




Research Faculty Summit 2018

Systems | Fueling future disruptions




The Good, the Bad, and the Ugly of ML for Networked Systems

Keith Winstein

Assistant Professor of Computer Science

Assistant Professor of Electrical Engineering

(by courtesy) Stanford University



The Good, the Bad, and the Ugly of ML for Networked Systems

Keith Winstein

Assistant Professor of Computer Science
Assistant Professor of Electrical Engineering (by courtesy)
Stanford University




Three uses of machine learning

- ▶ **Paradigm 1** [learn then deploy]:
“ML produces an artifact. People deploy it in real life.”
- ▶ **Paradigm 2:** [deploy and learn]:
“People deploy an artifact that learns in real life.”
- ▶ **Paradigm 3:** [learn from the machines]:
“ML teaches us about our own thinking.”

- ▶ **Paradigm 1** [learn then deploy]:
Often harder than we expect.
(“Past performance is no guarantee of future results.”)
- ▶ **Paradigm 2:** [deploy and learn]:
Worth researching, but hard because of the nature of networks.
- ▶ **Paradigm 3:** [learn from the machines]:
Old-fashioned AI view—still valuable!

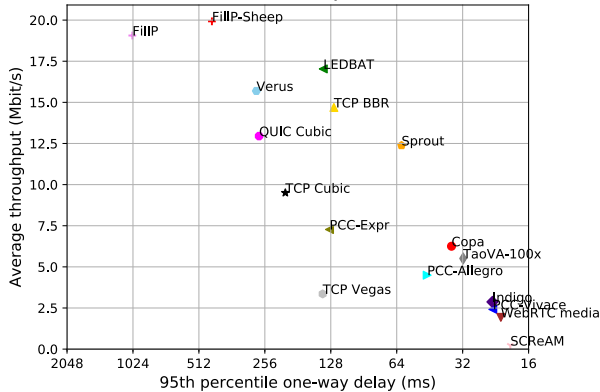
Paradigm 1: Sprout (NSDI 2013) in publication



Sprout, NSDI 2013 (figure 7)


Paradigm 1: Sprout in real life in America

test from AWS California 1 to Stanford ppp0, 3 runs of 30s each per scheme
(mean of all runs by scheme)



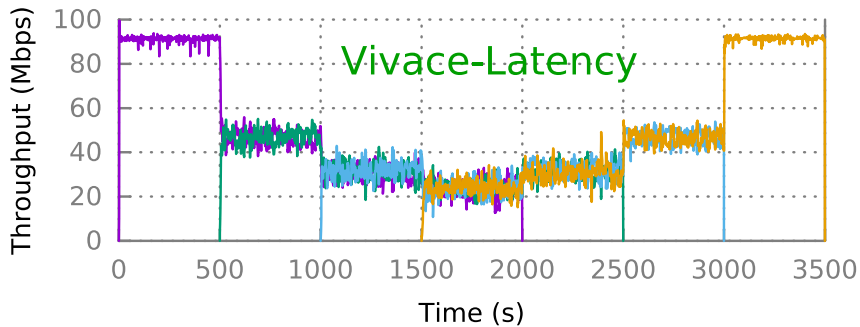
Stanford Pantheon result (July 31, 2018, T-Mobile in California),
<https://pantheon.stanford.edu/result/3455/>

Paradigm 1: Sprout in real life in India




Stanford Pantheon result (August 1, 2018, Airtel in New Delhi),
<https://pantheon.stanford.edu/result/3474/>

Paradigm 1: Vivace (NSDI 2018) in publication



Vivace, NSDI 2018 (figure 7)

Paradigm 1: Vivace in real life



Stanford Pantheon result (August 1, 2018, AWS Brazil-HostDime Colombia),
<https://pantheon.stanford.edu/result/3470/>

Paradigm 1: Pensieve (SIGCOMM 2017) in publication




Figure 11: Comparing Pensieve with existing ABR algorithms in the wild. Results are for the QoE_{lin} metric and were collected on the Verizon LTE cellular network, a public WiFi network, and the wide area network between Shanghai and Boston. Bars list averages and error bars span \pm one standard deviation from the average.

Paradigm 1: Pensieve in reproduction




Figure 1: Comparison of Pensieve with other ABR algorithms across 10 tests on real world networks

Stanford CS244 student project,

https://reproducingnetworkresearch.files.wordpress.com/2018/07/recreating_pensieve.pdf


Paradigm 1: BBR (ACM Queue 2016) in publication

“**BBR converges toward a fair share** of the bottleneck bandwidth whether competing with other BBR flows or with loss-based congestion control. [...] Unmanaged router buffers exceeding several BDPs, however, cause long-lived loss-based competitors to bloat the queue and grab more than their fair share.”

<https://queue.acm.org/detail.cfm?id=3022184>

Paradigm 1: BBR in independent evaluation

Cubic vs BBR over a 12ms RTT 10G circuit



Geoff Huston, *TCP and BBR*, RIPE 76 (May 2018)

<https://ripe76.ripe.net/presentations/10-2018-05-15-bbr.pdf>

Paradigm 1: BBR in independent evaluation

Why **not** use BBR?

- Because it **over achieves!**
- The classic question for many Internet technologies is scaling
 - “what if everyone does it?”
 - BBR is not a scalable approach
 - It works so well while it is used by just a few users, some of the time
 - But when it is active, BBR has the ability to slaughter concurrent loss-based flows
 - Which sends all the wrong signals to the TCP ecosystem
 - The loss-based flows convert to BBR to compete on equal terms
 - The network is then a BBR vs BBR environment, which is unstable

Geoff Huston, TCP and BBR, RIPE 76 (May 2018)

<https://ripe76.ripe.net/presentations/10-2018-05-15-bbr.pdf>

Paradigm 1: BBR in independent evaluation

Is this BBR experiment a failure?

Is it just too 'greedy' and too 'insensitive' to other flows to be allowed out on the Internet to play?

- Many networks have been provisioned as a response to the aggregate behaviours of loss-based TCP congestion control
- BBR changes all those assumptions, and could potentially push many networks into sustained instability
- We cannot use the conventional network control mechanisms to regulate BBR flows
 - Selective packet drop just won't create back pressure on the flow

Geoff Huston, TCP and BBR, RIPE 76 (May 2018)

<https://ripe76.ripe.net/presentations/10-2018-05-15-bbr.pdf>

Paradigm 1: BBR in independent evaluation

Where now?

BBR 2.0

- Alter BBR's 'sensitivity' to loss rates, so that it does not persist with an internal bandwidth delay product (BDP) that exceeds the uncongested BDP
 - This measure would moderate BBR 1.0's ability to operate for extended periods with very high loss levels
- Improve the dynamic sharing fairness by moderating the Bandwidth Delay Product by using an estimated 'fair' proportion of the path BDP
- Accommodate the signal distortion caused by ACK stretching middleware
- Place an upper bound on the volume of in-flight data
- Alter the +/- 25% probe factors dynamically (i.e. allow this to be less than 25% overload)

Geoff Huston, TCP and BBR, RIPE 76 (May 2018)

<https://ripe76.ripe.net/presentations/10-2018-05-15-bbr.pdf>

Paradigm 1: Google Flu Trends

Proposal (2008): **train** a model to predict flu incidence from historical search engine queries. **Then deploy** the model to predict flu in advance of the government.

Nov. 11, 2008 announcement

"All the News That's Fit to Print"

The New York Times

VOL. CLVIII . . . No. 54,492

©2008 The New York Times

NEW YORK, WEDNESDAY, NOVEMBER 12, 2008

\$1.50



DAY OF HONOR (AND SNACKS) Members of the military were among some 20,000 people in Manhattan for the Veterans Day parade.

ACHES, A SNEEZE, A GOOGLE SEARCH

Data on Web May Warn of Outbreaks of Flu

By MIGNOUL HELPT
SAN FRANCISCO — There is a new common symptom of the flu, in addition to the usual aches, coughs, fevers and sore throats. "You're not a lot of other Ameri-

Veterans' Families Seek Aid for Caregiver Role

By LESLIE KAUFMAN

They Kell met her husband, Matt, in August 2003 between his first and second tours of duty in Iraq. They married in January 2005. Six weeks later, Staff Sergeant Kell was shot in the neck while on patrol in Ramadi, Iraq, and remained a quadriplegic.

Because her husband, now 27, could no longer take care of himself, she even to get a drink of water. Mrs. Kell, 31, quit her job as an accountant to take care of him.

through four workers in two months she gave up. She said many of the caregivers from contractors on the government-provided list "were awful." One did not know how to use the lift system that hoists Mr. Kell out of bed; another gossiped about the family's private business.

But the real problem was that even the good caregivers could not help Mr. Kell live as he wanted. Regulations, for example, do not permit them to take him out of the house. "Matt is back to his

doctor's, went just a baby stiter." While she has never regretted leaving her job, the financial repercussions have been serious. Although Mr. Kell gets a full disability pension of \$6,800 a month and their home in Parker, Colo., was donated to them, they have lost Mrs. Kell's salary of \$58,000 a year, as well as employer contributions to her retirement account, and her dental plan.

Mrs. Kell has joined a growing group of veterans' families who are asking to be compensated in place of a caregiver. She sees it

Buying Binge Slams to Halt

Crisis of Confidence For U.S. Consumers

Just an one crisis of confidence may be ending, another may be coming.

The panic on Wall Street has cooled in the last few weeks, and banks have become somewhat more willing to make loans. But in these

DAVID LEONHART
ECONOMIST
American households appear to have fallen into their own defensive crouch.

Suddenly, our consumer society is doing a lot less consuming. The numbers are pretty intractable. Sales of new vehicles have dropped 22 percent in the third quarter. Consumer spending appears likely to fall next year for the first time since 1982 and perhaps by the largest amount since 1942.

With Wall Street edging back from the brink, this crisis of consumer confidence has become the No. 1 short-term issue for the economy. Nobody doubts that families need to start saving more than they saved over the last two decades. But if they change their behavior too quickly, it could be very painful.

Already, Circuit City has filed for bankruptcy, and General Motors has said that it's in danger of running out of cash. If the consumer slump continues, there is a potential for a dangerous feedback loop, in which spending cuts and layoffs reinforce each other.

It's a scary time. "Liz Allen, 20, a nursing student in Atlanta, told one of The Times reporters who lunched out across the country last weekend to ask people about the economy. "They can make the economy worse. If people worry too much they won't

Late Edition
Today, dispersed sunshine, high 51. Tonight, thickening clouds, low 43. Tomorrow, cloudy, a few scattered showers, some heavy late, high 54. Weather map is on Page A18.

DEMOCRATS SEEK EMERGENCY HELP FOR AUTOMAKERS

CALL FOR AID PACKAGE

Leaders May Try to Use Lame-Duck Session to Press Bush

By DAVID M. HERSHENHORN and CARL HULSE

WASHINGTON — Democratic Congressional leaders said Tuesday that they were ready to push emergency legislation to aid the imperiled auto industry when lawmakers return to Washington next week for the first time after the election, setting the stage for one last showdown with President Bush.

"Next week, during the lame-duck session of Congress, we are determined to pass legislation that will save the jobs of millions of workers whose livelihoods are on the line," the majority leader, Harry Reid of Nevada, said in a statement.

His call for the session came shortly after the House speaker, Nancy Pelosi, said Congress and the administration "must take immediate action" to avert off a possible collapse of the American auto industry.

Mrs. Pelosi stopped short of saying Congress would adopt legislation to provide emergency financial aid to the automakers, giving the Treasury Department the option of using money from the \$300 billion bailout program instead.

But with the White House insisting that the bailout money be reserved for financial institutions, that option seemed unlikely, leading a senior Democratic official to say Democrats would

Google Flu Trends plot as of today

Historical estimates


See data for:

United States

United States Flu Activity

Influenza estimate

● Google Flu Trends estimate ● United States data



United States: Influenza-like illness (ILI) data provided publicly by the [U.S. Centers for Disease Control](http://www.cdc.gov).

(<http://www.google.org/flutrends/about/how.html>)



Most of plot is training data

Historical estimates


See data for:

United States

United States Flu Activity


Influenza estimate

● Google Flu Trends estimate ● United States data



United States: Influenza-like illness (ILI) data provided publicly by the [U.S. Centers for Disease Control](#).

Large divergence (3.7×) in New England (HHS region 1)



Paradigm 1: spam filtering

SpamAssassin (spam filtering engine):

- ▶ Anybody can propose a spam-filtering algorithm.
- ▶ Central party **learns** best weights based on predictive power of each algorithm.
- ▶ Weights are **then deployed** in the field.

Paradigm 1: spam filtering

- ▶ **2007:** a rule is added
- ▶ Rule: “Does the year match 200x?”
- ▶ Catches a lot of spam.
- ▶ **Extremely low false-positive rate!**
- ▶ *(but... big surprise on 1/1/2010)*

The lesson? (paradigm 1)

Learn-then-deploy is a challenging pattern, and empirically it's easy to fool ourselves into premature declarations of success.

Paradigm 2: deploy **and** learn

- ▶ **Proposal:** Build systems that learn continuously, online.
- ▶ Directly observe operational figure-of-merit over time.
- ▶ React quickly to real-world changes.

QUIC, SIGCOMM 2017 (figure 6)




Figure 6: Search Latency reduction for users in the QUIC experiment over an 18-month period. Numbered events are described in Section 5.2.

Real systems learn over time: baby robot, car, etc.

Challenges of *network* learning

But classical results indicate challenges to learning when

- ▶ ... information is distributed. Dec-POMDP is undecidable.
- ▶ ... compute and data are in different places.
- ▶ ... agents are adversarial (congestion control, routing, traffic engineering, security).

All of these scenarios are characteristic of **networks**.

What if information is distributed?

ECONOMETRICA

VOLUME 47

NOVEMBER, 1979

NUMBER 6

THE IMPOSSIBILITY OF BAYESIAN GROUP DECISION MAKING WITH SEPARATE AGGREGATION OF BELIEFS AND VALUES

BY AANUND HYLLAND AND RICHARD ZECKHAUSER¹

Bayesian theory for rational individual decision making under uncertainty prescribes that the decision maker define independently a set of beliefs (probability assessments for the states of the world) and a system of values (utilities). The decision is then made by maximizing expected utility. We attempt to generalize the model to group decision making. It is assumed that the group's belief depends only on individual beliefs and the group's values only on individual values, that the belief aggregation procedure respects unanimity, and that the entire procedure guarantees Pareto optimality. We prove that only trivial (dictatorial) aggregation procedures for beliefs are possible.

1. INTRODUCTION


MANY DECISIONS MADE under uncertainty, indeed many important ones, are

What if the scenario is adversarial?

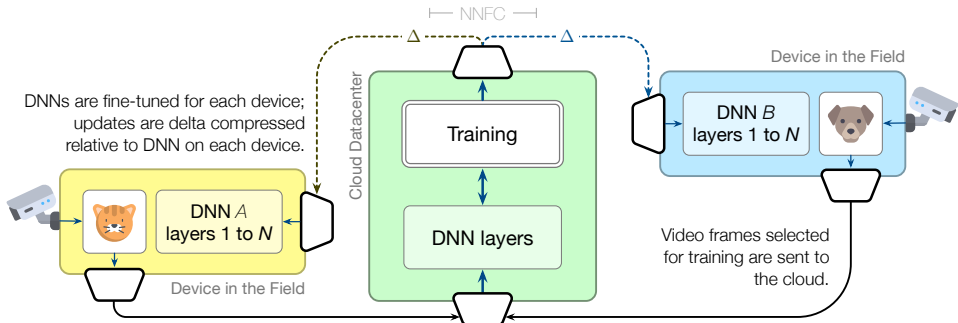
Burr's conjecture (Schapira and Winstein, HotNets 2017)

It is impossible for a decentralized congestion-control scheme *that greedily optimizes an objective function whose only input is the fate of its own traffic* to be globally asymptotically stable over a network with shared DropTail queues.

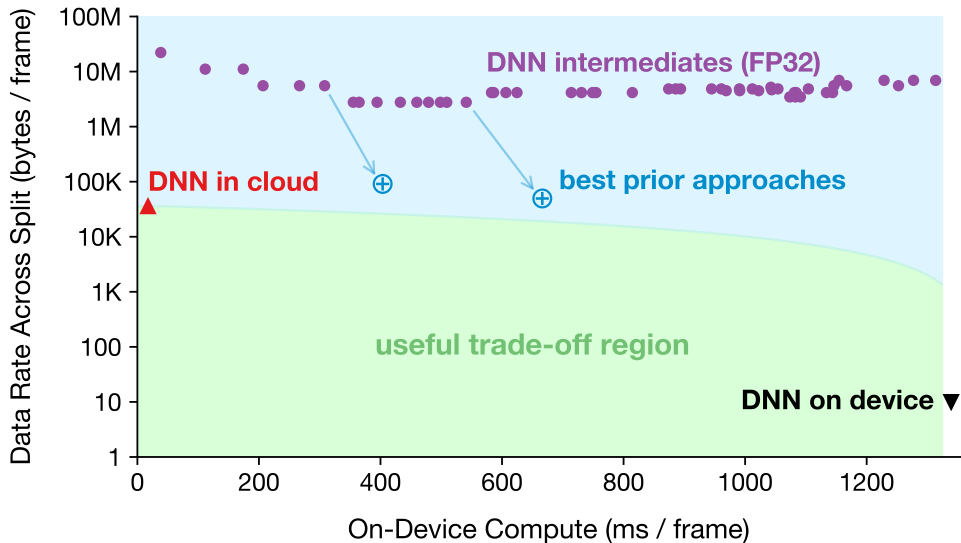
What if compute and data are separated?



What if compute and data are separated?



What if compute and data are separated?



The lesson? (paradigm 2)

Deploy-and-learn can be great, but **networked systems** present unique and interesting challenges worthy of research.

Paradigm 3: learn from the machines

Independent of ML's utility in deployment, machine learning can help us understand why systems ought to be the way they are.

Human: These are our requirements and objectives and design rules—what's the best system?

Machine: How about this?

Human: That's crazy! But, it does meet the requirements.
Hmm...

The lesson? (paradigm 3)



- ▶ Teaching something is the best way to learn anything.
- ▶ The dumber the student, the better the teacher learns.
- ▶ Machines are very dumb. Therefore...

Teaching machines to learn to design systems is the best way for us to learn to design systems.

- ▶ **Paradigm 1** [learn then deploy]:
Often harder than we expect.
(“Past performance is no guarantee of future results.”)
- ▶ **Paradigm 2:** [deploy and learn]:
Worth researching, but hard because of the nature of networks.
- ▶ **Paradigm 3:** [learn from the machines]:
Old-fashioned AI view—still valuable!


Slide palette info

The PowerPoint palette for this template has been built for you and is shown below. Avoid using too many colors in your presentation.




Use **Accent 1** as the main accent color.
Use **Accent 2** and **Accent 3** when additional colors are needed.

Use **Accents 4-6** sparingly – only when more colors are necessary.



Text with bullet points—adjusting list levels

- Main topic 1: size 28pt
 - Size 24pt for second level
 - Size 20pt for third level
 - Size 20pt for fourth level




Use the “Decrease List Level” and “Increase List Level” tools on the Home Menu to change text levels.

Try this:


1. Place your cursor in any row of text to the left that says “Size 20pt for subtopics”
2. Next click the Home tab, and then on the “Decrease List level” tool. Notice how the line moves up one level.
3. Now try placing your cursor in one of the “Main topic...” lines of text. Click the “Increase List Level” tool and see how the text is pushed in one level

Use these 2 tools to adjust your text levels as you work



Headline goes here

- Click to add text
 - Click to add text
 - Click to add text
 - Click to add text




Headline goes here

Click to add text

Click to add text

Click to add text

Click to add text



Headline goes here

- Click to add text
 - Click to add text
 - Click to add text
 - Click to add text

Headline goes here

Click to add text

Click to add text

Click to add text

Click to add text

Headline goes here

Transition or demo slide option 1

Subhead can go here

Transition or demo slide option 2

Subhead can go here

Transition or demo slide option 3

Subhead can go here

Thank you!

