




1950

FUTURE VISION



21 MARCH 2019

COMPUTER VISION

Project 4: GCP and such

- Quota emails
 - We won't hold this against you
- Out of memory
 - Anyone still having that?
- What other showstoppers?

Newsflash: ML Drama



NeurIPS (Neural Information Processing Systems) 2017.

Ali Rahimi (Google) just won a 'test of time' award for work on applying random functions as a basis for training; gave a speech.

"Artificial Intelligence is the new electricity" – Ng, 2017

"ML has become alchemy" – Rahimi, 2017



Rahimi, NeurIPS 2017 test of time

“If you’re building photo-sharing systems, alchemy is ok, but we’re beyond that now.

We’re building systems that govern healthcare and mediate our civic dialog; we influence elections.

I would like to live in a society whose systems are build on verifiable, rigorous, thorough knowledge, and not on alchemy.”

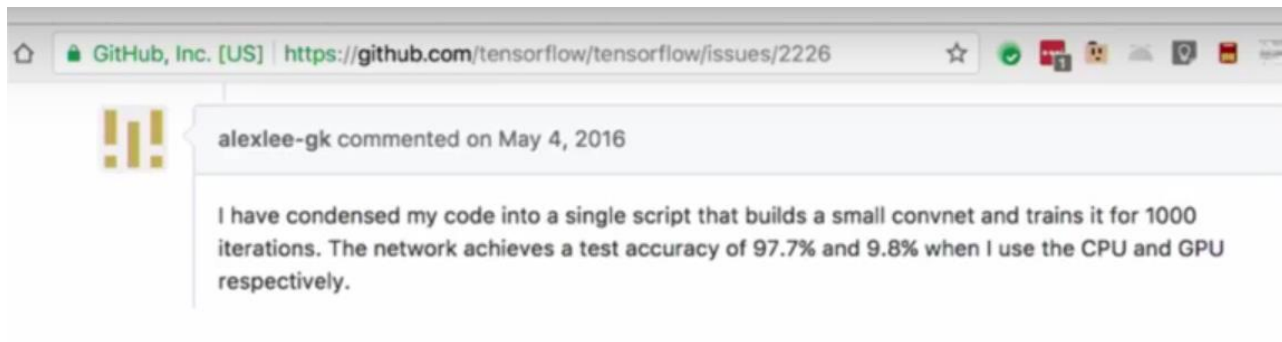
“This is happening because we apply brittle optimization strategies to loss surfaces we don’t understand.”

From: Boris

To: Ali

"On Friday, someone on another team changed the default rounding mode of some Tensorflow internals (from "truncate toward 0" to "round to even").

Our training broke. Our error rate went from <25% error to ~99.97% error (on a standard 0-1 binary loss)."



The screenshot shows a web browser window with the address bar displaying "GitHub, Inc. [US] | https://github.com/tensorflow/tensorflow/issues/2226". The page content shows a comment from user "alexlee-gk" dated "May 4, 2016". The comment text reads: "I have condensed my code into a single script that builds a small convnet and trains it for 1000 iterations. The network achieves a test accuracy of 97.7% and 9.8% when I use the CPU and GPU respectively."

alexlee-gk commented on May 4, 2016

I have condensed my code into a single script that builds a small convnet and trains it for 1000 iterations. The network achieves a test accuracy of 97.7% and 9.8% when I use the CPU and GPU respectively.

Yann LeCun retorts:



Yann LeCun

Yesterday at 11:57 · 🌐

In the history of science and technology, the engineering artifacts have almost always preceded the theoretical understanding: the lens and the telescope preceded optics theory, the steam engine preceded thermodynamics, the airplane preceded flight aerodynamics, radio and data communication preceded information theory, the computer preceded computer science.

Why? Because theorists will spontaneously study "simple" phenomena, and will not be enticed to study a complex one until there a practical importance to it.

Why dangerous? It's exactly this kind of attitude that lead the ML community to abandon neural nets for over 10 years, *despite* ample empirical evidence that they worked very well in many situations.

Yoshua Bengio chimes in:



Yoshua Bengio Ali Rahimi, many of us want to make deep learning understandable, and I have worked hard for more than two decades towards that goal in a number of papers. The problem with your talk is the attitude, which could be wrongly interpreted and lead to a repetition of past mistakes, as Yann wrote. Using the word Alchemy was not a good idea. However, I agree with you that I would like to see more science, more specifically more experimental science by opposition to engineering (see Yann's distinction above). Both have their role, but focusing the experiments on testing hypotheses which help us understand phenomena is not getting as much rewarded (by reviewers) as beating some benchmark. And THAT is unfortunate and needs a cultural change.

Like · Reply ·   33 · 18 hrs

Compute cost / human cost

- Recent CVPR best paper: 50,000 GPU hours
- Google's commercial object recognizer: continually training for years.

“JFT is an internal Google dataset that has 100 million labeled images with 15,000 labels. When we did this work, Google's baseline model for JFT was a deep convolutional neural network that had been trained for about six months using asynchronous stochastic gradient descent on a large number of cores.”

TO COMPLETE YOUR REGISTRATION, PLEASE TELL US
WHETHER OR NOT THIS IMAGE CONTAINS A STOP SIGN:



NO YES

ANSWER QUICKLY—OUR SELF-DRIVING
CAR IS ALMOST AT THE INTERSECTION.

SO MUCH OF "AI" IS JUST FIGURING OUT WAYS
TO OFFLOAD WORK ONTO RANDOM STRANGERS.

XKCD; thanks to **Iuliu Balibanu**

TO COMPLETE YOUR REGISTRATION, PLEASE TELL US
WHETHER OR NOT THIS IMAGE CONTAINS A STOP SIGN:



NO YES

ANSWER QUICKLY—OUR SELF-DRIVING
CAR IS ALMOST AT THE INTERSECTION.

SO MUCH OF "AI" IS JUST FIGURING OUT WAYS
TO OFFLOAD WORK ONTO RANDOM STRANGERS.

*Alt-text: "Crowdsourced steering" doesn't
sound quite as appealing as "self driving".*

Fast.ai – Aug 8th 2018

“...managed to train [ImageNet](#) to 93% accuracy in just 18 minutes, using 16 machines each with 8 [NVIDIA V100](#) GPUs, running the [fastai](#) and [PyTorch](#) libraries. This is a new speed record for training Imagenet to this accuracy...”

NVIDIA V100 -> \$8500k

$16 * 8 * 8500 = \$\$\$$ a million bucks

“... on publicly available public [AWS](#) cloud instances infrastructure, and costs around \$40 to run.”

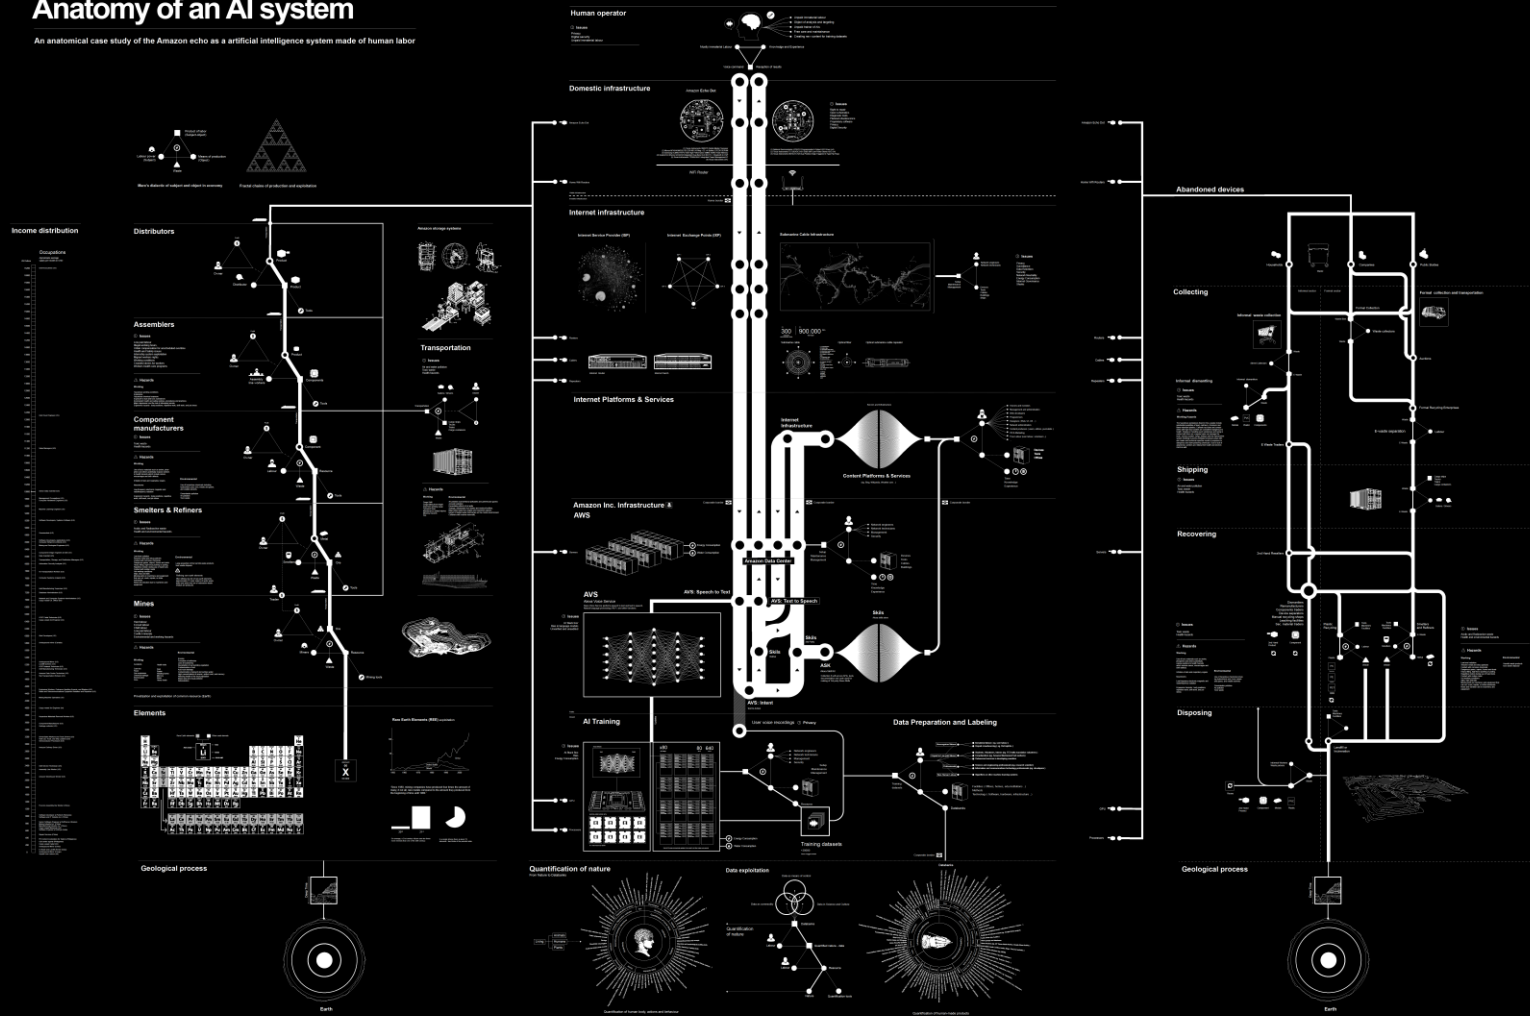
Anatomy of an AI System

- <https://anatomyof.ai/>

By Kate Crawford ¹ and Vladan Joler ²
(2018)

Anatomy of an AI system

An anatomical case study of the Amazon echo as a artificial intelligence system made of human labor



Rodney Brooks – The Mistakes We Make

- <https://rodneybrooks.com/the-seven-deadly-sins-of-predicting-the-future-of-ai/>
- Quote in notes.

Rodney Brooks – Artificial General Intelligence Tests

2, 4, 6, 8 year old competency tasks:

object recognition, language understanding, manual dexterity, social understanding.

“A two year old child can know that something is deliberately meant to function as a chair even if it is unlike any chair they have seen before. It can have a different number of legs, it can be made of different material, its legs can be shaped very oddly, it can even be a toy chair meant for dolls. A two year old child is not fazed by this at all. Despite having no visual features in common with any other chair the child has ever seen before the child can declare a new chair to be a chair. This is completely different from how a neural network is able to classify things visually.”

“But wait, there is more! A two year old can do one-shot visual learning from multiple different sources. Suppose a two year old has never been exposed to a giraffe in any way at all. Then seeing just one of a hand drawn picture of a giraffe, a photo of a giraffe, a stuffed toy giraffe, a movie of a giraffe, or seeing one in person for just a few seconds, will forever lock the concept of a giraffe into that two year old’s mind. That child will forever be able to recognize a giraffe as a giraffe, whatever form it is represented in. Most people have never seen a live giraffe, and none have ever seen a live dinosaur, but the are easy for anyone to recognize.”

But James...

...I thought we could treat machine learning like a magical black box? I like black boxes.

Deep learning is:

- a black box
- *also a black art.*
- Grad student
gradient descent : (



CV and society

CV / ML is a tool

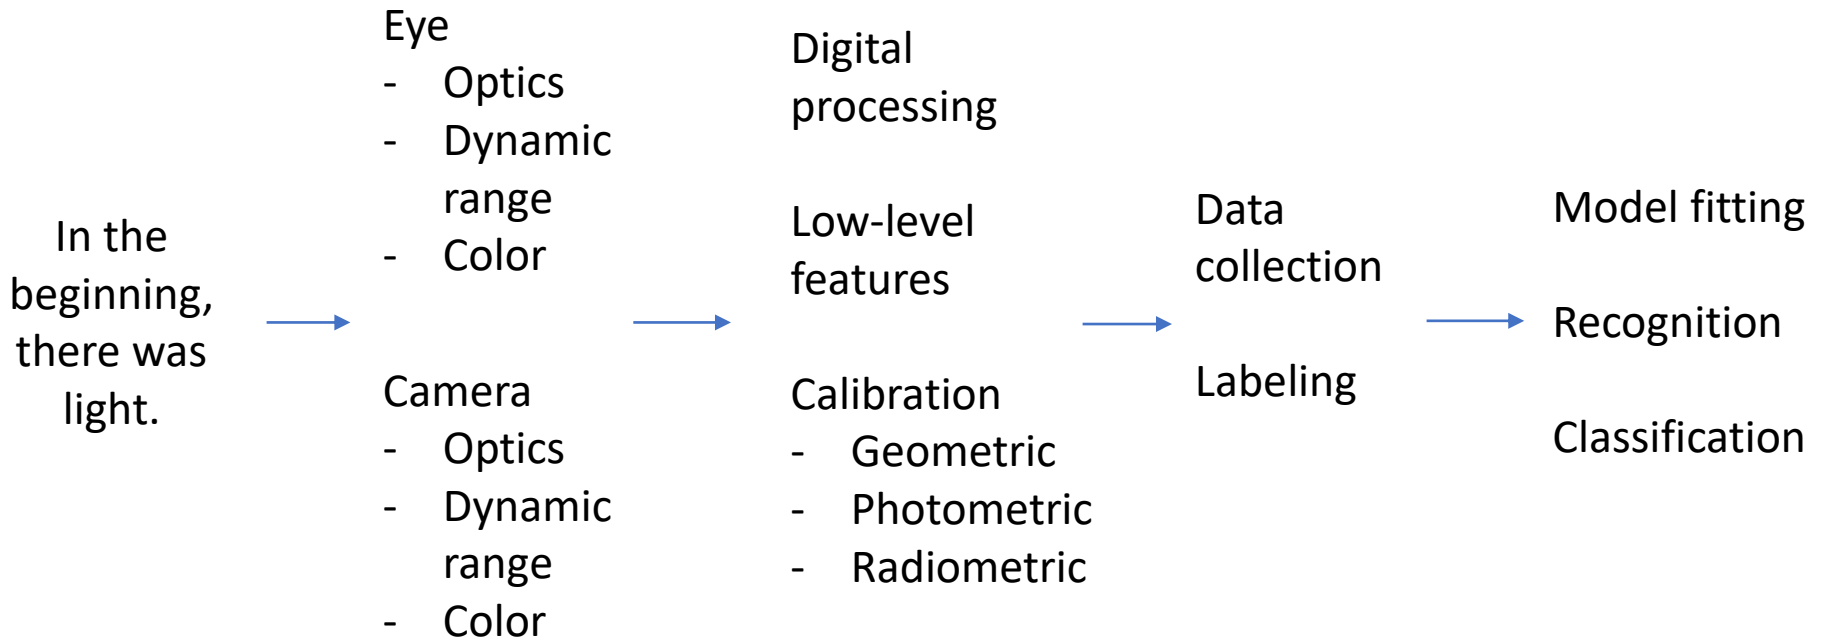
Any tool can be used for good or bad.

- Definition of good/bad *can* vary with point of view.

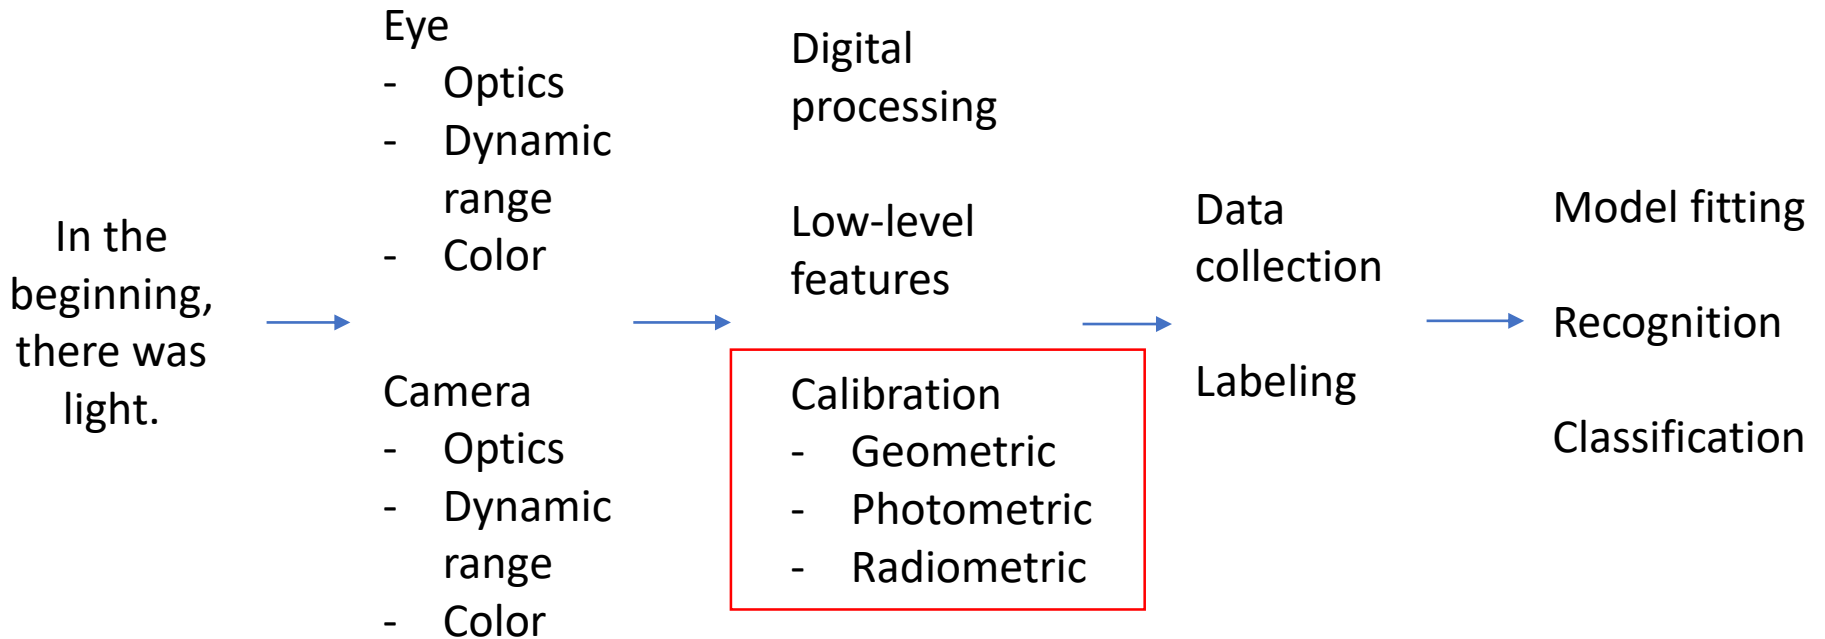
Tools are created and used in the real world.

- Time/money trade-offs
- Different agendas
- Advertent or inadvertent
- With or without awareness

Computer vision domain





Computer vision domain





Light response curves

Eye
Sensor:

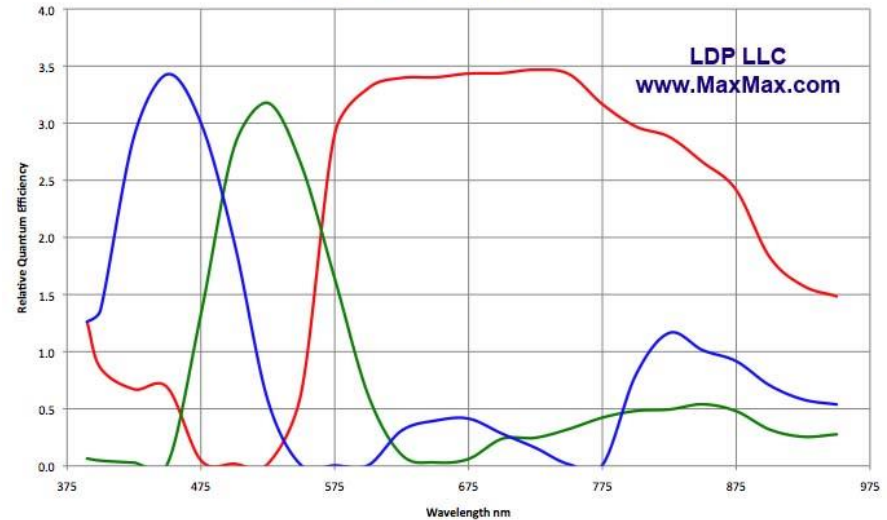


© Stephen E. Palmer, 2002

Camera
Sensor:




Canon 450D Quantum Efficiency



Light/reflectance output curves



250-500:1 contrast ratio (OLED = inf.)
6 / 8 / 10 bit dynamic range
3 / 4 additive primaries (RGB, rarely +yellow)
Defines a gamut



50-150:1 contrast ratio
??? dynamic range
4 subtractive primaries (CMYK)
Defines a gamut

We want:

Colors we see with our eyes in the world

=

Colors we see with our eyes in the reproduction

How do we calibrate these?

Time Warp: Photo Film processing

Price of film roll included film development costs.
Actually a monopoly!
1954 – broken up; independent film developers.





Not actually a Kodak shop...

An independent chain!

New problem:
How to control quality of output across these different developers?



Here's the process:

1. Kodak produces a photosensitive chemical film.
 - Different compounds gave different spectral responses.
 - Kodak calibrates these responses using photometers and using test scenes.
2. You buy Kodak film from the store and put it in your camera.
 - Play with your family in the park
 - Take photos.
3. You take the film back to the store/lab for development.
 - Kodak needs the lab to reproduce the film's intended spectral response
 - Needs to *calibrate* the chemical development/printer to a target

Solution:

Kodak sends test negatives and test prints to printer technicians to check that their output matches the intended response of the film!

Kodak's test input + output

- 'Shirley cards' – 1950s/60s
- Shirley was photographed hundreds of times by Kodak.
- One negative was processed as per Kodak specifications.
- A new unexposed negative + processed output was sent to each printer lab.
- Colors were calibrated on site to match the target Shirley card.



Circa 1960

Kodak's test input + output

- 'Shirley cards' – 1950s/60s
- Any issues with this approach?



Circa 1960


Over time

- 1978: Filmmaker Jean-Luc Godard refuses to use Kodachrome film in Mozambique.
- 1980s: Chocolate and furniture manufacturers complain.
- 1986: Kodacolor VR-G (or Gold) – film for dark browns.
 - “Photograph the details of a dark horse in low light.”




Shirley card, 1996

1980s – adverts



The Four Tops!



Bill Cosby!

Some other issues here now too : (

What are the underlying
technical problems?

*How do they relate to egalitarianism
...and how might we overcome them?*

Think-pair-share.

Issues

- Dynamic range: not enough!
- Color balance:

So digital fixes this, right?

- Well...

“The hardest part of being in a biracial relationship is taking a picture together.”



So digital fixes this, right?

...it's a lot better.

- 14-bit sensors (\approx eye's static range)
- High-dynamic range by combining low-dynamic range
- Digital post-processing for color balance

References

Canadian Journal of Communication:

Roth et al., Looking at Shirley, the Ultimate Norm: Colour Balance, Image Technologies, and Cognitive Equity

<http://www.cjc-online.ca/index.php/journal/article/view/2196>

<http://www.npr.org/2014/11/13/363517842/for-decades-kodak-s-shirley-cards-set-photography-s-skin-tone-standard/>

<https://priceconomics.com/how-photography-was-optimized-for-white-skin/>

<https://www.buzzfeed.com/syreetamcfadden/teaching-the-camera-to-see-my-skin/>

Word of warning

- Around 2013/2014 there were a lot of articles about this issue.
- Many articles rewrite the same few sources.
- Some do not have a technical background, and sometimes technical issues are confused.
- ‘Internet comments’ warning.
- ‘Take care.’

CV as making bank

- Intel buys Mobileye!
- \$15 billion
- Mobileye:
 - Spin-off from Hebrew University, Israel
 - 450 engineers
 - 15 million cars installed
 - 313 car models

NEWS

Intel buys driverless car technology firm Mobileye

🕒 2 hours ago | Business | 📄 138

🔗 Share



US chipmaker Intel is taking a big bet on driverless cars with a \$15.3bn (£12.5bn) takeover of specialist Mobileye.

Intel will pay \$63.54 a share in cash for the Israeli company, which develops "autonomous driving" systems.

Mobileye and Intel are already working together, along with German carmaker BMW, to put 40 test vehicles on the road in the second half of this year.

Intel expects the driverless market to be worth as much as \$70bn by 2030.

Technology companies are racing to launch driverless cars.

June 2016 - Tesla left Mobileye

- Fatal crash – car ‘autopilot’ ran into a tractor trailer.

“What we know is that the vehicle was on a divided highway with Autopilot engaged when a tractor trailer drove across the highway perpendicular to the Model S. Neither Autopilot nor the driver noticed the white side of the tractor trailer against a brightly lit sky, so the brake was not applied.” – [Tesla blog](#).

What computer vision problems
does this sound like?

Tesla crash: how it happened

A preliminary investigation into 25,000 Tesla Model S cars has been opened after a driver of one of the vehicles was killed while operating in Autopilot mode in a crash in Williston, Florida. Here is how the fatal accident occurred according to authorities.




Diagram not to scale.

Sources: Florida Highway Patrol Troop; U.S. National Highway Traffic Safety Administration

C. Chan, 30/06/2016



June 2016 - Tesla left Mobileye

- Fatal crash – car ‘autopilot’ ran into a tractor trailer.

“What we know is that the vehicle was on a divided highway with Autopilot engaged when a tractor trailer drove across the highway perpendicular to the Model S. Neither Autopilot nor the driver noticed the white side of the tractor trailer against a brightly lit sky, so the brake was not applied.” – [Tesla blog](#).

What computer vision problems
does this sound like?

What HCI problems does
this sound like?

Autosteer

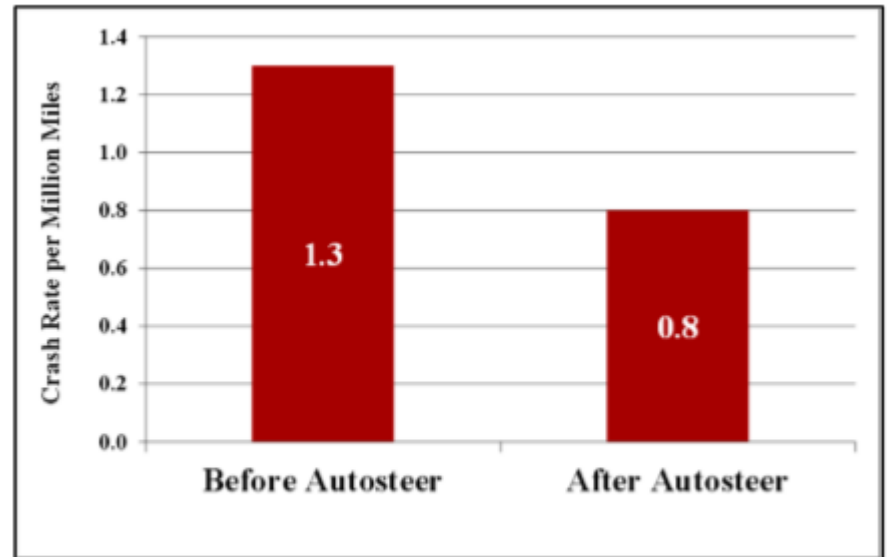


Figure 11. Crash Rates in MY 2014-16 Tesla Model S and 2016 Model X vehicles Before and After Autosteer Installation.

- <https://www.bbc.com/news/technology-47468391>
- **Uber 'not liable' for self-driving death**

Instagram filters

- Filters that brighten
- Filters that darken
- Filters can do anything!



Snapchat



select bitch @caseyjohnston · 20 Apr 2016
oh god @snapchat you didn't pic.twitter.com/IBZUHZKODg

5 57 74



grace
@tequilafunrise

Follow

.@Snapchat wanna tell me why u thought this yellowface was ok??



RETWEETS 155
LIKES 196



10:13 AM - 9 Aug 2016


86 155 196

“Anime inspired”

FaceApp

- Learning-based face transformations

Make them smile



Meet your future self



Look younger



Change gender





Got a tip? [Let us know.](#)

Follow Us [f](#) [ig](#) [tw](#) [yt](#) [fb](#) [in](#) [g+](#) [rss](#)

[News](#) [Video](#) [Events](#) [Crunchbase](#)

[Message Us](#)



DISRUPT NY Mike Einziger of Incubus And Pharrell Williams Are Coming To Disrupt NY To Debut New Audio Tech [Find Out More](#) ▶

neural networks

FaceApp

algorithmic bias

algorithmic accountability

Artificial Intelligence

Popular Posts



Doug finds the best Amazon deals
4 days ago



Elon Musk's Neuralink wants to turn cloud-based AI into an extension of our brains
4 days ago



Oculus co-founder Palmer Luckey donated \$100,000 to Trump's inauguration
5 days ago



FTC tells 'influencers' to quit trying to hide the fact that they're shilling for brands
5 days ago



Uber gets sued over alleged 'Hell' program to track Lyft drivers
a day ago

FaceApp apologizes for building a racist AI

Posted 45 minutes ago by [Natasha Lomas \(@riptari\)](#)



2017/04/25



If only all algorithmic bias were as easy to spot as this: FaceApp, a photo-editing app that [uses a neural network for editing selfies](#) in a photorealistic way, has apologized for building a racist algorithm.

The app lets users upload a selfie or a photo of a face, and offers a series of filters that can then be applied to the image to subtly or radically alter its appearance — its appearance-shifting effects include aging and even changing gender.

The problem is the app also included a so-called "hotness" filter, and this filter was racist. As [users pointed out](#), the filter was lightening skin tones to achieve its mooted "beautifying" effect. You can see the filter pictured above in a before and after shot of President Obama.

In an emailed statement apologizing for the racist algorithm, FaceApp's founder and CEO Yaroslav Goncharov told us: "We are deeply sorry for this unquestionably serious issue. It is an unfortunate side-effect of the underlying neural network caused by the training set bias, not intended behaviour. To mitigate the issue, we have renamed the effect to exclude any positive connotation associated with it. We are also working on the complete fix that should arrive soon."

NEWSLETTER SUBSCRIPTIONS

- The Daily Crunch**
Get the top tech stories of the day delivered to your inbox
- TC Weekly Roundup**
Get a weekly recap of the biggest tech stories
- Crunchbase Daily**
The latest startup funding announcements

Enter your email [SUBSCRIBE](#)

protected by reCAPTCHA
[Privacy Terms](#)



[SEE ALL NEWSLETTERS](#) >>

LATEST CRUNCH REPORT



Uber Responds to iPhone Tracking Report | Crunch Report


Lena and Fabio

Examples: Controversy and Appropriateness



‘Lena’

Alexander Sawchuk @ USC, 1973



‘Fabio’

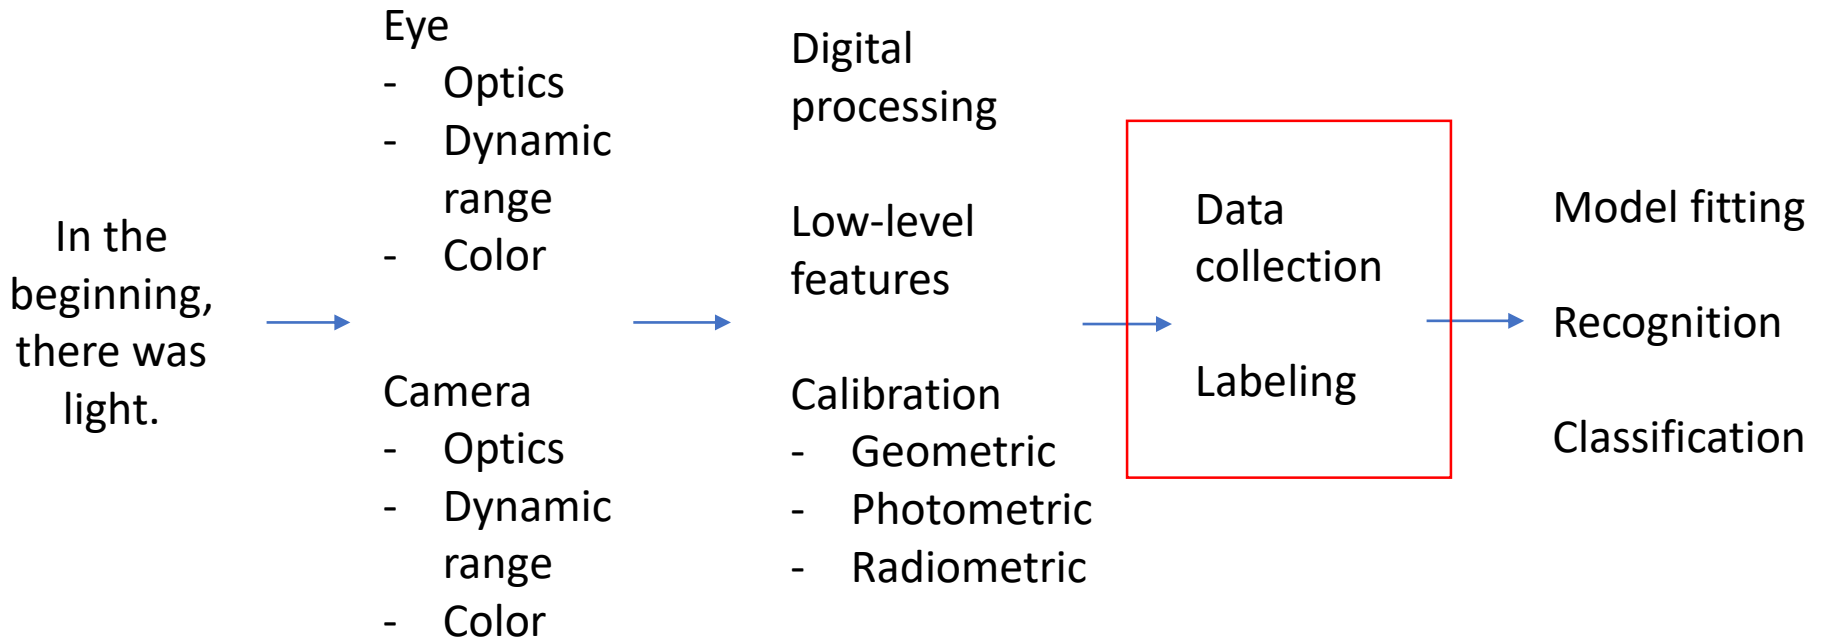
Deanna Needell @ Claremont McKenna, 2012

Wired article explaining it all


<https://www.wired.com/story/finding-lena-the-patron-saint-of-jpegs/>

Dataset Bias

Computer vision domain



Bias/variance trade-off



Bias = accuracy
Variance = precision

Unbiased Look at Dataset Bias

Torralba and Efros, CVPR 2011

“The authors would like to thank the Eyjafjallajokull volcano as well as the wonderful kirs at the Buvette in Jardin du Luxembourg for the motivation (former) and the inspiration (latter) to write this paper.”

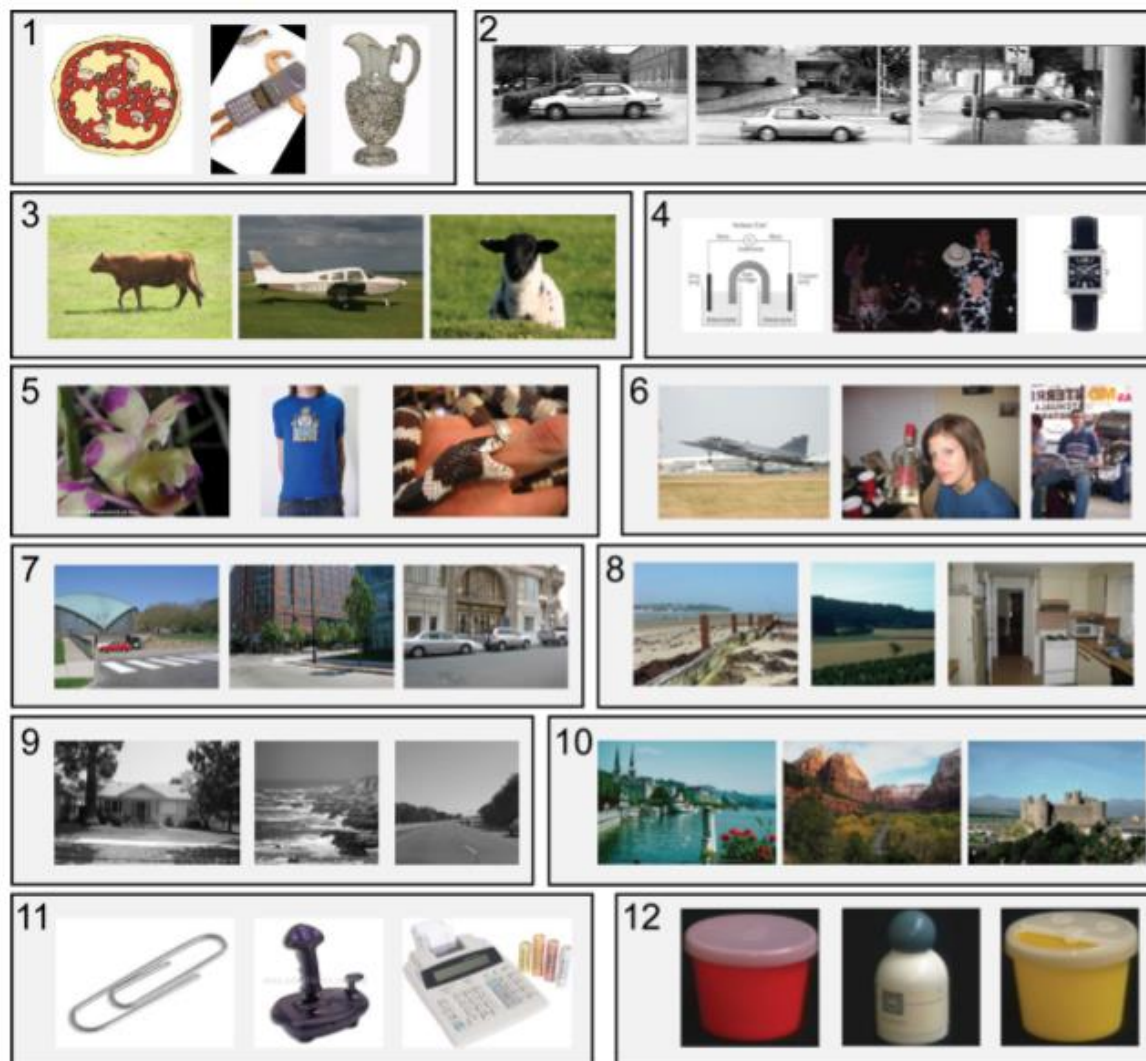
Next few slide contents are from the paper

Progression of dataset complexity

- COIL-100:



- 15 scenes: Out of the lab, backgrounds
- Caltech-101: Google-mined, single object in middle.
- LabelMe: Multiple objects, anywhere
- PASCAL VOC: More rigorous testing standards
- ImageNet: Internet-scale, real-world



Caltech101 Tiny LabelMe 15 Scenes
 MSRC Corel COIL-100 Caltech256
 UIUC PASCAL 07 ImageNet SUN09

Figure 1. Name That Dataset: Given three images from twelve popular object recognition datasets, can you match the images with the dataset? (answer key below)

PASCAI cars



SUN cars



Caltech101 cars



ImageNet cars



LabelMe cars



Figure 4. Most discriminative cars from 5 datasets

Measuring Dataset Bias

- Idea: cross-dataset generalization
- Train an object classifier on one dataset
- Test on the same object class on another dataset
- Observe performance as measure of bias

<i>task</i>	Test on:		SUN09	LabelMe	PASCAL	ImageNet	Caltech101	MSRC	Self	Mean others	Percent drop
	Train on:										
<i>“person” detection</i>	SUN09		69.6	56.8	37.9	45.7	52.1	72.7	69.6	53.0	24%
	LabelMe		58.9	66.6	38.4	43.1	57.9	68.9	66.6	53.4	20%
	PASCAL		56.0	55.6	56.3	55.6	56.8	74.8	56.3	59.8	-6%
	ImageNet		48.8	39.0	40.1	59.6	53.2	70.7	59.6	50.4	15%
	Caltech101		24.6	18.1	12.4	26.6	100	31.6	100	22.7	77%
	MSRC		33.8	18.2	30.9	20.8	69.5	74.7	74.7	34.6	54%
	Mean others		44.4	37.5	31.9	38.4	57.9	63.7	71.1	45.6	36%
<i>“person” classification</i>	SUN09		16.1	11.8	14.0	7.9	6.8	23.5	16.1	12.8	20%
	LabelMe		11.0	26.6	7.5	6.3	8.4	24.3	26.6	11.5	57%
	PASCAL		11.9	11.1	20.7	13.6	48.3	50.5	20.7	27.1	-31%
	ImageNet		8.9	11.1	11.8	20.7	76.7	61.0	20.7	33.9	-63%
	Caltech101		7.6	11.8	17.3	22.5	99.6	65.8	99.6	25.0	75%
	MSRC		9.4	15.5	15.3	15.3	93.4	78.4	78.4	29.8	62%
	Mean others		9.8	12.3	13.2	13.1	46.7	45.0	43.7	23.4	47%

Different kinds of bias

- *Selection bias*
 - Retrieve different kinds of images; keywords/search engines can bias.
- *Capture bias*
 - Objects photographed in similar ways that do not generalize, e.g., object always in center, race track car vs. street car, mugs.



mug



All

Shopping

Images

Maps

Videos

More

Settings

Tools



Different kinds of bias

- *Selection bias*
 - Retrieve different kinds of images; keywords/search engines can bias.
- *Capture bias*
 - Objects photographed in similar ways that do not generalize, e.g., object always in center, race track car vs. street car, mugs.
- *Category/label bias*
 - Poorly-defined classes, e.g., painting vs. picture
- *Negative set bias*
 - In one vs. all classification, 'all' or "the rest of the world" is not well represented.
 - "Are features which helps classify 'boat' object really the boat, or are they the water it sits on?"
 - Low bias negative set would include many boat-free images of rivers and lakes.

Measuring Negative Set Bias

- Take negative examples from other datasets and add to superset; train against this.

<i>task</i>			Positive Set:					Mean	
	Negative Set:		SUN09	LabelMe	PASCAL	ImageNet	Caltech101		MSRC
<i>“car” detection</i>	self		67.6	62.4	56.3	60.5	97.7	74.5	70.0
	all		53.8	51.3	47.1	65.2	97.7	70.0	64.1
	percent drop		20%	18%	16%	-8%	0%	6%	8%
<i>“person” detection</i>	self		67.4	68.6	53.8	60.4	100	76.7	71.1
	all		52.2	58.0	42.6	63.4	100	71.5	64.6
	percent drop		22%	15%	21%	-5%	0%	7%	9%

- Drop in performance of ‘all’ suggests negative examples are being misclassified

Overcoming bias at collection time

- *Selection bias*
 - Multiple keywords, search engines, countries.
 - Collect unknown images and label them by crowd-sourcing.
- *Capture bias*
 - Better sampling
 - Different transforms: noise, flips, rotations, affine, crops.

Overcoming bias at collection time

- *Category/label bias*

- Clear instruction to turkers; unambiguous classes (possible?)
- Pre-label clustering, or multiple acceptable answers.

- *Negative set bias*

- Cross-dataset mining
- Mine for hard negatives from unlabeled set using a reliable algorithm and high threshold.

Undoing the Damage of Dataset Bias

Khosla et al., ECCV 2012

“While it remains in question whether creating an unbiased dataset is possible given limited resources, we propose a discriminative framework that directly exploits dataset bias during training.”

More info

Kate Crawford @ NIPS 2017

https://www.youtube.com/watch?v=fMym_BKWQzk

The Trouble with Bias

More examples


<https://www.quora.com/What-are-examples-of-computer-vision-bugs-related-to-race>

<http://www.telegraph.co.uk/technology/2016/12/07/robot-passport-checker-rejects-asian-mans-photo-having-eyes/>

Viola-Jones with a bad training database



Google Photos (2015)



Google Photos (2015)

- What do you think the problem was?
- How could you fix it?

- Has it been fixed? Anyone use Google Photos?

Google Photos (2015)



(((Yonatan Zunger)))

@yonatanzunger



@jackyalcine Quick update: we shouldn't be making piles with that label anymore, and searches are mostly fixed, but they can still turn up.. [in]



(((Yonatan Zunger)))

@yonatanzunger



@jackyalcine ..photos where we failed to recognize that there was a face there at all. We're working on that issue now.

LIKES

8



11:17 AM - 29 Jun 2015



2



8

Not just a vision problem

Text embeddings also suffer:

<https://gist.github.com/rspeer/ef750e7e407e04894cb3b78a82d66aed>

‘Sentiment analysis’ ->

```
In [12]: text_to_sentiment("this example is pretty cool")
```

```
Out[12]: 3.889968926086298
```

```
In [13]: text_to_sentiment("this example is okay")
```

```
Out[13]: 2.7997773492425186
```

```
In [14]: text_to_sentiment("meh, this example sucks")
```

```
Out[14]: -1.1774475917460698
```

Not just a vision problem

Text embeddings also suffer:

<https://gist.github.com/rspeer/ef750e7e407e04894cb3b78a82d66aed>

‘Sentiment analysis’ ->

```
In [15]: text_to_sentiment("Let's go get Italian food")
```

```
Out[15]: 2.0429166109408983
```

```
In [16]: text_to_sentiment("Let's go get Chinese food")
```

```
Out[16]: 1.4094033658140972
```

```
In [17]: text_to_sentiment("Let's go get Mexican food")
```

```
Out[17]: 0.38801985560121732
```

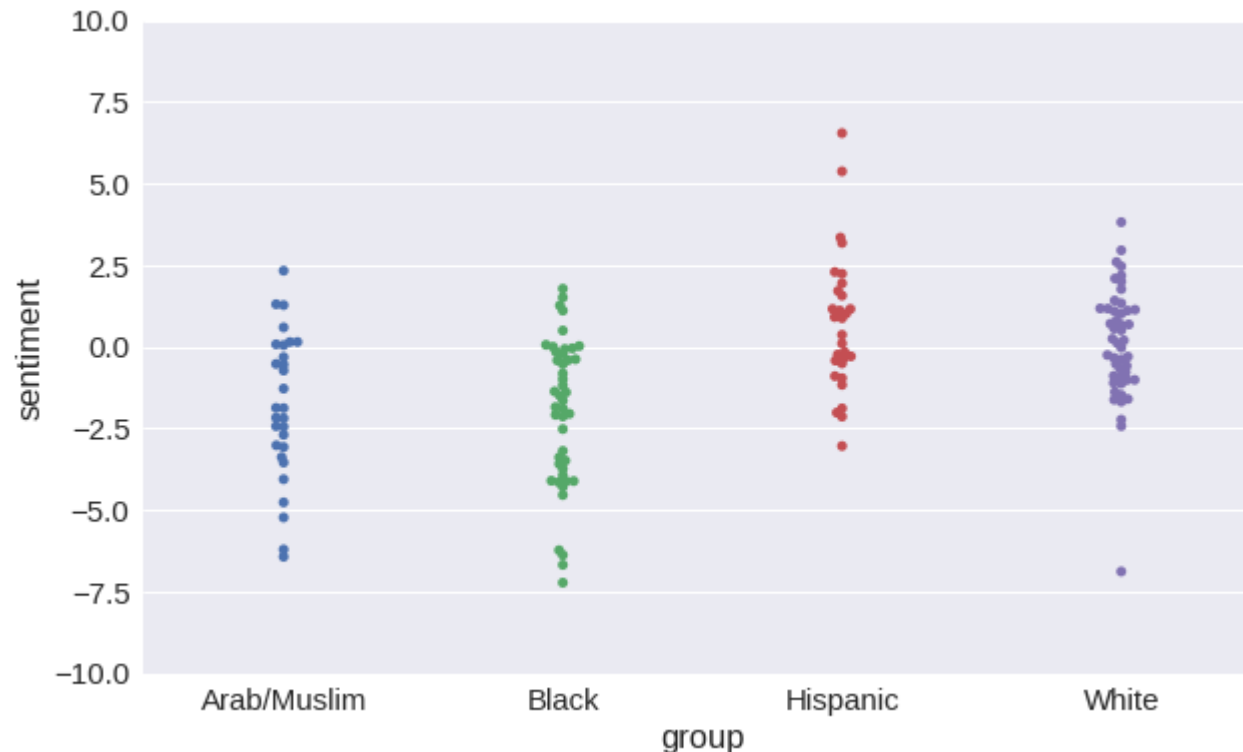
Word embedding trained on Google News – word2vec

```
In [20]: text_to_sentiment("My name is Yvette")
```

```
Out[20]: 0.98463802132985556
```

```
In [21]: text_to_sentiment("My name is Shaniqua")
```

```
Out[21]: -0.47048131775890656
```



AI 'Safety'

Concrete Problems in AI Safety

- <https://arxiv.org/abs/1606.06565>

In context of robots, but promising ideas

- Regularizer based on expert 'risk' of class confusion

Criminality

- Wu and Zhang, *Automated Inference on Criminality using Face Images*, on arXiv 2016



(a) Three samples in criminal ID photo set S_c .



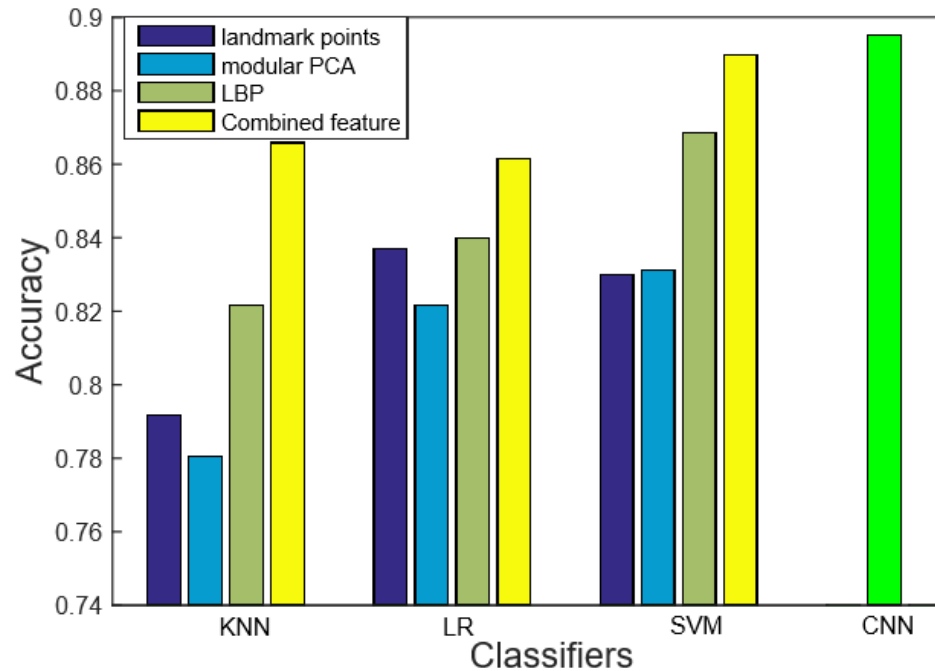
(b) Three samples in non-criminal ID photo set S_n

Figure 1. Sample ID photos in our data set.

“Unlike a human examiner/judge, a computer vision algorithm or classifier has absolutely no subjective baggages, having no emotions, no biases whatsoever due to past experience, race, religion, political doctrine, gender, age, etc., no mental fatigue, no preconditioning of a bad sleep or meal. The automated inference on criminality eliminates the variable of meta-accuracy (the competence of the human judge/examiner) all together.”

Criminality

- 1100 non-criminal, 730 criminal Chinese face photos
- Tested various features + classifiers



Criminality

K-means, averaging clusters

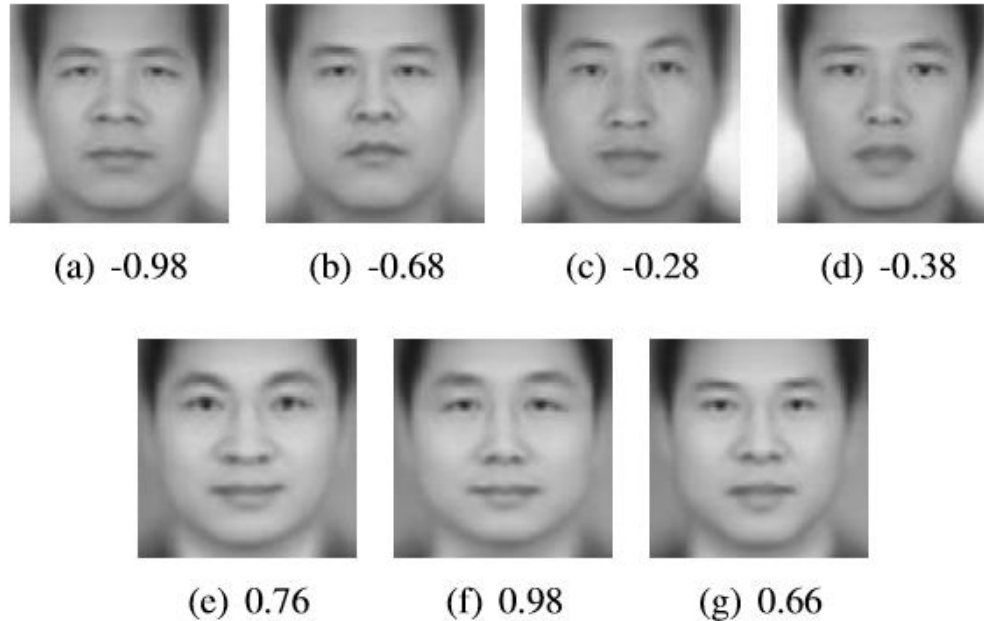


Figure 13. (a), (b), (c) and (d) are the four subtypes of criminal faces corresponding to four cluster centroids on the manifold of S_c ; (e), (f) and (g) are the three subtypes of non-criminal faces corresponding to three cluster centroids on the manifold of S_n . The number associated with each face is the average score of human judges (-1 for criminals; 1 for non-criminals).

What biases might exist? Discuss!

- Selection bias
- Capture bias
- Category/label bias
- Negative set bias

Is this real?

Whatever the case, it needs care! Significant ramifications.

Humans *might* be able to do this:

- *Small but statistically significant ability to tell criminal from non-criminal in photo.*

Valla, J., Williams, W., & Ceci, S. J. (2011).

The accuracy of inferences about criminality based on facial appearance.

Journal of Social, Evolutionary, and Cultural Psychology, 5(1), 66-91.

MIT Technology Review has a good overview:

<https://www.technologyreview.com/s/602955/neural-network-learns-to-identify-criminals-by-their-faces/>

1



2



3



4




5



6



Th
Fac
See



OUR CLASSIFIERS



High IQ



Academic Researcher



Professional Poker
Player



Terrorist

Utilizing advanced machine learning techniques we developed and continue to evolve an array of classifiers. These classifiers represent a certain persona, with a unique personality type, a collection of personality traits or behaviors. Our algorithms can score an individual according to their fit to these classifiers.

[Learn More>](#)

Individuals who could present a threat to public safety. According ...

[Click Here](#)



Facial analysis software is being used to predict sexuality and security risks

February 5, 2018

[Click Here](#)

ETHICS & FAIRNESS USAGE

The Faception Advantage: Privacy and Fairness

- **Decision Support System** - Provides actionable intelligence as predicted traits and behaviors of individuals
- **Operates anonymously** – no need for identity, phone, social media data
- **Objective** – not biased by race, class, gender, age
- We store only classifiers (like a mathematical equation), **no “big data”** on people



“Guns don't kill people, *people* kill people!”


“Machine learning doesn't kill people,
training data kills people!”

- *ML community, all the time.*

Dataset improvement: MS COCO



an elephant standing on top of a basket being held by a woman.
a woman standing holding a basket with an elephant in it.
a lady holding an elephant in a small basket.
a lady holds an elephant in a basket.
an elephant inside a basket lifted by a woman.



What is COCO?



COCO is a new image recognition, segmentation, and captioning dataset. COCO has several features:

- ✓ Object segmentation
- ✓ Recognition in Context
- ✓ Multiple objects per image
- ✓ More than 300,000 images
- ✓ More than 2 Million instances
- ✓ 80 object categories
- ✓ 5 captions per image
- ✓ Keypoints on 100,000 people

Decent Pew Overview on Big Picture

Rainie and Anderson *Code-Dependent: Pros and Cons of the Algorithm Age*

The screenshot shows the top portion of a web page from the Pew Research Center. At the top, there is a navigation bar with the text "NUMBERS, FACTS AND TRENDS SHAPING YOUR WORLD" on the left and "ABOUT | FOLLOW US" on the right, followed by a search box. Below this is the Pew Research Center logo and the subtitle "Internet, Science & Tech". A horizontal menu contains various categories: U.S. POLITICS, MEDIA & NEWS, SOCIAL TRENDS, RELIGION, INTERNET & TECH (which is highlighted), SCIENCE, HISPANICS, and GLOBAL. Below the menu are links for PUBLICATIONS, TOPICS, PRESENTATIONS, INTERACTIVES, FACT SHEETS, DATASETS, and EXPERTS. The main content area features the word "REPORT" and the date "FEBRUARY 8, 2017". The article title is "Code-Dependent: Pros and Cons of the Algorithm Age", with a sub-headline: "Algorithms are aimed at optimizing everything. They can save lives, make things easier and conquer chaos. Still, experts worry they can also put too much control in the hands of corporations and governments, perpetuate bias, create filter bubbles, cut choices, creativity and serendipity, and could result in greater unemployment". The authors are listed as "BY LEE RAINIE AND JANNA ANDERSON". Below the text is a code block showing C++ code for a Levenshtein distance algorithm. On the right side, there are sections for "REPORT MATERIALS" (with a link to "Complete Report PDF") and "TABLE OF CONTENTS" (with an "Overview" section listing six themes).

NUMBERS, FACTS AND TRENDS SHAPING YOUR WORLD ABOUT | FOLLOW US Search

PewResearchCenter *Internet, Science & Tech*

U.S. POLITICS MEDIA & NEWS SOCIAL TRENDS RELIGION INTERNET & TECH SCIENCE HISPANICS GLOBAL

PUBLICATIONS TOPICS PRESENTATIONS INTERACTIVES FACT SHEETS DATASETS EXPERTS

REPORT

FEBRUARY 8, 2017

Code-Dependent: Pros and Cons of the Algorithm Age

Algorithms are aimed at optimizing everything. They can save lives, make things easier and conquer chaos. Still, experts worry they can also put too much control in the hands of corporations and governments, perpetuate bias, create filter bubbles, cut choices, creativity and serendipity, and could result in greater unemployment

BY LEE RAINIE AND JANNA ANDERSON

```
19 template<typename T>
20 unsigned int levenshtein(const T& s1, const T& s2) {
21     const size_t len1 = s1.size(), prevCol(len2+1), i++;
22     vector<unsigned int> col(len2+1);
23     for (unsigned int i = 0; i < prevCol.size(); i++)
24         prevCol[i] = i;
25     for (unsigned int i = 0; i < len1; i++) {
26         col[0] = i+1;
27         for (unsigned int j = 0; j < len2; j++)
28             col[j+1] = std::min( std::min(prevCol[j] + j + 1, col[j] +
29                                     prevCol[j] + (s1[i]==s2[j] ? 0 : 1) ));
30         col.swap(prevCol);
31     }
32     return prevCol[len2];
33 }
```

REPORT MATERIALS

Complete Report PDF

TABLE OF CONTENTS

Overview

Themes illuminating concerns and challenges

Key experts' thinking about the future impacts of algorithms

About this canvassing of experts

Theme 1: Algorithms will continue to spread everywhere

Theme 2: Good things lie ahead

Theme 3: Humanity and human judgment are lost when data and predictive modeling become paramount

Theme 4: Biases exist in algorithmically-organized systems

Theme 5: Algorithmic categorizations deepen divides

Theme 6: Unemployment will rise

<http://www.pewinternet.org/2017/02/08/code-dependent-pros-and-cons-of-the-algorithm-age/>

Help Do Something About It

Joy Buolamwini

<https://www.theguardian.com/technology/2017/may/28/joy-buolamwini-when-algorithms-are-racist-facial-recognition-bias/>

Founded 'Algorithmic Justice League'

<https://www.ajlunited.org/>



Look at Joy's paper on face bias in commercial softwares

Amazon Is Pushing Facial Technology That a Study Says Could Be Biased

In new tests, Amazon's system had more difficulty identifying the gender of female and darker-skinned faces than similar services from IBM and Microsoft.

By [Natasha Singer](#)

Jan. 24, 2019



Over the last two years, Amazon has aggressively marketed its facial recognition technology [to police departments](#) and federal agencies as a service to help law enforcement identify suspects more quickly. It has done so as another tech giant, Microsoft, has [called on Congress to regulate](#) the technology, arguing that it is too risky for companies to oversee on their own.

Now a new study from researchers at the M.I.T. Media Lab has found that Amazon's system, Rekognition, had much more difficulty in telling the gender of female faces and of darker-skinned faces in photos than similar services from IBM and Microsoft. The results raise questions about potential bias that could hamper Amazon's drive to popularize the technology.

Effect of public shaming on face recognition databases

<https://www.engadget.com/2019/01/25/amazon-rekognition-facial-analysis-gender-race-bias-mit/>

The paper: http://www.aies-conference.com/wp-content/uploads/2019/01/AIES-19_paper_223.pdf

“However, Amazon has disputed the results of MIT's tests, which took place over the course of 2018. It argued the researchers hadn't used the current version of Rekognition, and said the gender identification test used facial analysis (which picks out faces from images and assigns generic attributes to them), rather than facial recognition, which looks for a match for a specific face. It noted that they are distinct software packages.”

"Using an up-to-date version of Amazon Rekognition with similar data downloaded from parliamentary websites and the Megaface dataset of [1 million] images, we found exactly zero false positive matches with the recommended 99 [percent] confidence threshold," Matt Wood, general manager of deep learning and AI at Amazon Web Services, [told](#) *VentureBeat*.

More info on the technical issues here:

<https://itif.org/publications/2019/01/27/note-press-facial-analysis-not-facial-recognition>

Predicting Financial Crime: Augmenting the Predictive Policing Arsenal

Brian Clifton¹, Sam Lavigne¹, and Francis Tseng¹

¹ The New Inquiry
<https://thenewinquiry.com/>

Abstract. Financial crime is a rampant but hidden threat. In spite of this, predictive policing systems disproportionately target “street crime” rather than white collar crime. This paper presents the White Collar Crime Early Warning System (WCCEWS), a white collar crime predictive model that uses random forest classifiers to identify high risk zones for incidents of financial crime.

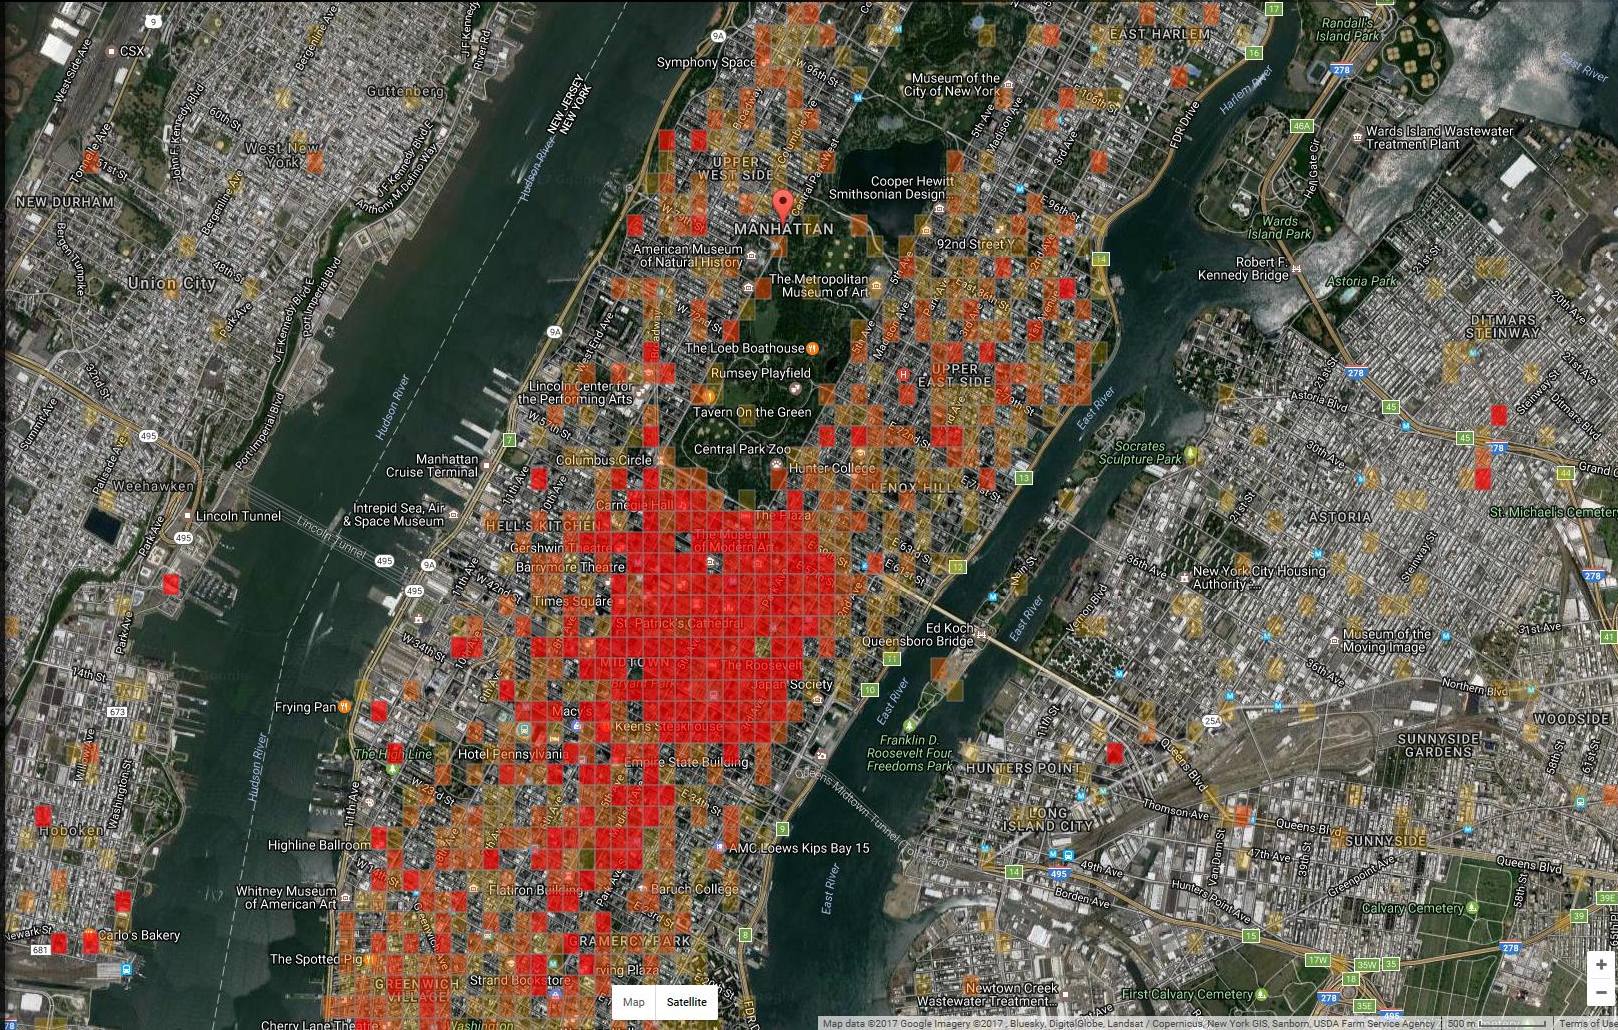
Keywords: Criminal justice; crime models; capitalism, financial malfeasance; white collar crime; police patrol.

WHITE COLLAR CRIME RISK ZONES

White Collar Crime Risk Zones uses machine learning to predict where financial crimes are mostly likely to occur across the US. To learn about our methodology, read our [white paper](#).

By Brian Clifton, Sam Lavigne and Francis Tseng for *The New Inquiry Magazine*, Vol. 59: **ABOLISH!**

THE NEW INQUIRY



Recently researchers have demonstrated the effectiveness of applying machine learning techniques to facial features to quantify the “criminality” of an individual²¹.

²¹ X. Wu and X. Zhang, “Automated inference on criminality using face images,” *CoRR*, vol. abs/1611.04135, 2016.




Figure 13. (a), (b), (c) and (d) are the four subtypes of criminal faces corresponding to four cluster centroids on the manifold of S_c ; (e), (f) and (g) are the three subtypes of non-criminal faces corresponding to three cluster centroids on the manifold of S_n . The number associated with each face is the average score of human judges (-1 for criminals; 1 for non-criminals).

Recently researchers have demonstrated the effectiveness of applying machine learning techniques to facial features to quantify the “criminality” of an individual²¹.

²¹ X. Wu and X. Zhang, “Automated inference on criminality using face images,” *CoRR*, vol. abs/1611.04135, 2016.

We therefore plan to augment our model with facial analysis and psychometrics to identify potential financial crime at the individual level. As a proof of concept, we have downloaded the pictures of 7000 corporate executives whose LinkedIn profiles suggest they work for financial organizations, and then averaged their faces to produce generalized white collar criminal subjects unique to each high risk zone. Future efforts will allow us to predict white collar criminality through real-time facial analysis.

Face detection + facial landmark detection +
image warping + averaging/PCA!



Fig. 7: Predicted White Collar Criminal for 40.7087811, -74.0064149

WHITE COLLAR CRIME RISK ZONES

White Collar Crime Risk Zones uses machine learning to predict where financial crimes are most likely to occur across the US. To learn about our methodology, read our [white paper](#).

By [Brian Clifton](#), [Sam Lavigne](#) and [Francis Tseng](#) for *The New Inquiry Magazine*, Vol. 59: **ABOLISH!**

Manhattan

Search

Most Likely Suspect



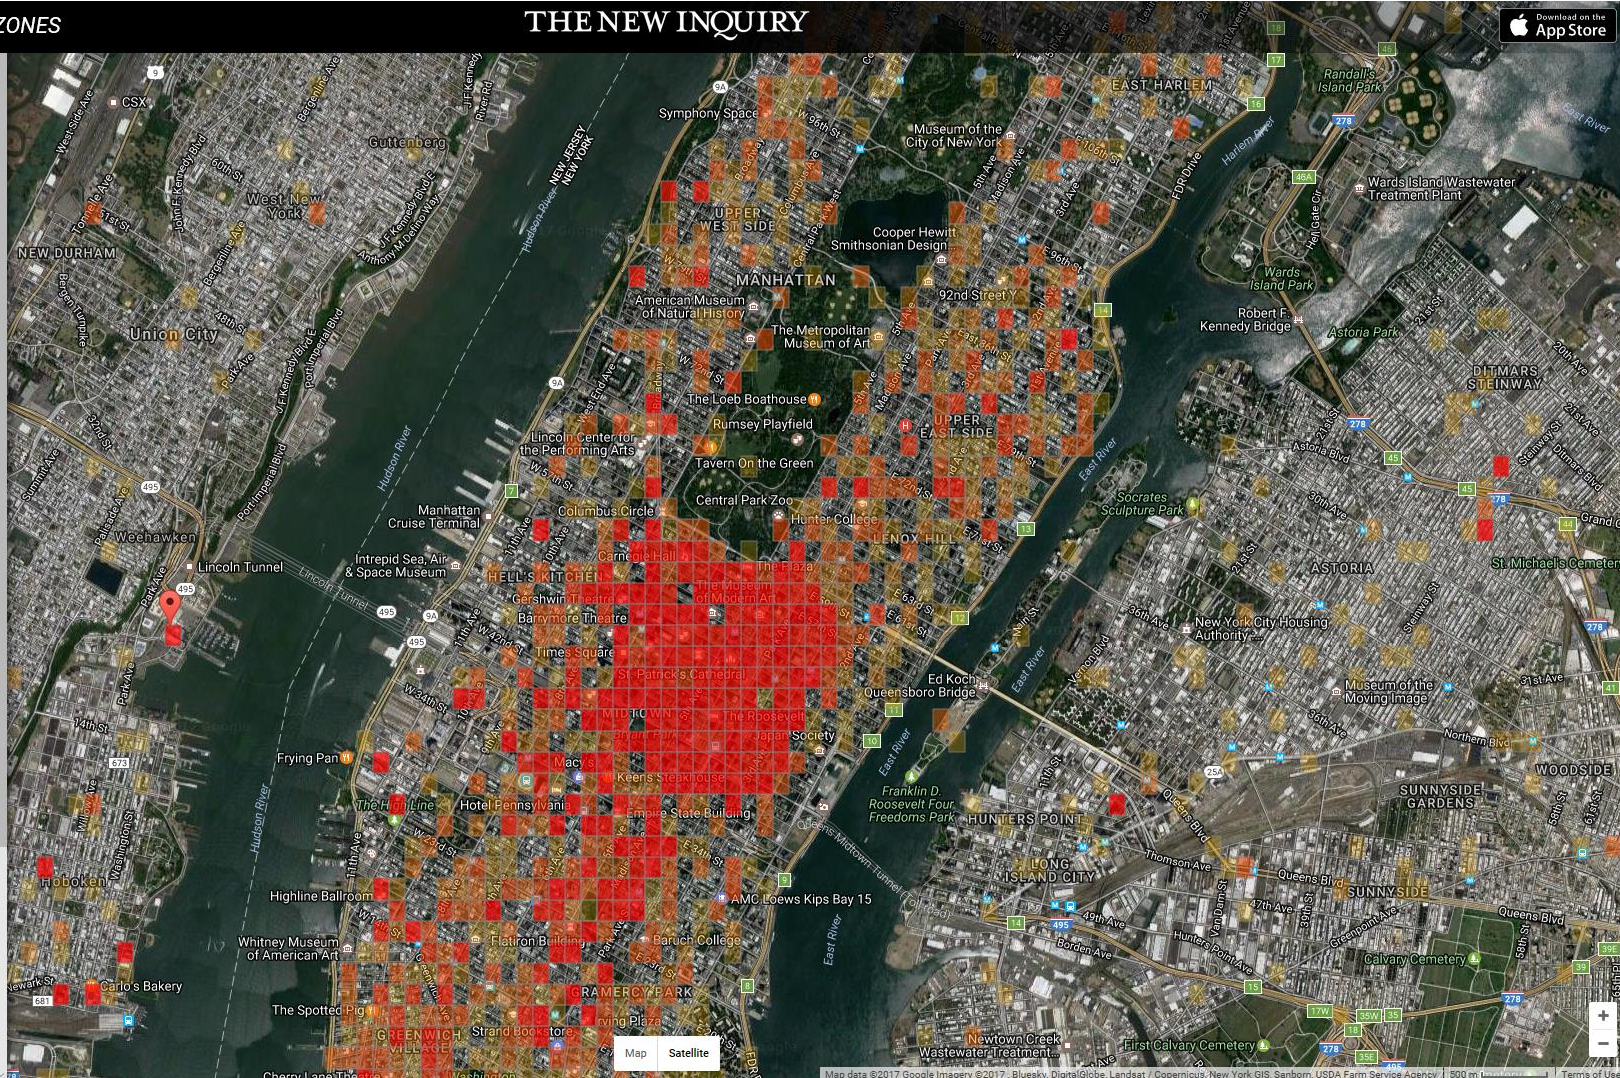
Top Risk Likelihoods

- FAILURE TO SUPERVISE (18.75%)
- BREACH OF FIDUCIARY DUTY (18.57%)
- EMPLOYMENT DISCRIMINATION BASED ON AGE (14.12%)

Approx. Crime Severity (in USD)



THE NEW INQUIRY



Map data ©2017 Google Imagery ©2017, Bluesky, DigitalGlobe, Landsat / Copernicus, New York GIS, Sanborn, USDA Farm Service Agency | 500 m | Terms of Use

[<https://whitecollar.thenewinquiry.com/>]

WHITE COLLAR CRIME RISK ZONES

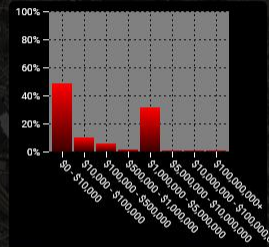
White Collar Crime Risk Zones uses machine learning to predict where financial crimes are mostly likely to occur across the US. To learn about our methodology, read our white paper.

By Brian Clifton, Sam Lavigne and Francis Tseng for *The New Inquiry Magazine*, Vol. 59: ABOLISH.

Top Risk Likelihoods

- BREACH OF FIDUCIARY DUTY (26.17%)**
- BUY IN TRADING DISPUTE (24.69%)**
- EMPLOYMENT DISCRIMINATION BASED ON AGE (18.6%)**

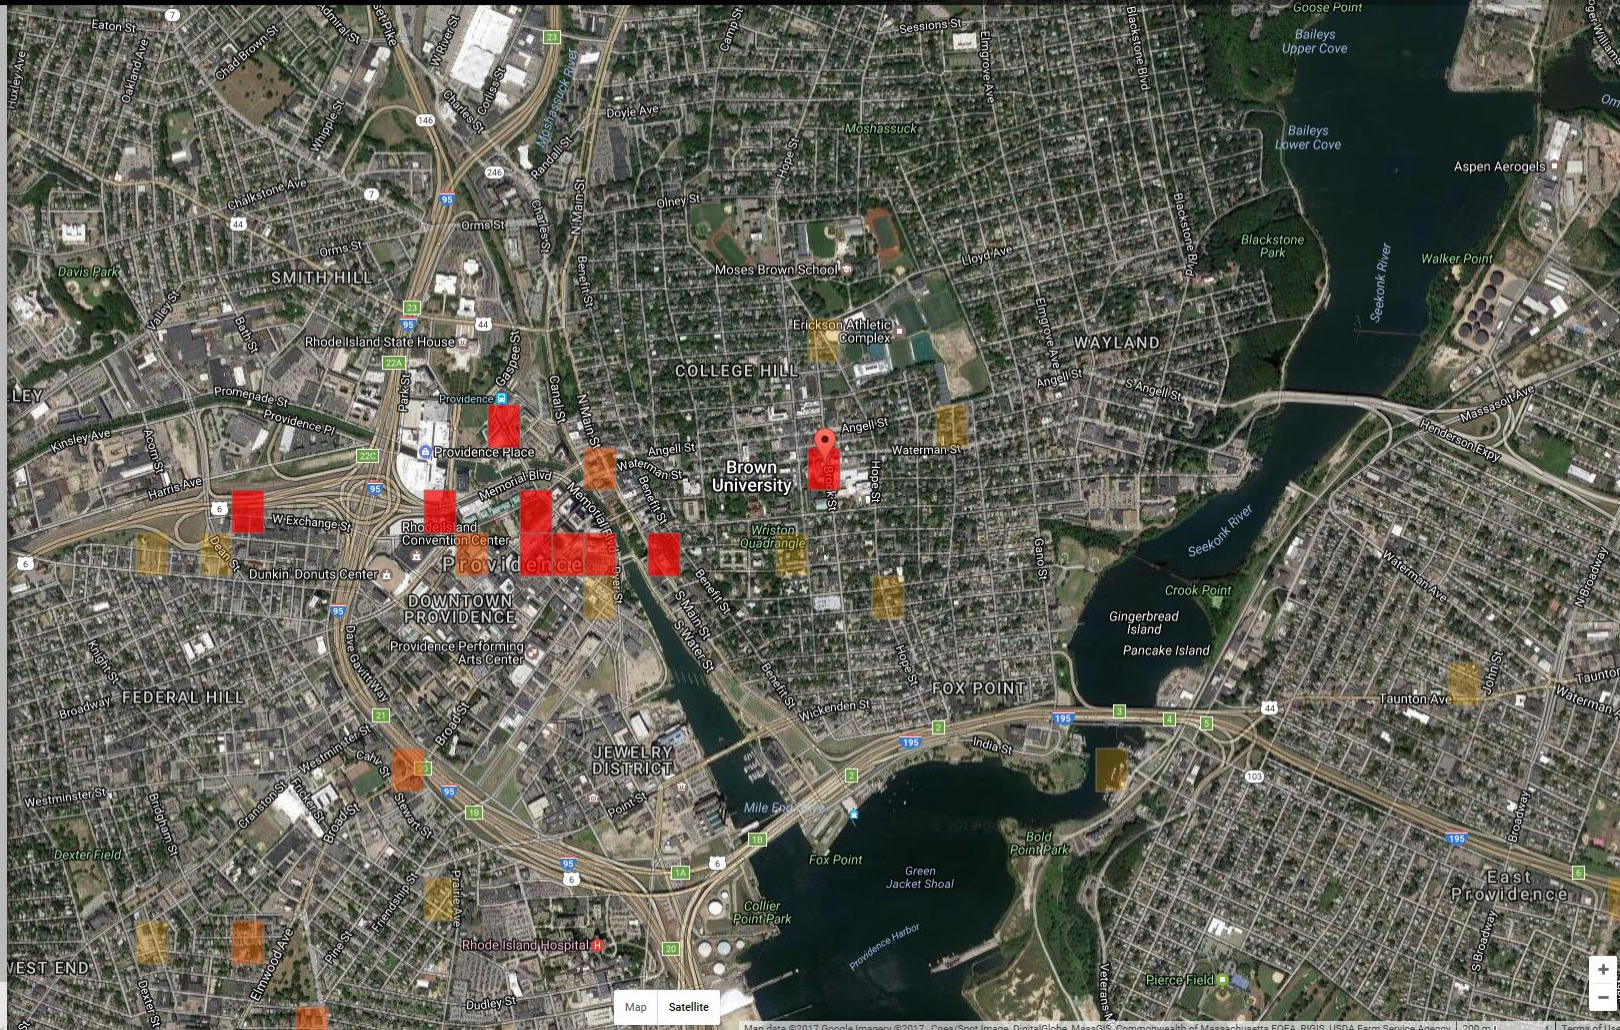
Approx. Crime Severity (in USD)



Nearby Financial Firms

- Citizens Bank
- Atlas ATM
- Santander Bank ATM
- Santander Bank
- ATM
- WRG Services Inc.

THE NEW INQUIRY



Map data ©2017 Google Imagery ©2017, Ches/Spot Image, DigitalGlobe, Maxar, Commonwealth of Massachusetts EOEA, RIGIS, USDA Farm Service Agency | 200 m | Terms of Use

Facebook translation (+image)

HAARETZ Friday, December 01, 2017. Kislev 13, 5778 Time in Israel: 9:55 PM

Israel News | All sections | Israel - Iran - Syria | Israel - Saudi Arabia | Avi Gabbay | Palestinian killed | Gaza - Israel

Israel Arrests Palestinian Because Facebook Translated 'Good Morning' to 'Attack Them'

No Arabic-speaking police officer read the post before arresting the man, who works at a construction site in a West Bank settlement

Yotam Berger | Oct 22, 2017 1:36 PM

6566 | Tweet | 3 | Zen



The Facebook post that mistranslated 'good morning' to 'hurt them'



Dataset openness

Rights on data

Later on could be used in previously inconceivable ways

E.g., Flickr scrape for Creative Commons images is legal

<https://www.theinquirer.net/inquirer/news/3072504/ibm-flicker-photos-facial-recognition-privacy-scandal>