

CS Research for Practitioners

Lessons from The Morning Paper

Speaker Adrian Colyer



CS Research for practitioners:

lessons from The Morning Paper

Adrian Colyer



@adriancolyer

the morning paper

an interesting/influential/important paper from the world of CS every weekday morning, as selected by Adrian Colyer

Home

A Year in Papers

DECEMBER 14, 2015

We've reached the end of term again, and I'm taking a break from writing up papers over the holidays – a chance to replenish my backlog and start planning for 2016 too! I want to see what I can do to improve the readability of the site as well. The Morning Paper will resume on the 4th January.

In a moment I'll share with you the **top 10 most read** and **most tweeted** papers, plus **some of my own picks**. But first a quick look back over the year. Through the course of 2015 I've posted **206** paper write-ups on The Morning Paper plus a few original pieces and other miscellaneous posts. That means I'm now at over 300 paper reviews in total since #themorningpaper began. It's amazing how a little every day adds up over time!

I'd like to say a huge thank-you to everyone who's been following along, I love all the interaction that the papers lead to. And if you're not yet subscribed to The Morning Paper and you're looking for a *New Years Resolution*, **signing up to the mailing list will get you**

SUBSCRIBE



never miss an issue! The Morning Paper delivered straight to your inbox.

SEARCH

ARCHIVES

Select Month ▼

MOST READ IN THE LAST FEW DAYS

> [ARIES: A Transaction Recovery Method Supporting Fine-Granularity Locking](#)



blog.acolyer.org

650




Foundations



Frontiers



One hot / 1-of-N



1-of-N
encoding

Distributed representation

	King	Queen	Woman	Princess	...
Royalty	0.99	0.99	0.02	0.98	
Masculinity	0.99	0.05	0.01	0.02	
Femininity	0.05	0.93	0.999	0.94	
Age	0.7	0.6	0.5	0.1	
...	...				

Finding meaning in context

... an efficient method for learning high quality distributed vector ...

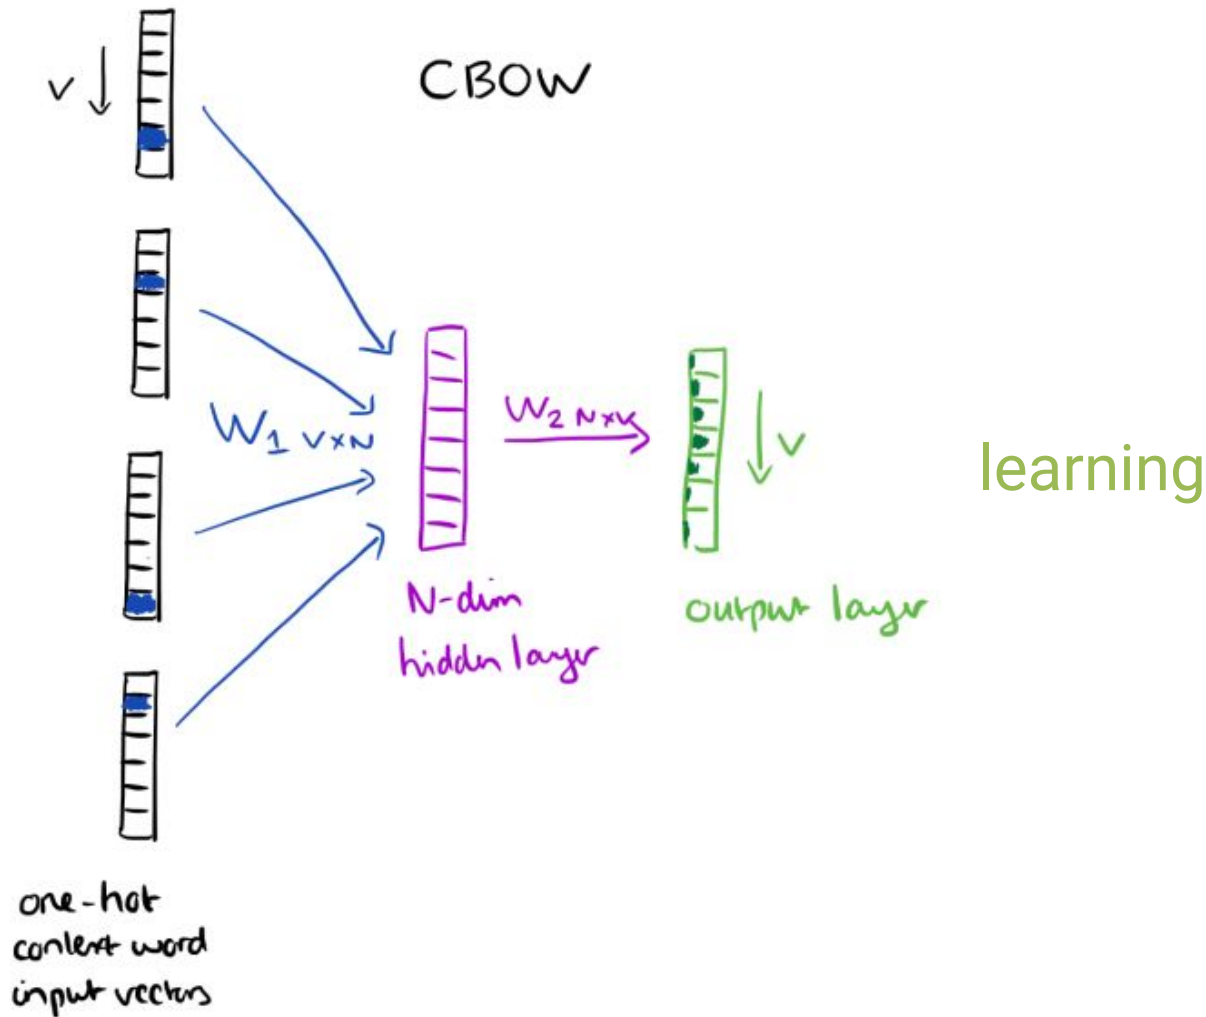
The diagram illustrates the concept of finding meaning in context for word embeddings. It features the sentence "... an efficient method for learning high quality distributed vector ...". The words "an efficient method" are grouped by a green bracket underneath and labeled "context" in green. The word "learning" is highlighted in yellow and has a blue arrow pointing up to it from the label "focus word" written in blue. The words "high quality distributed vector" are grouped by a green bracket underneath and labeled "context" in green.

method


for

high

quality



learning




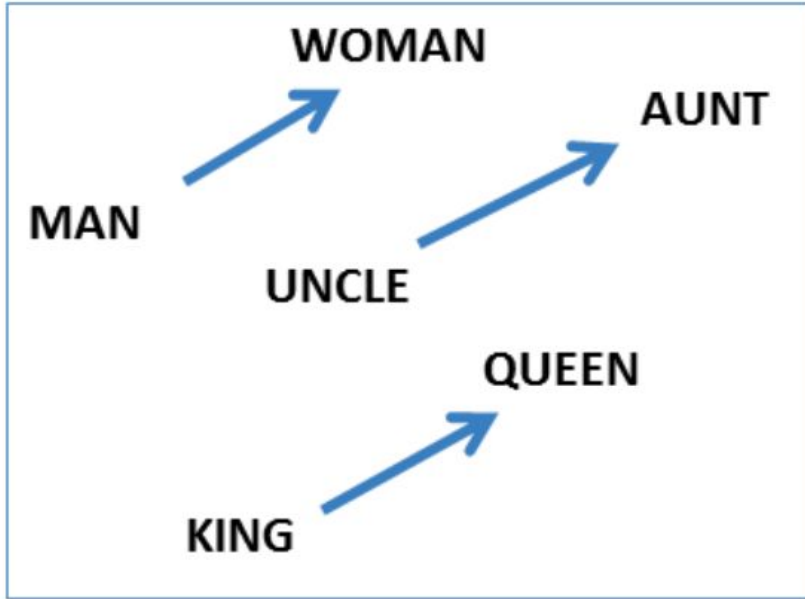
method

for

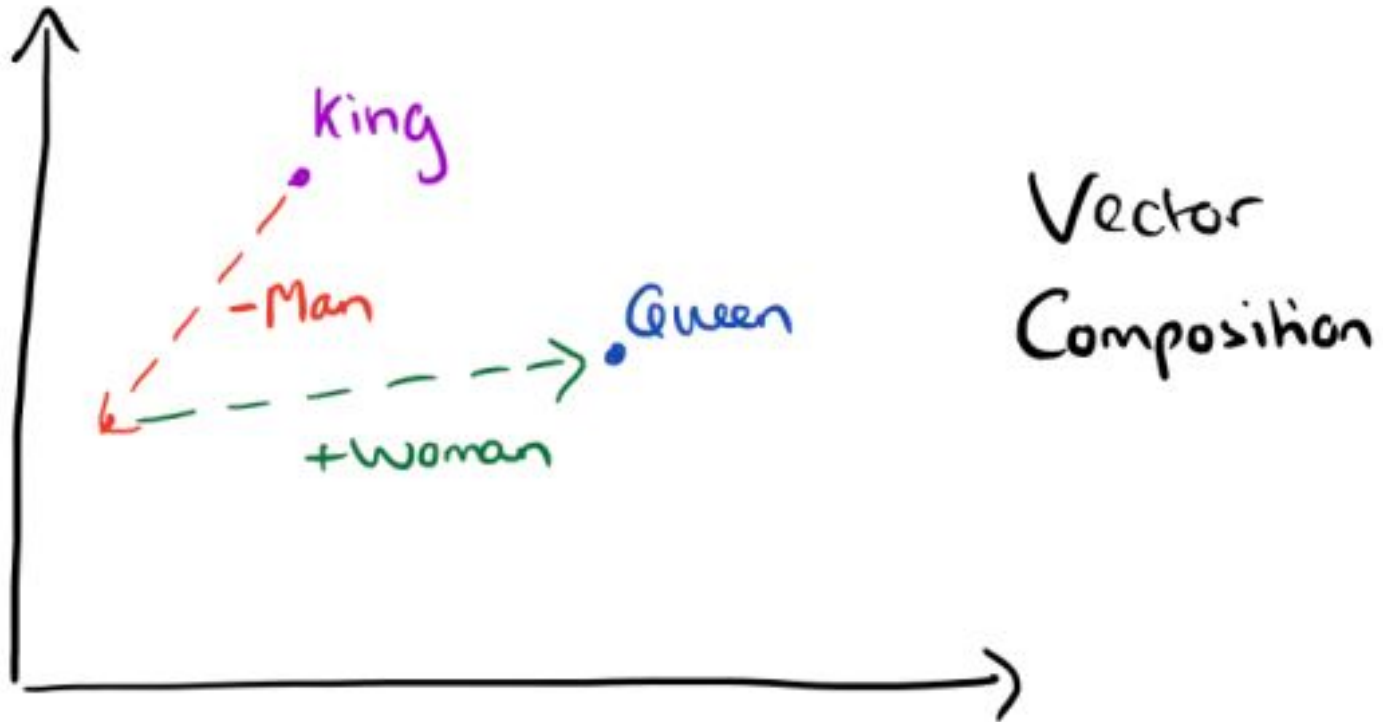
high

quality

Vector offsets



King - Man + Woman = ?



More examples

Relationship	Example 1	Example 2	Example 3
France - Paris	Italy: Rome	Japan: Tokyo	Florida: Tallahassee
Einstein - scientist	Messi: midfielder	Mozart: violinist	Picasso: painter
big - bigger	small: larger	cold: colder	quick: quicker

Czech + currency = Koruna

Vietnam + capital = Hanoi

German + airlines = Lufthansa

Russian + river = Volga

Papers so far...

- Efficient estimation of word representations in vector space, *Mikolov et al. 2013*
- Distributed representations of words and phrases and their compositionality, *Mikolov et al. 2013*
- Linguistic regularities in continuous space word representations, *Mikolov et al. 2013*
- word2vec parameter learning explained, *Rong 2014*
- word2vec explained: deriving Mikolov et al's negative sampling word-embedding method, *Goldberg & Levy 2014*
- See also: GloVe: Global vectors for word representation, *Pennington et al. 2014*

Using word embedding to enable semantic queries on relational databases, Bordawekar & Shmeuli, *DEEM*'17

Relation (table)

Word	Word	Word	Word

→ Sentence

Document

Find similar customers based on purchased items

```
SELECT X.custID, X.name, Y.custID, Y.name,  
similarityUDF(X.purchase, Y.purchase) AS sim  
FROM sales X, sales Y  
similarityUDF(X.purchase, Y.purchase) > 0.5  
ORDER BY X.name, sim  
LIMIT 10
```

Customers that have purchased allergenic items

```
SELECT X.number, X.name,  
similarityUDF(X.purchase, 'allergenic') AS sim  
FROM sales X  
similarityUDF(X.purchase, 'allergenic') > 0.3  
ORDER BY X.name, sim  
LIMIT 10
```


Accelerating innovation through analogy mining, Hope et al., *KDD'17*

How does the product work? What is the product good for?

Amazing Pillow

*A pillow combined with alarm clock, bluetooth, sensors and more features to improve and monitor sleep.

*Wake up comfortably with built in alarm clock

*Track sleep patterns

*Built in blind fold with led lights and sensors

*Full support for any kind of sleeper

*Alarm includes led lighting, vibrations and built in headphones for comfort.

Amazing Pillow

*A pillow combined with alarm clock, bluetooth, sensors and more features to improve and monitor sleep .

*Wake up comfortably with built in alarm clock

*Track sleep patterns

*Built in blind fold with led lights and sensors

*Full support for any kind of sleeper

*Alarm includes led lighting, vibrations and built in headphones for comfort .

*Near purpose,
Far mechanism.*

Figure 1: Collecting purpose and mechanism annotations from the crowd.

*“there is rich
meaning in
context”*



Are these ideas actually any good?





*“despite having data, **the number of companies that successfully transform into data-driven organisations stays low**, and how this transformation is done in practice is little studied.”*

The evolution of continuous experimentation in software product development, Fabijan et al., *ICSE'17*

Agile,
Lean,
CI,
CD, [2-way exchange]
CE




Image credit: Martin Fowler, "Microservices prerequisites"

Tech.

Crawl

Walk

Run

Fly

Metrics
Platform
Pervasiveness



Org.

Engineering team self-sufficiency
Experimentation team role




Biz.

OEC



A dirty dozen: twelve common metric interpretation pitfalls in online controlled experiments, Dmitriev et al., *KDD'17*

Logs: debug -> signals
Signals -> metrics



Seven rules of thumb for website experimenters, Kohavi et al., *KDD'14*

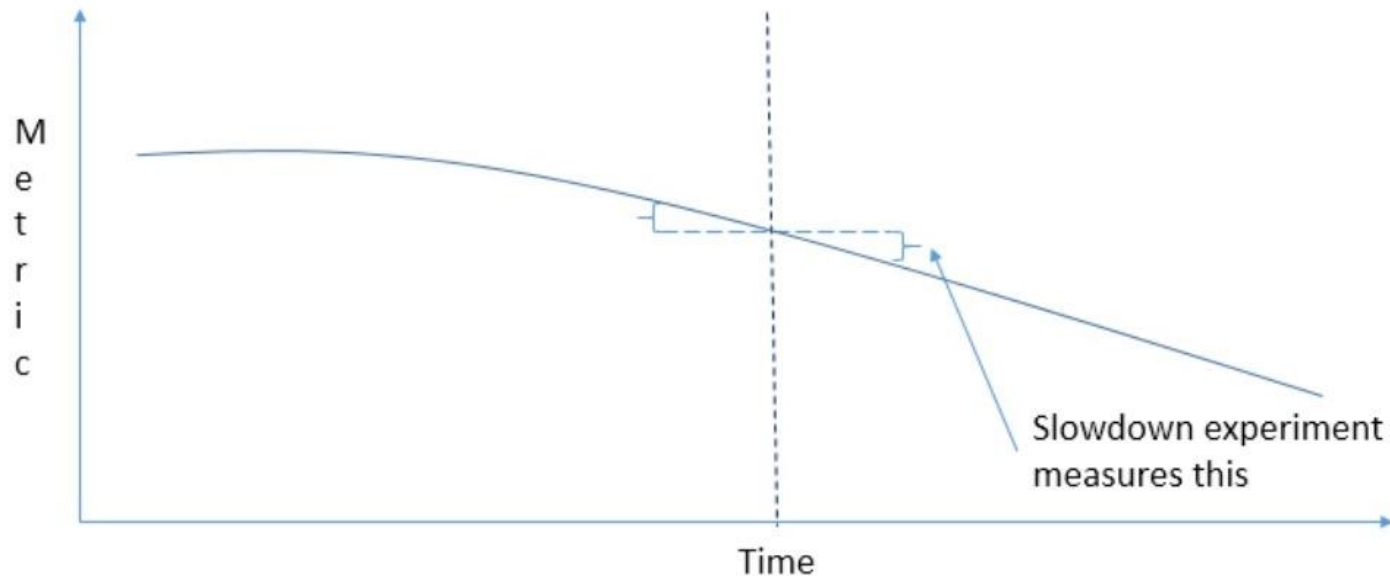



Figure 3: Typical relationship between performance (time) and a metric of interest



Randy Shoup @randyshoup · 5h

"If you want to move the needle on key metrics, ... invest in improving your site performance. Nothing else seems to come close."



Adrian Colyer @adriancolyer

Perf. and simplicity are both hugely important when you're trying to improve your key metrics. 7 rules of thumb: blog.acolyer.org/2017/09/26/sev...



1



4



9



Randy Shoup

@randyshoup

Follow

Performance improvements to the eBay search pages in 2010 increased overall revenue *for the entire company* by 2% - \$120M / year

3:26 PM - 26 Sep 2017

3 Retweets 7 Likes



2



3



7



“Any sufficiently complex system acts as a black box when it becomes easier to experiment with than to understand. Hence, black-box optimization has become increasingly important as systems become more complex.”

Google Vizier: a service for black-box optimization, Golovin et al., *KDD'17*



$$f: \mathcal{X} \rightarrow \mathbb{R}$$

(n^d)

Image credit: <https://pixabay.com>



Randy Shoup @randyshoup · 3h

Performance improvements to the eBay search pages in 2010 increased overall revenue *for the entire company* by 2% - \$120M / year



2



3



7



Randy Shoup

@randyshoup

Follow

Equivalent to our simultaneous introduction of machine-learned ranking in search, which gave another 2% - so another 120M / year

3:28 PM - 26 Sep 2017

3 Retweets 3 Likes



1



3



3





Randy Shoup @randyshoup · 3h

Equivalent to our simultaneous introduction of machine-learned ranking in search, which gave another 2% - so another 120M / year



1



3



3



Randy Shoup

@randyshoup

Follow

How did we prove that quarter billion in yearly revenue was due to our work, and not other factors?

We A/B tested, of course.

3:30 PM - 26 Sep 2017

5 Retweets 4 Likes



1



5



4



TFX: A TensorFlow-based production scale machine learning platform, Baylor et al., *KDD'17*

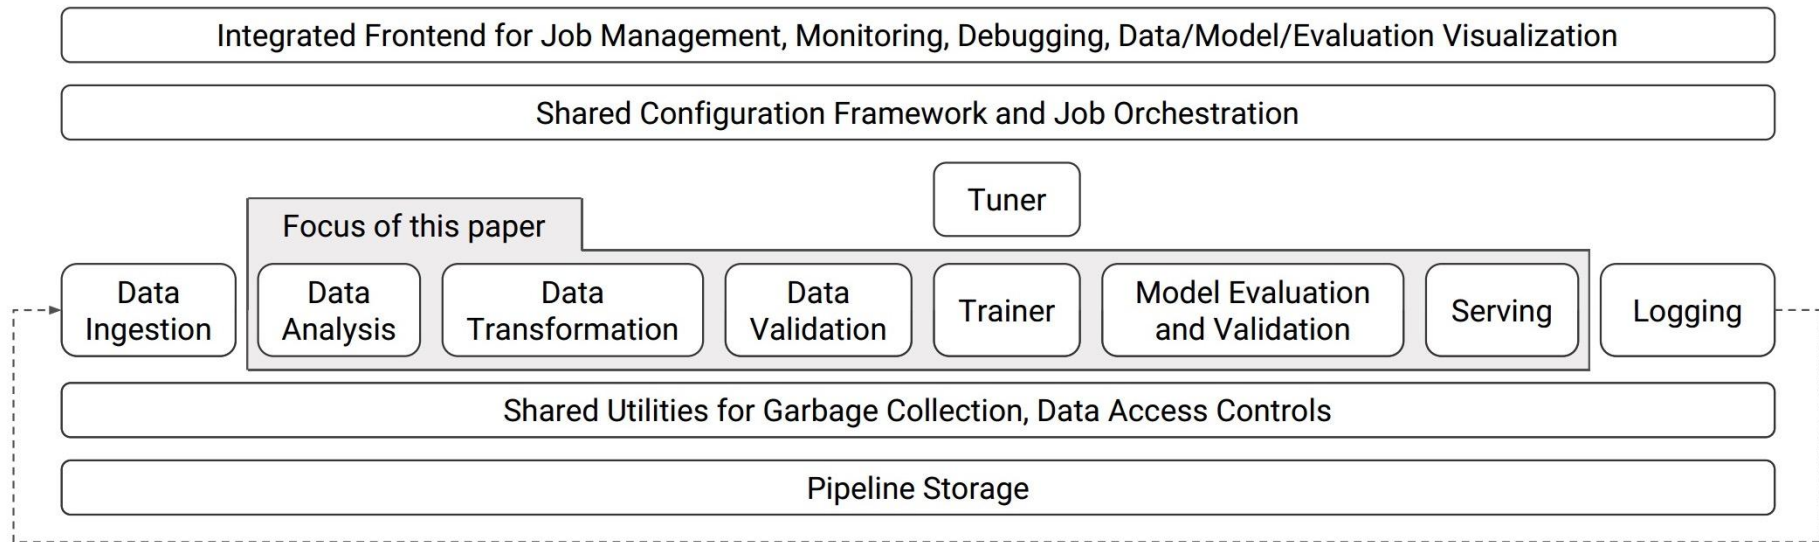




Figure 1: High-level component overview of a machine learning platform.


ActiveClean: Interactive data cleaning for statistical modeling, Krishnan et al., *VLDB'16*



(a) Systematic Error



(b) Mixed Dirty and Clean



(c) Sampled Clean Data

Neural Architecture Search with reinforcement learning, Zoph et al., *ICLR*'17


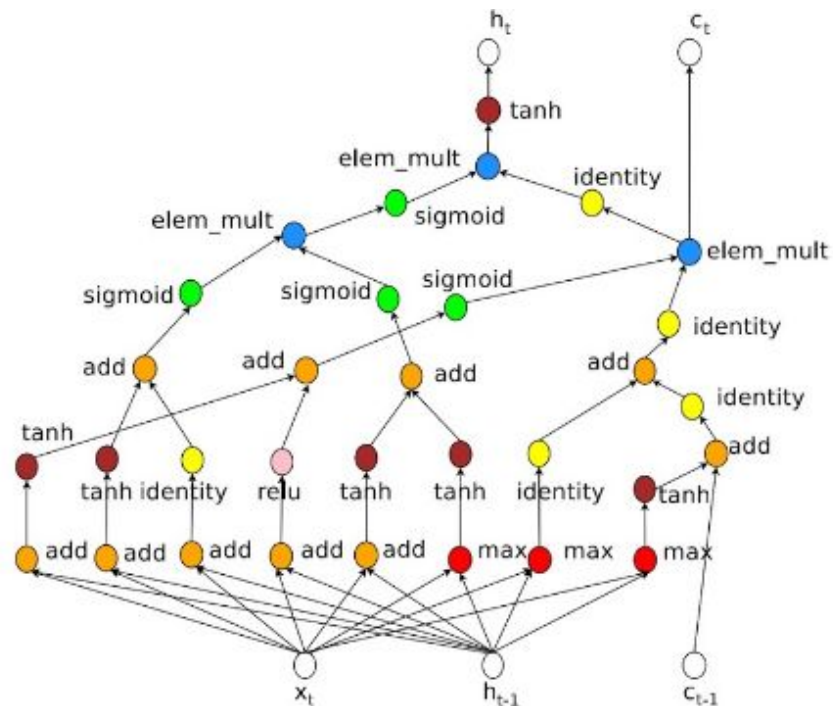
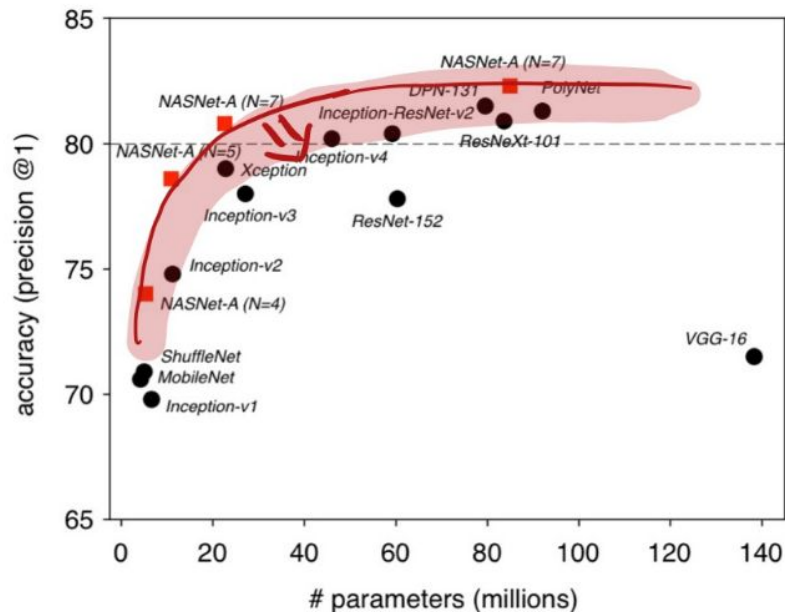



Figure 1: An overview of Neural Architecture Search.



Learning transferable architectures for scalable image recognition, Zoph et al., *ArXiv*'17





Ashwin Baskaran

@AshwinBaskaran

Follow



Great :-/ Now you have to be in the top 0.01% to make an impact [#deeplearning](#)
[blog.acolyer.org/2017/09/11/lea](#) ... via
[@adriancolyer](#) [#themorningpaper](#)

Will the *human expert* frontier move from designing learning systems to designing meta-learning systems?

6:21 PM - 11 Sep 2017

Neurosurgeon: collaborative intelligence between the cloud and the mobile edge, Kang et al., *ASPLOS'17*





Figure 5: The per layer execution time (the light-colored left bar) and size of data (the dark-colored right bar) after each layer's execution (input for next layer) in AlexNet. Data size sharply increases then decreases while computation generally increases through the network's execution.



Distributed deep neural networks over the cloud, the edge, and end devices, Teerapittayanon et al., *ICDCS'17*


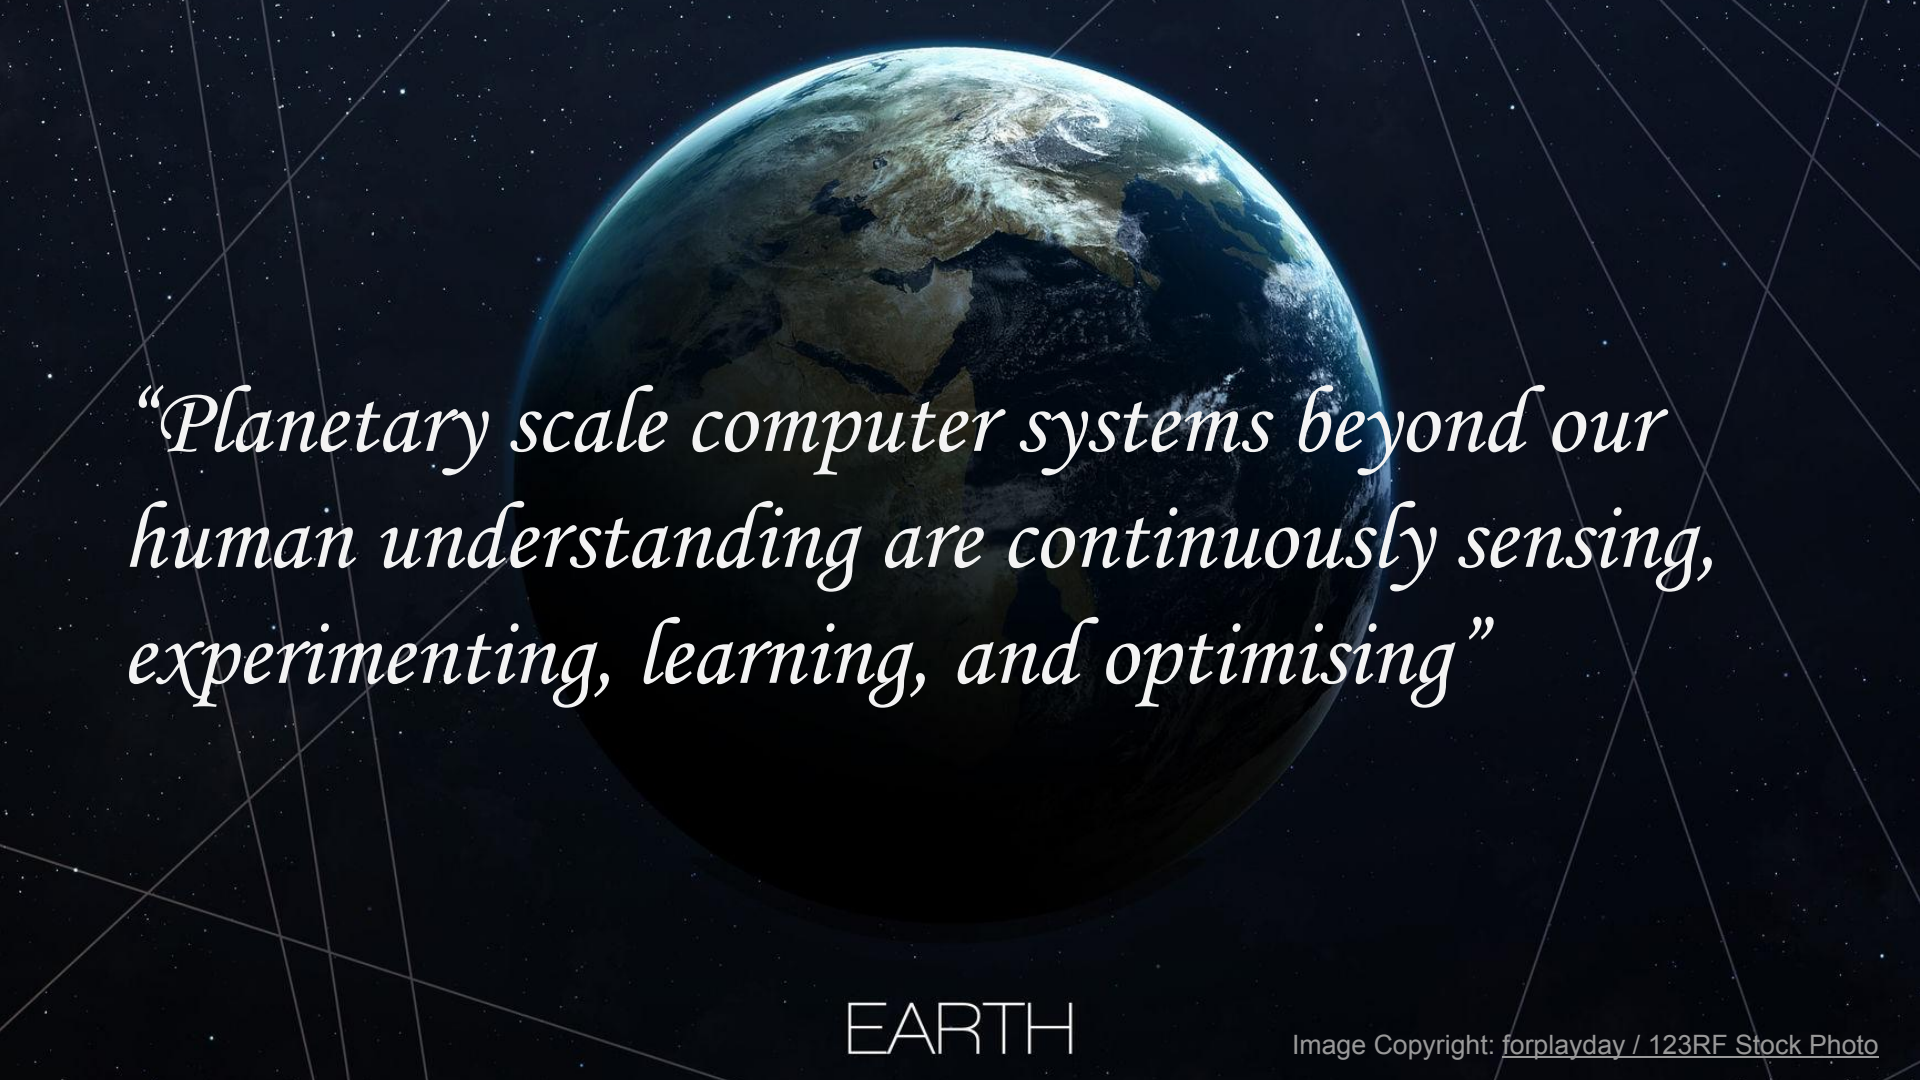




Figure 5. Example images of three objects (person, bus, car) from the multi-view multi-camera dataset. The six devices (each with their own camera) capture the same object from different orientations. An all grey image denotes that the object is not present in the frame.




“Planetary scale computer systems beyond our human understanding are continuously sensing, experimenting, learning, and optimising”

EARTH

Image Copyright: [forplayday / 123RF Stock Photo](#)


European Union regulations on algorithmic decision making and a “right to explanation”, Goodman & Flaxman, 2016



Practical black-box attacks against deep learning systems using adversarial examples, Papernot et al., *CCS'17*



Universal adversarial perturbations, Moosavi-Dezfooli et al., *CVPR'17*



Indian elephant



Indian elephant

Adversarial examples for evaluating reading comprehension systems, Jia & Liang, *EMNLP'17*

Article: Super Bowl 50

Paragraph: “Peyton Manning became the first quarterback ever to lead two different teams to multiple Super Bowls. He is also the oldest quarterback ever to play in a Super Bowl at age 39. The past record was held by John Elway, who led the Broncos to victory in Super Bowl XXXIII at age 38 and is currently Denver’s Executive Vice President of Football Operations and General Manager. *Quarterback Jeff Dean had jersey number 37 in Champ Bowl XXXIV.*”

Question: “What is the name of the quarterback who was 38 in Super Bowl XXXIII?”

Original Prediction: John Elway

Prediction under adversary: Jeff Dean

	Match Single	Match Ens.	BiDAF Single	BiDAF Ens.
Original	71.4	75.4	75.5	80.0
ADDSSENT	27.3	29.4	34.3	34.2
ADDONESSENT	39.0	41.8	45.7	46.9
ADDANY	7.6	11.7	4.8	2.7
ADDCOMMON	38.9	51.0	41.7	52.6

Table 2: Adversarial evaluation on the Match-LSTM and BiDAF systems. All four systems can be fooled by adversarial examples.

Figure 1: An example from the SQuAD dataset. The BiDAF Ensemble model originally gets the answer correct, but is fooled by the addition of an adversarial distracting sentence (in blue).

IoT goes nuclear: creating a ZigBee chain reaction, Ronen et al., *IEEE Security & Privacy* 2017




Figure 2. The ZLL architecture.



Figure 3. Philips Hue bridge (gateway), lamps, and wireless switch.



Figure 9. ZigBee warflying scenario

“What we demonstrate in this paper is that even IoT devices made by companies with deep knowledge of security, which are backed by industry standard cryptographic techniques, can be misused by hackers and rapidly cause city-wide disruptions which are very difficult to stop.”

CLKSCREW: Exposing the perils of security-oblivious energy management, Tang et al., *USENIX Security 2017*

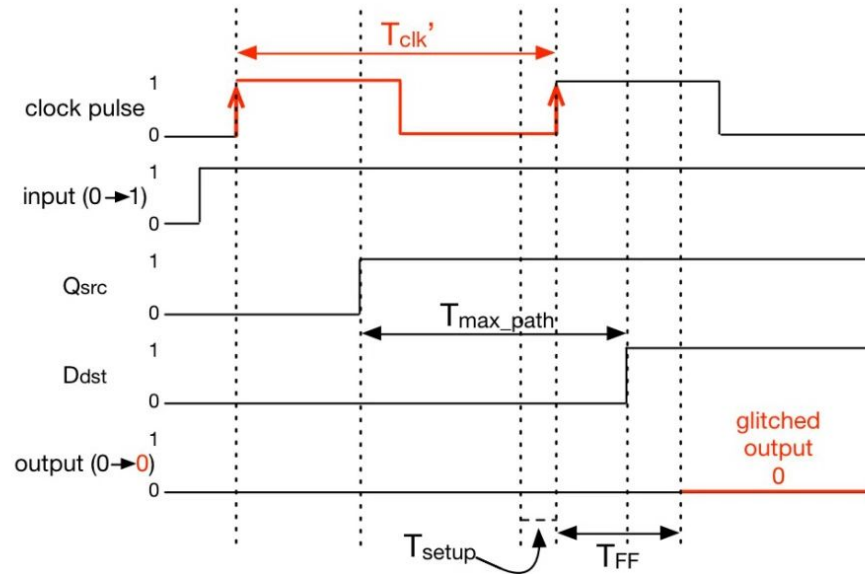


Figure 4: Bit-level fault due to overclocking: Reducing clock period $T_{clk} \rightarrow T_{clk}'$ results in a bit-flip in output $1 \rightarrow 0$.



Image Copyright: [sepavo](#) / 123RF Stock Photo

REM: Resource-efficient mining for blockchains, Zhang et al., *USENIX Security 2017*




Figure 1: Architecture overview of REM

I'm just building
a webapp!
Does any of this
research stuff
apply to me?



Feral concurrency control: an empirical investigation of modern application integrity, Bailis et al., *SIGMOD*'15

*“By shunning decades of work on native database concurrency control solutions, Rails has developed a set of primitives for handling application integrity in the application tier—building, from the underlying database system’s perspective, **a feral concurrency control system.**”*

ACIDRain: concurrency-related attacks on database backed web applications, Warszawski & Bailis, *SIGMOD*'17

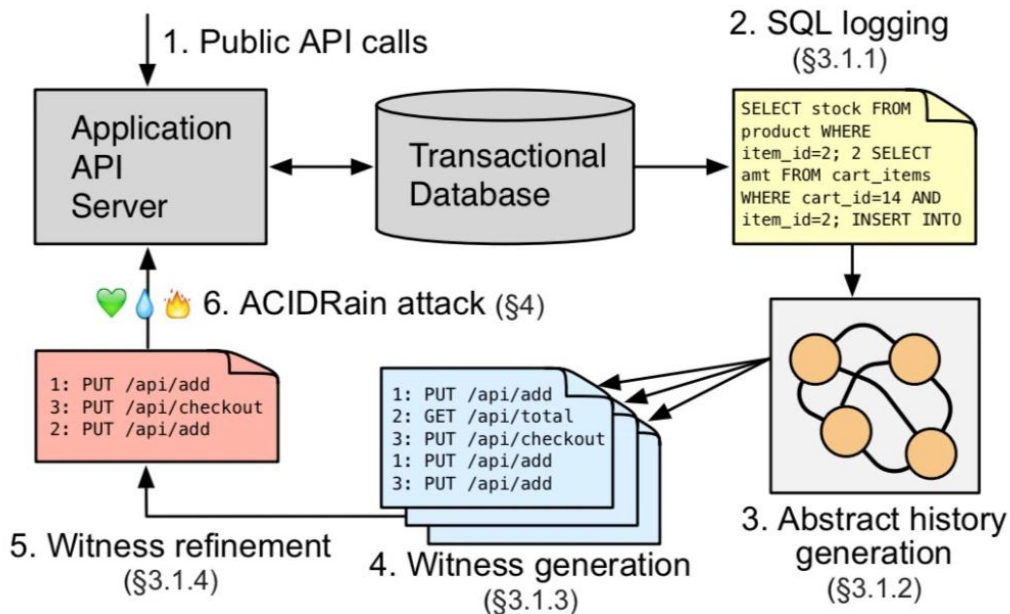


Figure 2: 2AD workflow to discover ACIDRain attacks.

12 eCommerce apps

60% top 1M Commerce sites

22 vulnerabilities

2 hours or less to craft
an exploit for each

Thou shalt not depend on me: analysing the use of outdated JavaScript libraries on the web, Launinger et al., *NDSS'17*

37% vulnerable

jQuery -> 36.7%,
Angular -> 40.1%




Fig. 2. Example causality tree.

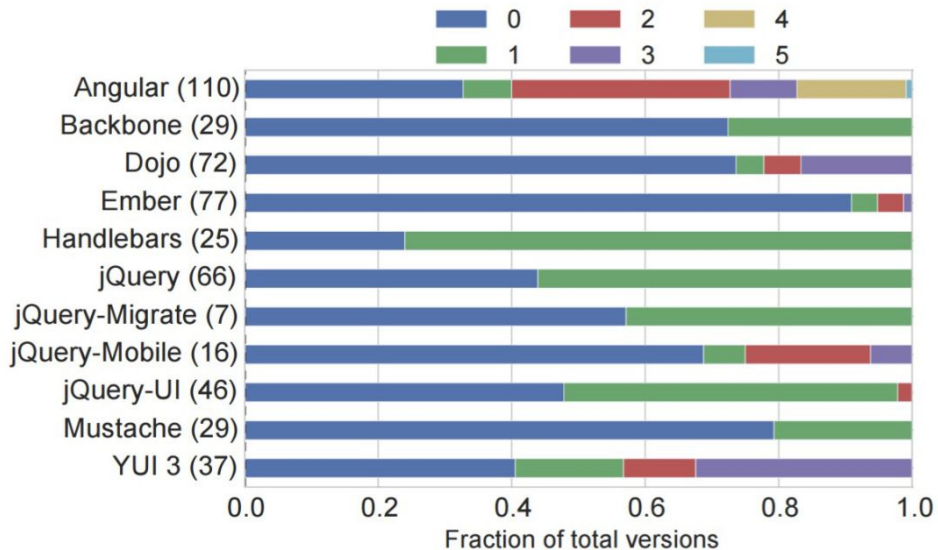


Fig. 1. Fraction of library versions with i distinct known vulnerabilities each (represented by colours), out of the total library versions in parentheses. Angular 1.2.0 has 5 known vulnerabilities and there are 110 versions overall.

To type or not to type: quantifying detectable bugs in JavaScript, Gao et al., *ICSE'17*




Fig. 1: The error model of this experiment.

Wrapping Up



*Welcome to the crazy,
wonderful, exciting,
sometimes terrifying, but
always fascinating world of
computer science research!*

the morning paper

an interesting/influential/important paper from the world of CS every weekday morning, as selected by Adrian Colyer

Home

Blurred Persistence: Efficient Transactions in Persistent Memory

JANUARY 21, 2016

Blurred Persistence: Efficient Transactions in Persistent Memory – Lu, Shu, & Sun, 2015

We had software transactional memory (STM), then hardware support for transactional memory (HTM), and now with persistent memory in which the **memory plays the role of stable storage**, we can have *persistent transactional memory*. And with persistent transactional memory, there's an issue that will surely make you smile with recognition: in-place of managing the relationship between volatile memory and disk, we now have to manage the relationship between the volatile CPU cache and memory! It's all the same considerations (forcing, stealing etc.) but in a new context and with a few new twists. Chief among those twists is that you have a lot less control over how and when the hardware moves data from cache to memory than you do over how and when you move data from memory to disk.

In case you find all these various permutations of non-volatile memory / storage confusing (I do!), then this might help:

SUBSCRIBE



never miss an issue! The Morning Paper delivered straight to your inbox.

SEARCH

ARCHIVES

Select Month ▼

MOST READ IN THE LAST FEW DAYS

- > [Petuum: A New Platform for Distributed Machine Learning on Big Data](#)
- > [Chimera: Large-Scale Classification Using Machine](#)

01

A new paper every weekday

Published at <http://blog.acolyer.org>.

02

Delivered Straight to your inbox

If you prefer email-based subscription to read at your leisure.

03

Announced on Twitter

I'm @adriancolyer.

04

Go to a Papers We Love Meetup

A repository of academic computer science papers and a community who loves reading them.

Papers We Love $f(x)=x$

05

Share what you learn

Anyone can take part in the great conversation.



THANK YOU !


@adriancolyer



TALK
soon



Cartoon images credit: Bitmoji



Please

**Remember to
rate this session**

Thank you!

