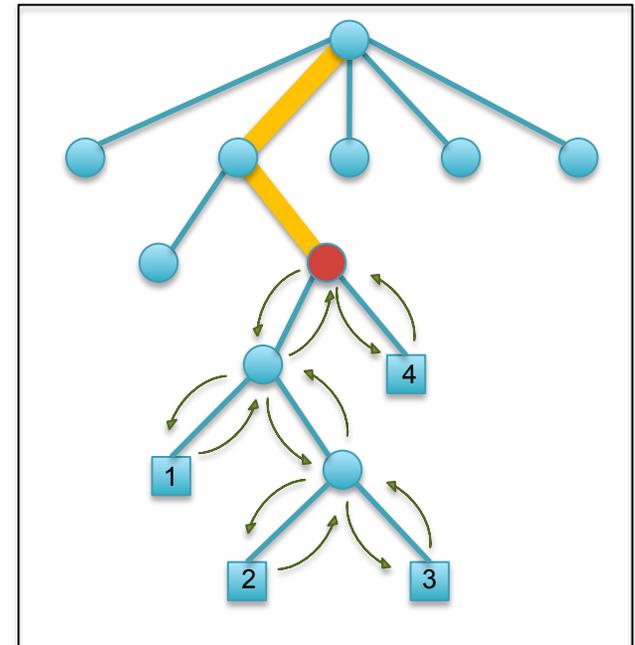


Graphs & MD simulations  
Constant Sync

# MUMmerGPU

<http://mummergpu.sourceforge.net>

- Index reference using a suffix tree
  - Each suffix represented by path from root
  - Reorder tree along space filling curve
- Map many reads simultaneously on GPU
  - Find matches by walking the tree
  - Find coordinates with depth first search
- Performance on nVidia GTX 8800
  - Match kernel was ~10x faster than CPU
  - Search kernel was ~4x faster than CPU
  - End-to-end runtime ~4x faster than CPU



- Cores are only part of the solution.
- Need storage, fast IO
- Locality is king

**High-throughput sequence alignment using Graphics Processing Units.**

Schatz, MC, Trapnell, C, Delcher, AL, Varshney, A. (2007) BMC Bioinformatics 8:474.

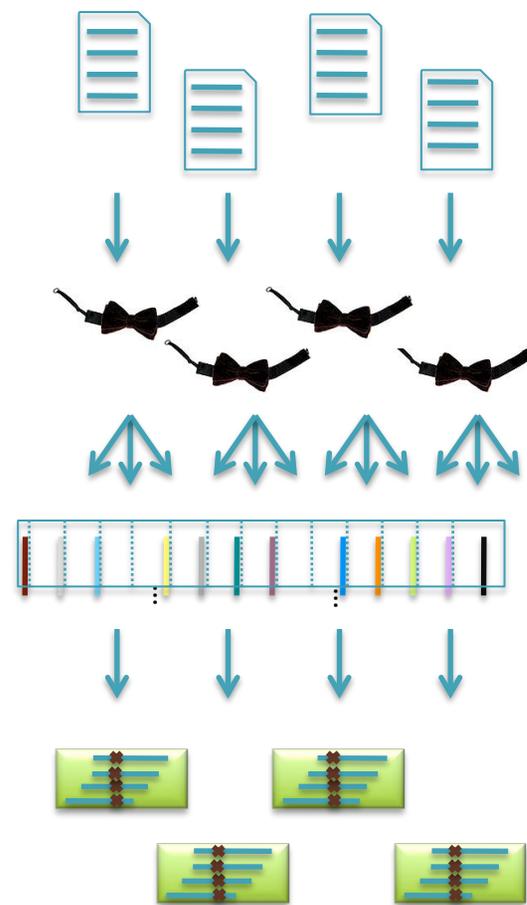


# Crossbow

<http://bowtie-bio.sourceforge.net/crossbow>

- Align billions of reads and find SNPs
  - Reuse software components: Hadoop Streaming
  - Mapping with Bowtie, SNP calling with SOAPsnp
- 4 hour end-to-end runtime including upload
  - Costs \$85; Today's costs <\$30

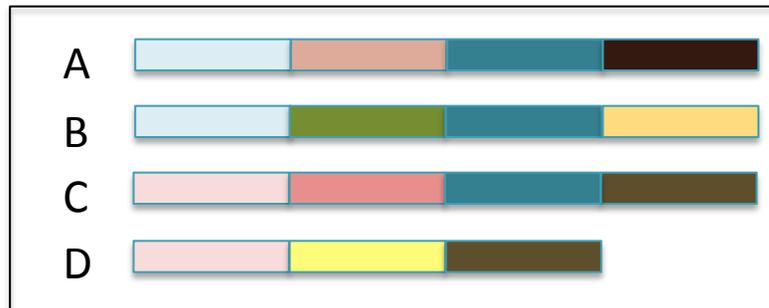
- Very compelling example of cloud computing in genomics
- Transfer takes time, but totally depends on institution
- Need more applications!



## Searching for SNPs with Cloud Computing.

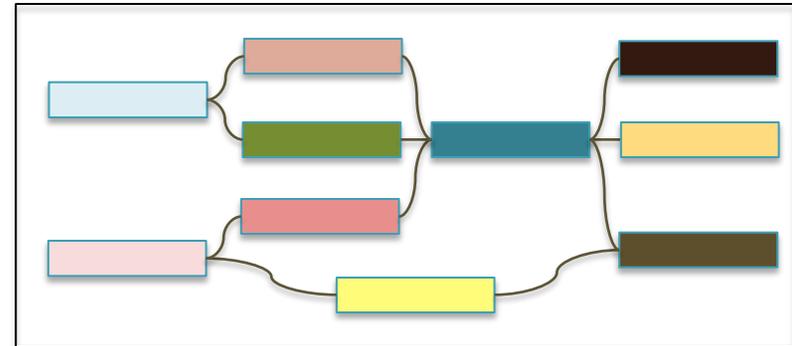
Langmead B, Schatz MC, Lin J, Pop M, Salzberg SL (2009) *Genome Biology*. **10**:R134

# Pan-Genome Alignment & Assembly



Time to start considering problems for which  $N$  complete genomes is the input to study the “pan-genome”

- Available today for many microbial species, near future for higher eukaryotes



Pan-genome colored de Bruijn graph

- Encodes all the sequence relationships between the genomes
- How well conserved is a given sequence?
- What are the pan-genome network properties?

**Rapid pan genome analysis with augmented suffix trees**

Marcus, S, Schatz, MC (2014) *In preparation*

# Compute & Algorithmic Challenges

**Expect to see many dozens of major informatics centers that consolidate regional / topical information**

- Clouds for Cancer, Autism, Heart Disease, etc
- Plus many smaller warehouses down to individuals
- Move the code to the data

**Parallel hardware and algorithms are required**

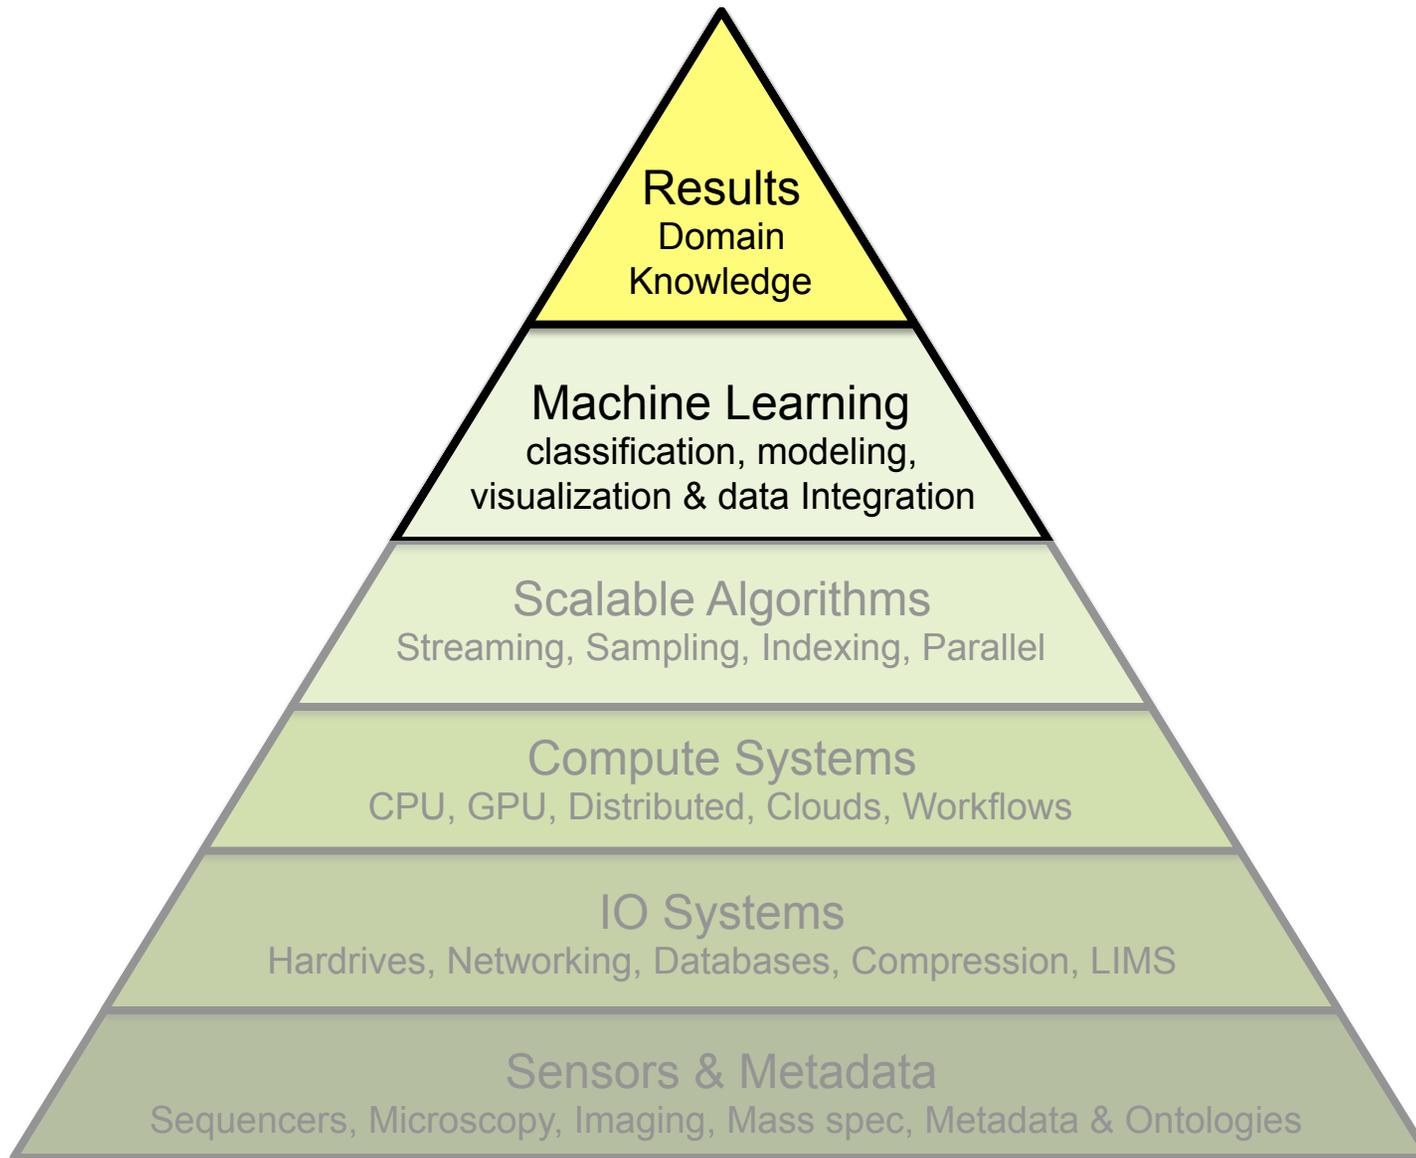
- Expect to see >1000 cores in a single computer
- Compute & IO needs to be considered together
- Rewriting efficient parallel software is complex and expensive

**Applications will shift from individuals to populations**

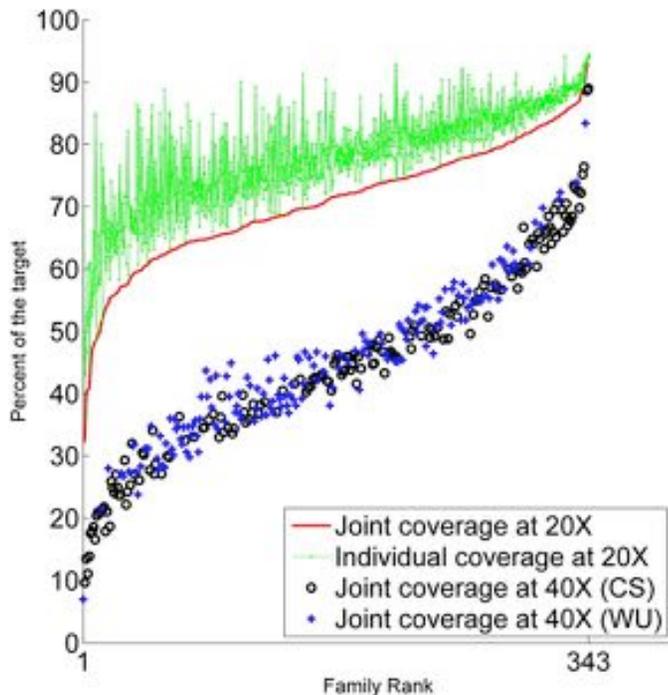
- Read mapping & assembly fade out
- Population analysis and time series analysis fade in
- Need for network analysis, probabilistic techniques



# Quantitative Biology Technologies



# Exome sequencing of the SSC



Last year saw 3 reports of >593 families from the Simons Simplex Collection

- Parents plus one child with autism and one non-autistic sibling
- All attempted to find “gene killing mutations” specific to the autistic children to find genes associated with the disease
- Iossifov (343) and O’Roak (50) used GATK, Sanders (200) didn’t attempt to identify indels

## De novo gene disruptions in children on the autism spectrum

Iossifov *et al.* (2012) *Neuron*. 74:2 285-299

## De novo mutations revealed by whole-exome sequencing are strongly associated with autism

Sanders *et al.* (2012) *Nature*. 485, 237–241.

## Sporadic autism exomes reveal a highly interconnected protein network of de novo mutations

O’Roak *et al.* (2012) *Nature*. 485, 246–250.

# Scalpel: Haplotype Microassembly

DNA sequence **micro-assembly** pipeline for accurate detection and validation of *de novo* mutations (SNPs, indels) within exome-capture data.



## Features

1. Combine **mapping** and **assembly**
2. Exhaustive search of **haplotypes**
3. **De novo mutations**



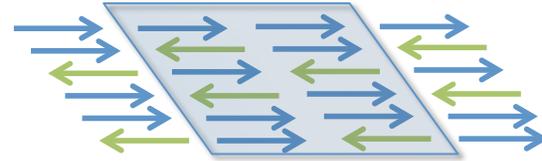
NRXN1 *de novo* SNP  
(auSSC12501 chr2:50724605)

**Accurate detection of *de novo* and transmitted INDELS within exome-capture data using micro-assembly**

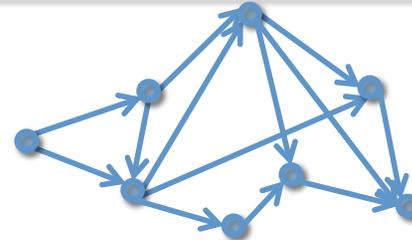
Narzisi, G, O'Rawe, J, Iossifov, I, Lee, Y, Wang, Z, Wu, Y, Lyon, G, Wigler, M, Schatz, MC (2014) *Under review.*

# Scalpel Pipeline

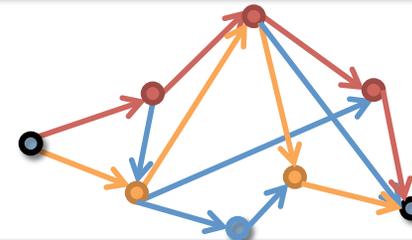
Extract reads mapping within the exon including (1) well-mapped reads, (2) soft-clipped reads, and (3) anchored pairs



Decompose reads into overlapping  $k$ -mers and construct de Bruijn graph from the reads



Find end-to-end haplotype paths spanning the region



Align assembled sequences to reference to detect mutations



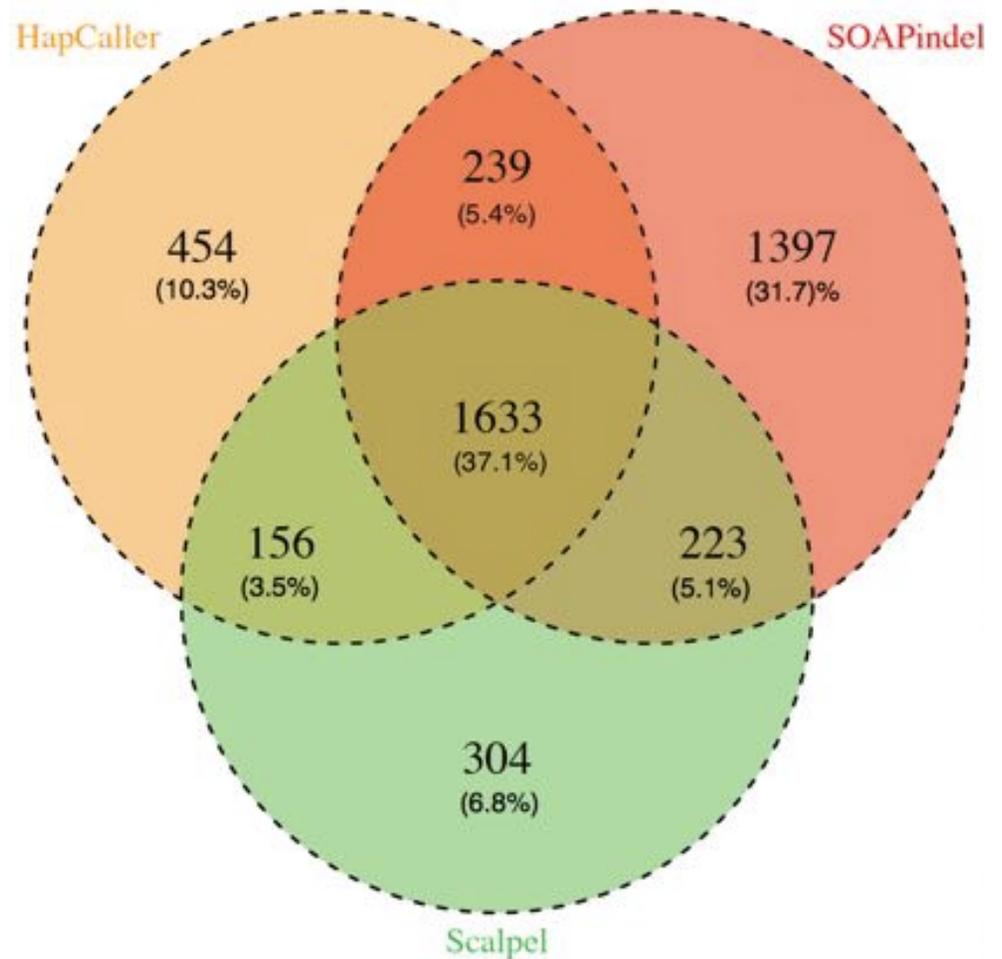
# Experimental Analysis & Validation

Selected one deep coverage exome for deep analysis

- Individual was diagnosed with ADHD
- 80% of the target at >20x coverage
- Evaluated with Scalpel, SOAPindel, and GATK Haplotype Caller

1000 indels selected for validation

- 200 Scalpel
- 200 GATK Haplotype Caller
- 200 SOAPindel
- 200 within the intersection
- 200 long indels (>30bp)



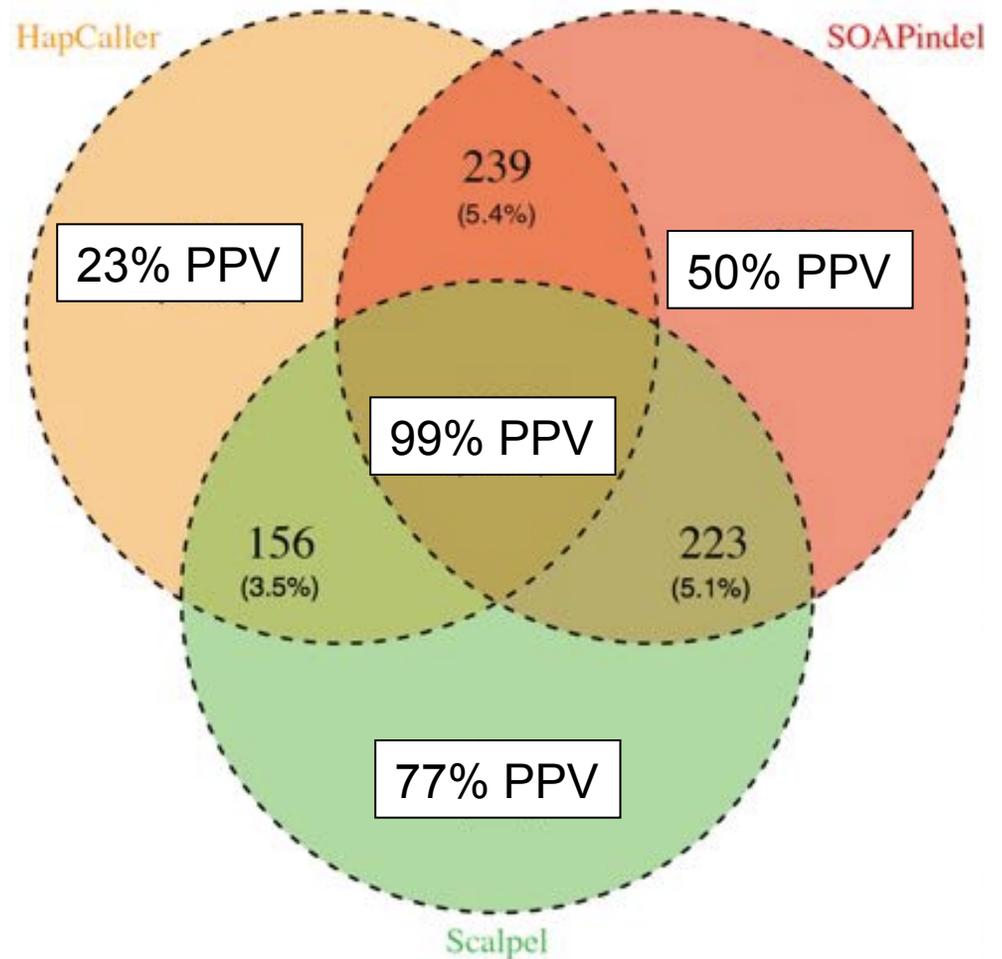
# Experimental Analysis & Validation

Selected one deep coverage exome for deep analysis

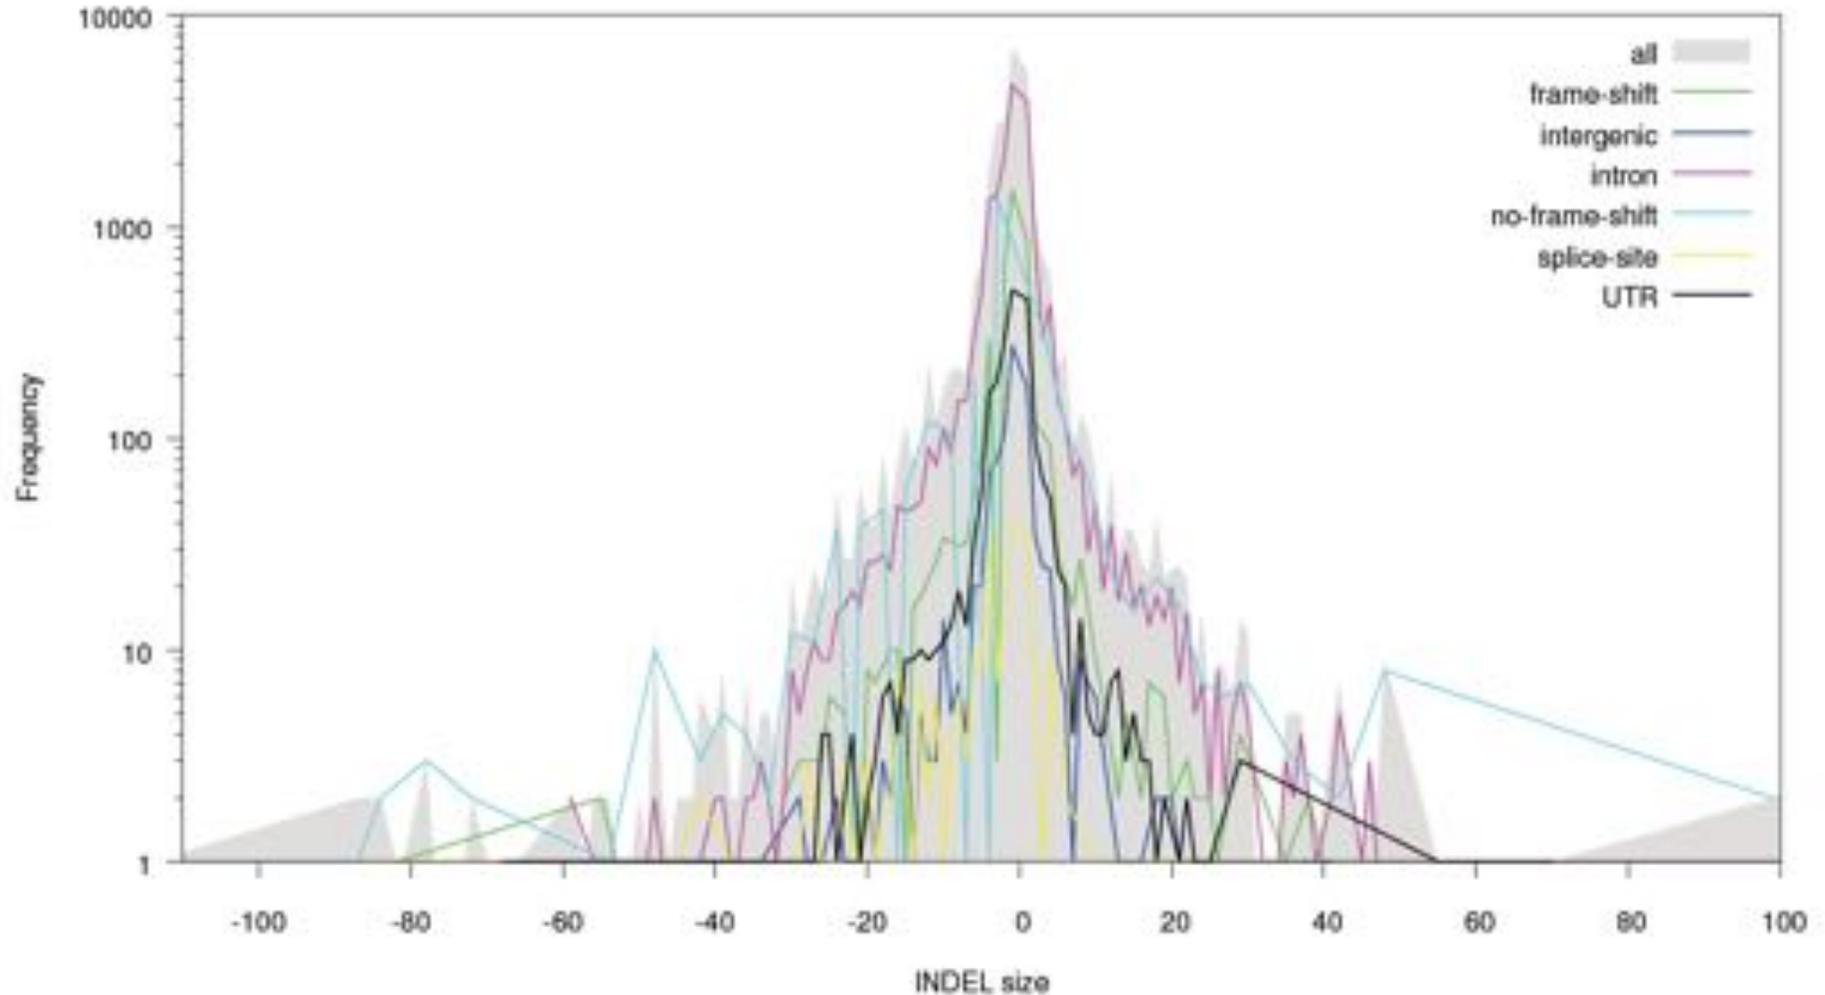
- Individual was diagnosed with ADHD (See Gholson for details)
- 80% of the target at >20x coverage
- Evaluated with Scalpel, SOAPindel, and GATK Haplotype Caller

1000 indels selected for validation

- 200 Scalpel
- 200 GATK Haplotype Caller
- 200 SOAPindel
- 200 within the intersection
- 200 long indels (>30bp)



# Revised Analysis of the SSC

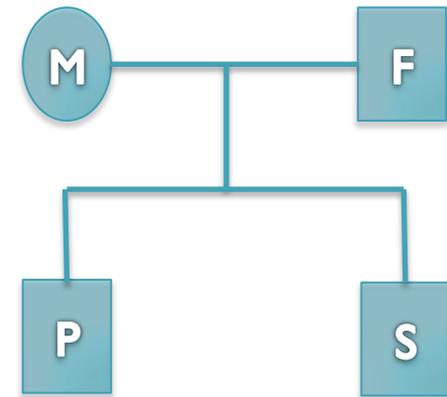


Constructed database of  $>1M$  transmitted and de novo indels  
Many new gene candidates identified, population analysis underway

# De novo mutation discovery and validation

**Concept:** Identify mutations not present in parents.

**Challenge:** Sequencing errors in the child or low coverage in parents lead to false positive de novos



**Reference:** . . . TCAAATCCTTTTAAATAAGAAGAGCTGACA . . .

**Father:** . . . TCAAATCCTTTTAAATAAGAAGAGCTGACA . . .

**Mother:** . . . TCAAATCCTTTTAAATAAGAAGAGCTGACA . . .

**Sibling:** . . . TCAAATCCTTTTAAATAAGAAGAGCTGACA . . .

**Proband(1):** . . . TCAAATCCTTTTAAATAAGAAGAGCTGACA . . .

**Proband(2):** . . . TCAAATCCTTTTAAAT\*\*\*\*AAGAGCTGACA . . .

4bp heterozygous deletion at chr15:9352406 | CHD2

# De novo Genetics of Autism

- In 593 family quads so far, we see significant enrichment in de novo *likely gene killers* in the autistic kids
  - Overall rate basically 1:1
  - 2:1 enrichment in nonsense mutations
  - 2:1 enrichment in frameshift indels
  - 4:1 enrichment in splice-site mutations
  - Most de novo originate in the paternal line in an age-dependent manner (56:18 of the mutations that we could determine)
- Observe strong overlap with the 842 genes known to be associated with fragile X protein FMR1
  - Related to neuron development and synaptic plasticity
  - Also strong overlap with chromatin remodelers



# The fallacy of big data?

810 DATA

## The Parable of Google Flu: Traps in Big Data Analysis

David Galet,\*\*\* Ryan Kennedy,\*\*\* Gary King,\*\*\* Alessandro Vespignani\*\*\*

In February 2013, Google Flu Trends (GFT) made headlines, but not for a reason that Google executives or the creators of the flu tracking system would have hoped. Alvarez reported that GFT was predicting more than double the proportion of doctor visits for influenza-like illness (ILI) than the Centers for Disease Control and Prevention (CDC), which bases its estimates on surveillance reports from laboratories across the United States (1, 2). This happened despite the fact that GFT was built to predict CDC reports. Given the GFT's often held up as an exemplary use of big data (3, 4), what lessons can we draw from this event?

The problems we identify are not limited to GFT. Research on whether search or social media can predict a flu becomes commonplace (5–7) and is often put in direct tension with traditional methods and hypotheses. Although these studies have shown the value of these data, we are far from a place where they can supplant more traditional methods or theories (8). We explore two issues that contributed to GFT's misadventure: data hubris and algorithm dynamics—and offer lessons for moving forward in the big data era.

### Big Data Hubris

“Big data hubris” is the often implicit assumption that big data are a substitute for, rather than a supplement to, traditional data collection and analysis. Elsewhere, we have asserted that there are enormous scientific possibilities in big data (9–11). However, quantity of data does not mean that one can ignore foundational issues of measurement and construct validity and reliability.



Large errors in flu predictions were largely avoidable, which offers lessons for the use of big data.

## Big Data Hubris

“Big data hubris” is the often implicit assumption that big data are a substitute for, rather than a supplement to, traditional data collection and analysis. Elsewhere, we have asserted that there are enormous scientific possibilities in big data (9–11). However, quantity of data does not mean that one can ignore foundational issues of measurement and construct validity and reliability.

ability and dependencies among data. The core change is that most big data are now being generated by a small number of individuals or organizations, which means that the data are often not representative of the population as a whole.

The initial version of GFT was based on a problematic marriage of search data. Essentially, the metric was the number of search terms that were related to the flu. The data of finding words to search for is not necessarily related to the flu, but a correlation was found. GFT developers, report waiting out seasons were compared to the flu but strong link to the CDC data, such as those that high school basketball (12).

There has been a warning that big data were overhyped. The number of cases—a common metric in data analysis. This is a method of showing out positive search terms failed when GFT completely missed the nonseasonal 2009 influenza A (H1N1) pandemic (13, 14). In short, the initial version of GFT was partly detection, partly winter detection. GFT engineers updated the algorithm in 2009, and this method has

being GFT work with new real-time health data (15, 16). For example, by combining GFT and lagged CDC data, as well as dynamically recalibrating GFT, we can substantially improve on the performance of GFT or the CDC alone (see the eText). This is no substitute for ongoing evaluation and improvement, but, by incorporating this information, GFT could have largely fixed itself and would have likely remained out of the headlines.

1. Alvarez, D. L. Google Flu Trends overestimates flu cases. *Nature* 494, 454–455 (2013).  
2. Alvarez, D. L. Google Flu Trends overestimates flu cases. *Nature* 494, 454–455 (2013).  
3. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
4. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
5. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
6. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
7. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
8. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
9. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
10. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
11. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
12. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
13. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
14. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
15. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).  
16. Anderson, A. C. et al. Google Flu Trends: A real-time public health surveillance system. *PLoS Med.* 9, e1001188 (2012).

# The risks of big data?

## Predicting Social Security numbers from public data

Alessandro Acquisti<sup>1</sup> and Ralph Gross

Carnegie Mellon University, Pittsburgh, PA 15213

Communicated by Stephen E. Fienberg, Carnegie Mellon University, Pittsburgh, PA, May 5, 2009 (received for review January 18, 2009)

Information about an individual's place and date of birth can be exploited to predict his or her Social Security number (SSN). Using only publicly available information, we observed a correlation between individuals' SSNs and their birth data and found that for younger cohorts the correlation allows statistical inference of private SSNs. The inferences are made possible by the public availability of the Social Security Administration's Death Master

File and the widespread accessibility of persons from multiple sources, such as data brokers or professional working sites. Our results highlight the unexpected consequences of the complex interactions among data sources in modern information economies and the risks associated with information revelation in

identity theft | online social networks | privacy | statistics

In modern information economies, sensitive personal information is often in plain sight amid transactions that rely on their unhindered circulation. Such is the case with Social Security numbers in the United States: Created as identifiers for tracking individual earnings (1), they have become authentication devices (2), becoming one of the most sought after pieces of information most often sought by identity thieves. The Social Security Administration (SSA), which issues them, has tried to keep SSNs confidential (3), coordinating with law enforcement to limit their public exposure (4).<sup>\*</sup> After embarrassing incidents in the private sector entities also have attempted to strengthen their consumers' and employees' data (7).<sup>†</sup> How have these efforts already left the harm: We demonstrate that

number (SN). The SSA openly provides information about the process through which ANs, GNs, and SNs are issued (1). ANs are currently assigned based on the zipcode of the mailing address provided in the SSN application form [RM00201.030] (1). Low-population states and certain U.S. possessions are allocated 1 AN each, whereas other states are allocated sets of ANs (for instance, an individual applying from a zipcode within

publish on social networking sites (10). Using this method, we identified with a single attempt the first 5 digits for 44% of DMF records of deceased individuals born in the U.S. from 1989 to 2003 and the complete SSNs with <1,000 attempts (making SSNs akin to 3-digit financial PINs) for 8.5% of those records. Extrapolating to the U.S. living population, this would imply the potential identification of millions of SSNs for individuals whose birth data were available. Such findings highlight the hidden privacy costs of widespread information dissemination and the complex interactions among multiple data sources in modern information economies (11), underscoring the role of public records as breeder documents (12) of more sensitive data.

**Keywords:**

SEE COMMENTARY

# Learning and Translation

## **Tremendous power from data aggregation**

- Observe the dynamics of biological systems
- Breakthroughs in medicine and biology of profound significance

## **Be mindful of the risks**

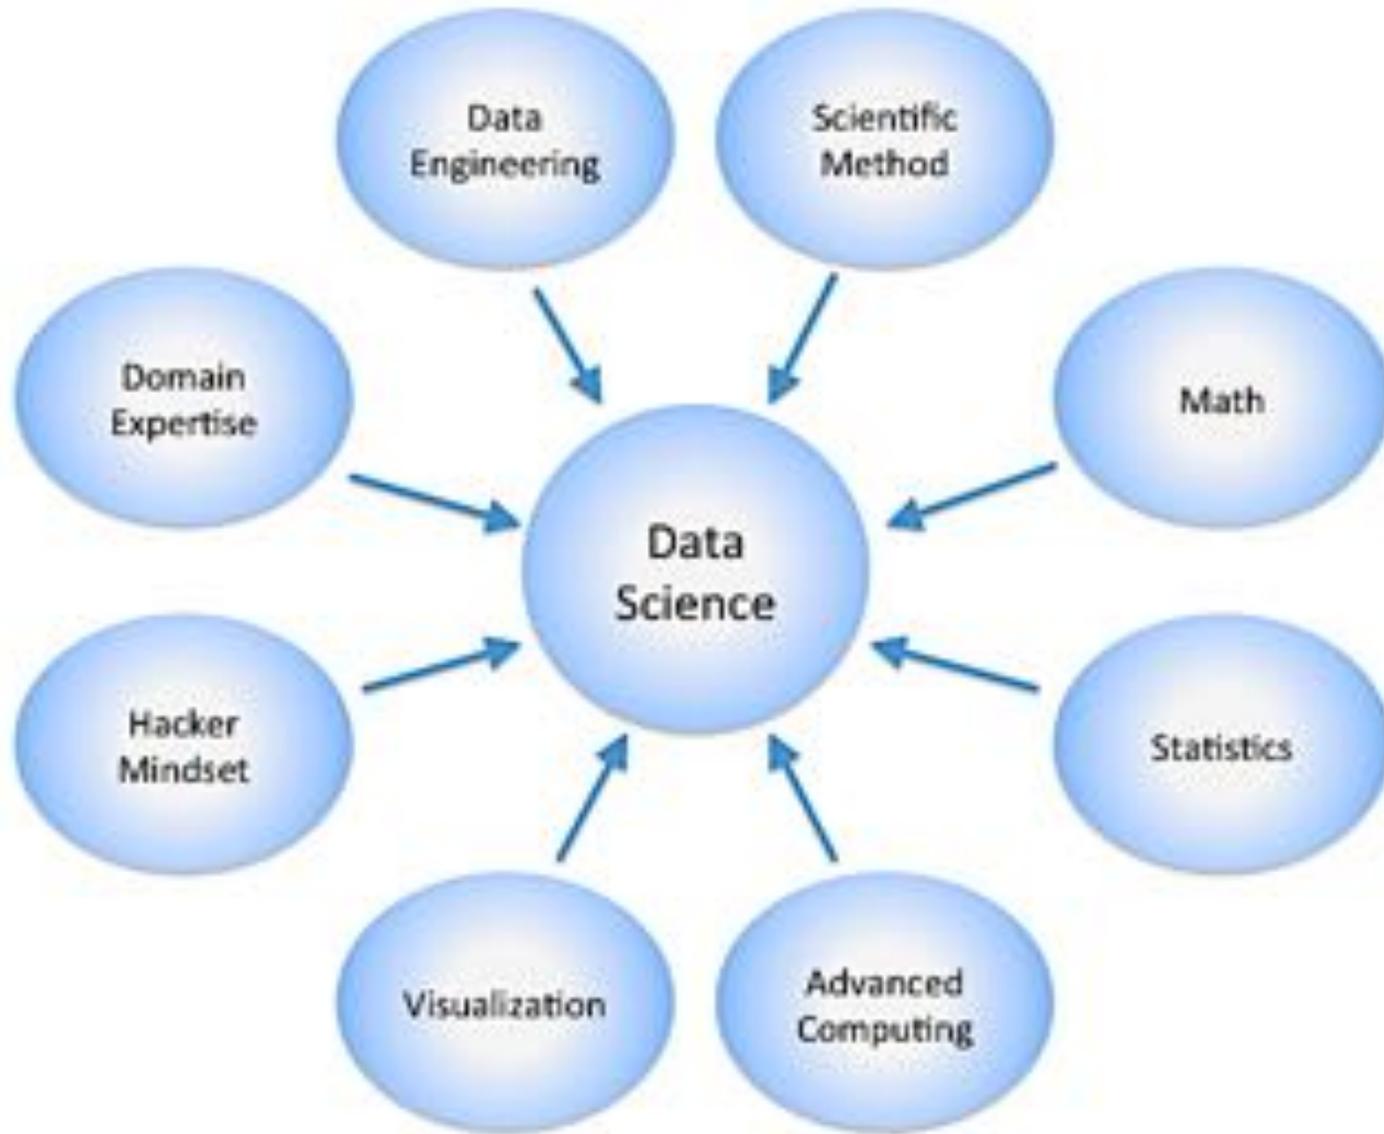
- The potential for over-fitting grows with the complexity of the data, statistical significance is a statement about the sample size
- Reproducible workflows, APIs are a must
- Caution is prudent for personal data

## **The foundations of biology will continue to be observation, experimentation, and interpretation**

- Technology will continue to push the frontier
- Feedback loop from the results of one project into experimental design for the next



# Who is a Data Scientist?



[http://en.wikipedia.org/wiki/Data\\_science](http://en.wikipedia.org/wiki/Data_science)

# Acknowledgements

Schatz Lab

**Giuseppe Narzisi**

Shoshana Marcus

James Gurtowski

Srividya

Ramakrishnan

Hayan Lee

Rob Aboukhalil

Mitch Bekritsky

Charles Underwood

Tyler Gavin

Alejandro Wences

Greg Vurture

Eric Biggers

Aspyn Palatnick

CSHL

Hannon Lab

Gingeras Lab

Jackson Lab

Iossifov Lab

Levy Lab

Lippman Lab

Lyon Lab

Martienssen Lab

McCombie Lab

Tuveson Lab

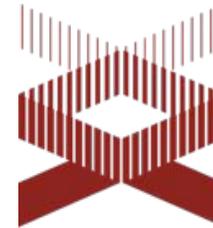
Ware Lab

Wigler Lab

IT Department

**SFARI**

SIMONS FOUNDATION  
AUTISM RESEARCH INITIATIVE



National Human  
Genome Research  
Institute



U.S. DEPARTMENT OF  
**ENERGY**



# ***Biological Data Sciences***

Cold Spring Harbor Laboratory, Nov 5 - 8, 2014

Michael Schatz, Anne Carpenter, Matt Wood



# Thank you

<http://schatzlab.cshl.edu>

@mike\_schatz / #KSBigData