


Deep Learning: What It Is and What It Can Do For You

diogomda@google.com 



ML Hipster

@ML_Hipster

 **Follow**

A machine learning researcher, a cryptocurrency expert, and an Erlang programmer walk into a bar. Facebook buys the bar for \$27 billion.

RETWEETS

5,873

LIKES

2,956



1:00 PM - 25 Feb 2014

 38

 5.9K

 3.0K



PHP CEO

@PHP_CEO

 Follow



HIRING FOR MACHINE LEARNING

DOES ANYONE KNOW SOMEONE CALLED
AL?

APPARENTLY WE REALLY NEED THIS GUY TO
GET ANYTHING DONE

RETWEETS

141

LIKES

195



1:25 PM - 22 Jan 2016



6



141



195

Conclusions

Deep neural networks and machine learning are starting to produce significant breakthroughs in healthcare and basic science

If you're not considering how to use deep neural nets to solve your problems, **you almost certainly should be**















What is Deep Learning?

Neural Networks


A mostly complete chart of architectures

-  Backfed Input Cell
-  Input Cell
-  Noisy Input Cell
-  Hidden Cell
-  Probabilistic Hidden Cell
-  Spiking Hidden Cell
-  Output Cell
-  Match Input Output Cell
-  Recurrent Cell
-  Memory Cell
-  Open Memory Cell
-  Scanning Filter
-  Convolution



Calculus

SILVANUS P. THOMPSON
AND MARTIN GARDNER
**CALCULUS
MADE EASY**



The first complete revision in over 75 years of the million-copy bestseller—including more than 20 new problems

CHAPTER I.

TO DELIVER YOU FROM THE PRELIMINARY TERRORS.

THE preliminary terror, which chokes off most fifth-form boys from even attempting to learn how to calculate, can be abolished once for all by simply stating what is the meaning—in common-sense terms—of the two principal symbols that are used in calculating.

These dreadful symbols are:

(1) d which merely means “a little bit of.”

Thus dx means a little bit of x ; or du means a little bit of u . Ordinary mathematicians think it more polite to say “an element of,” instead of “a little bit of.” Just as you please. But you will find that these little bits (or elements) may be considered to be indefinitely small.

(2) \int which is merely a long S , and may be called (if you like) “the sum of.”

Thus $\int dx$ means the sum of all the little bits of x ; or $\int dt$ means the sum of all the little bits of t . Ordinary mathematicians call this symbol “the

C.M.E.

A

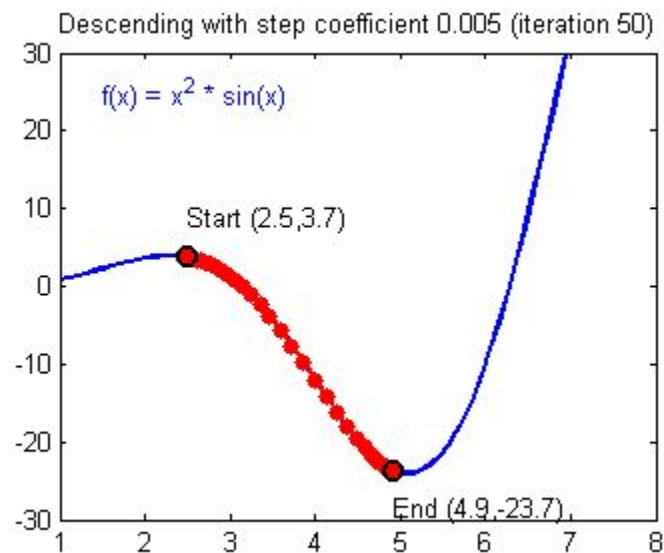
2 CALCULUS MADE EASY

integral of.” Now any fool can see that if x is considered as made up of a lot of little bits, each of which is called dx , if you add them all up together you get the sum of all the dx 's, (which is the same thing as the whole of x). The word “integral” simply means “the whole.” If you think of the duration of time for one hour, you may (if you like) think of it as cut up into 3600 little bits called seconds. The whole of the 3600 little bits added up together make one hour.

When you see an expression that begins with this terrifying symbol, you will henceforth know that it is put there merely to give you instructions that you are now to perform the operation (if you can) of totalling up all the little bits that are indicated by the symbols that follow.

That's all.

Gradient Descent





ML Hipster

@ML_Hipster

 Follow



Machine learning.

$$\Delta w_{ij} = -\alpha \frac{\partial E}{\partial w_{ij}}$$

RETWEETS

23

LIKES

76



4:54 PM - 16 Feb 2016




2



23




76



IS THIS A
CAT or DOG?



CAT DOG



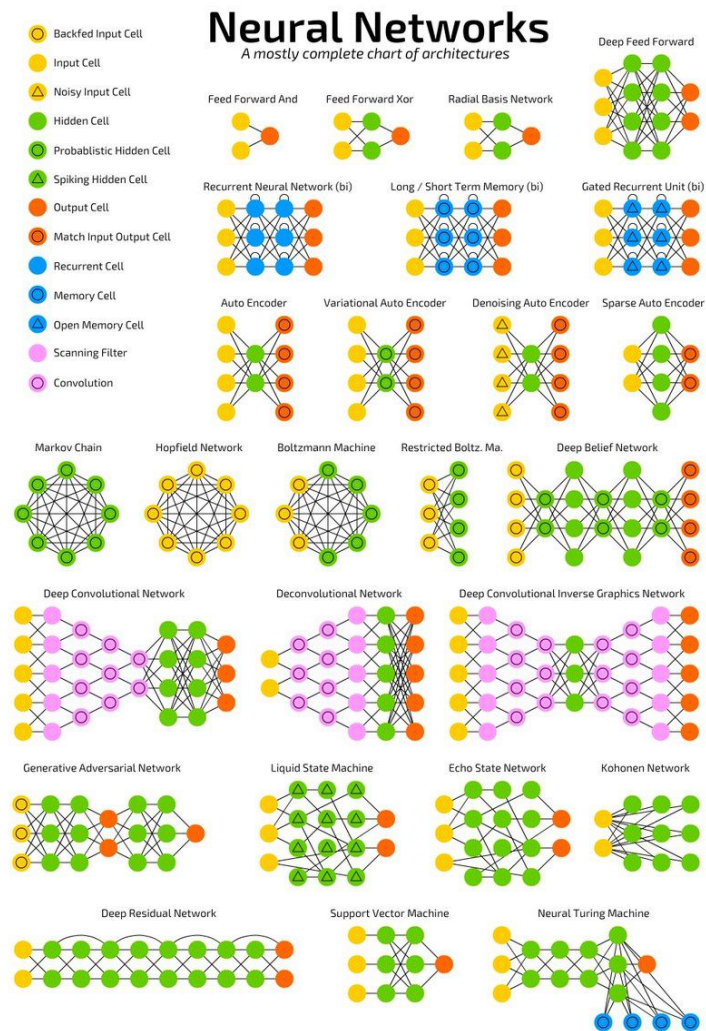
Is that it?

Is that it?



Is that it?

TensorFlow



Is that it?



What can it do for you?

When Will AI Exceed Human Performance? Evidence from AI Experts

Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, Owain Evans

(Submitted on 24 May 2017 (v1), last revised 30 May 2017 (this version, v2))

Advances in artificial intelligence (AI) will transform modern life by reshaping transportation, health, science, finance, and the military. To adapt public policy, we need to better anticipate these advances. Here we report the results from a large survey of machine learning researchers on their beliefs about progress in AI. Researchers predict AI will outperform humans in many activities in the next ten years, such as translating languages (by 2024), writing high-school essays (by 2026), driving a truck (by 2027), working in retail (by 2031), writing a bestselling book (by 2049), and working as a surgeon (by 2053). Researchers believe there is a 50% chance of AI outperforming humans in all tasks in 45 years and of automating all human jobs in 120 years, with Asian respondents expecting these dates much sooner than North Americans. These results will inform discussion amongst researchers and policymakers about anticipating and managing trends in AI.





Figure 2: Timeline of Median Estimates (with 50% intervals) for AI Achieving Human Performance. Timelines showing 50% probability intervals for achieving selected AI milestones. Specifically,

“If a typical person can do a mental task with less than one second of thought, **we can probably automate it using AI,** either now or in the near future.”

— **Andrew Ng**, chief scientist at Baidu



Andrej Karpathy, Machine Learning PhD student at Stanford

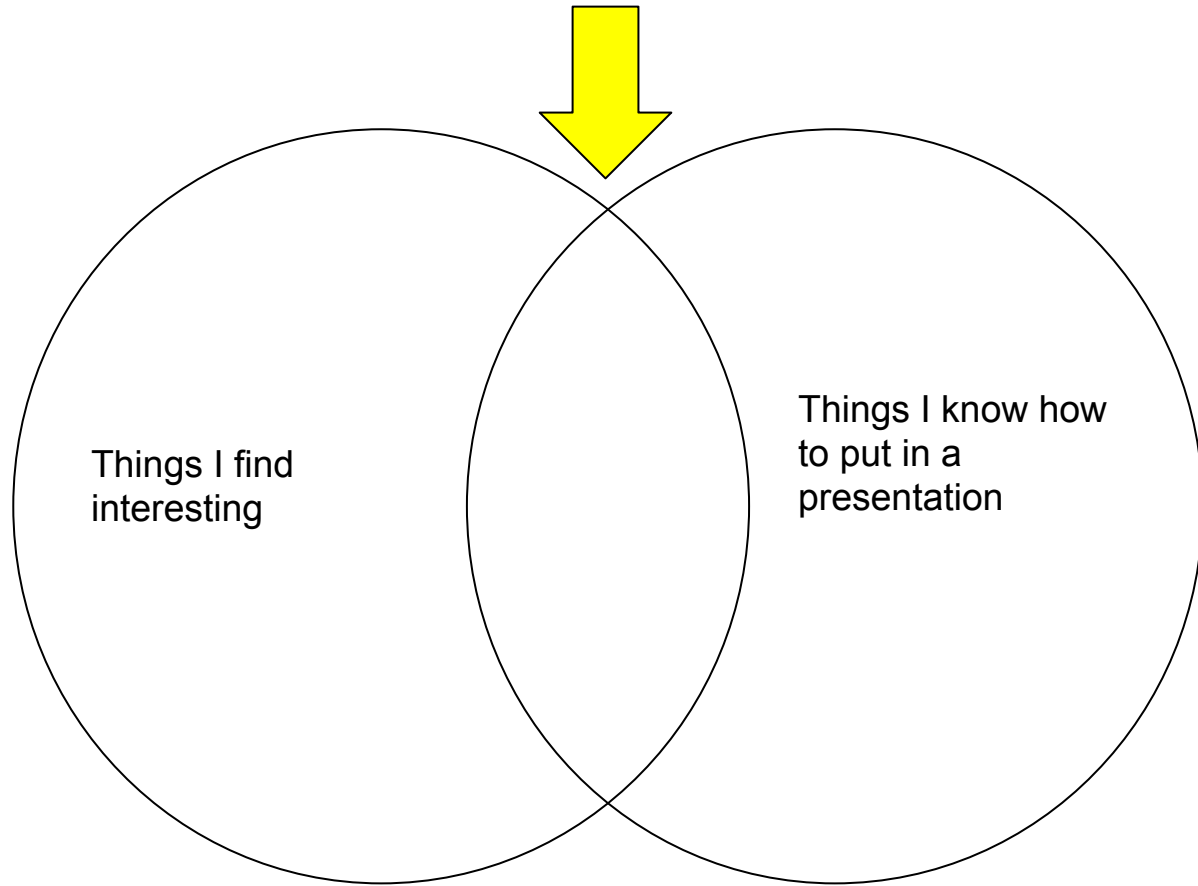
Written Sep 8 · Upvoted by Tao Xu, [Built ML systems at Airbnb, Quora, Facebook and Microsoft.](#) and Xavier Amatriain, [Former researcher, now leading ML and engineering teams](#)

I haven't found a way to properly articulate this yet but somehow everything we do in deep learning is **memorization** (interpolation, pattern recognition, etc) instead of **thinking** (extrapolation, induction, etc). I haven't seen a single compelling example of a neural network that I would say "thinks", in a very abstract and hard-to-define feeling of what properties that would have and what that would look like.

My general rule: deep learning is an appropriate tool for supervised direct pattern matching tasks. (Bonus points if you can design priors that are particularly suited to your problem.)



Let's talk specifics



Let's talk specifics

Very Yes: \$\$\$ goes up when some metric goes up



bradford cross

@bradfordcross

 Follow



standard results for a bread and butter AI use case: optimize an existing product or service like google, facebook, amazon and netflix do.

VentureBeat  @VentureBeat

DoorDash sees 25% lift from AI recommendations onvb.co/LuR3tBH

RETWEETS

12

LIKES

31



8:27 PM - 17 May 2017



12



31

No: Unsupervised Learning



Richard

@RichardSocher

 **Follow**



Rather than spending a month figuring out an unsupervised machine learning problem, just label some data for a week and train a classifier.

RETWEETS

291

LIKES

606



2:47 PM - 10 Mar 2017



18




291



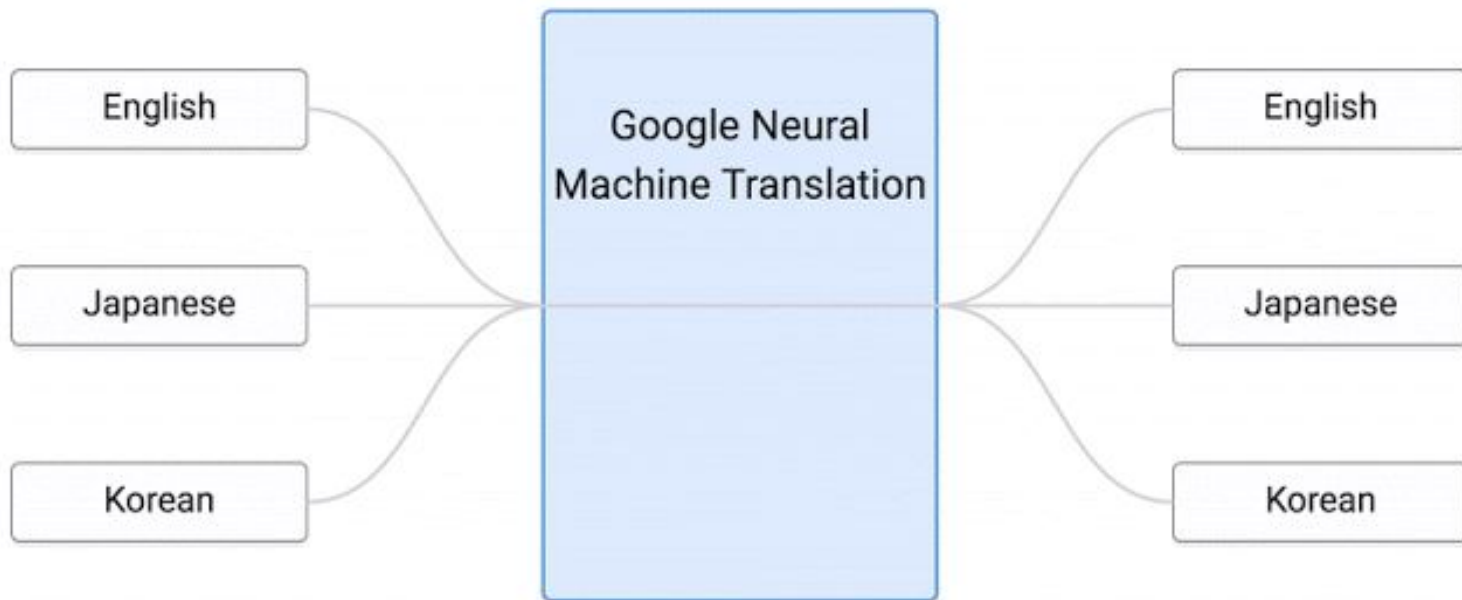
606

Yes: Speech Recognition




Yes: Translation

Training



Yes: Image Classification

Things



Google

my photos of siamese cats



Web

Images

Shopping

Videos

More ▾



Your photos

Only you can see these results



Not Yet: Unbiased Image Classification



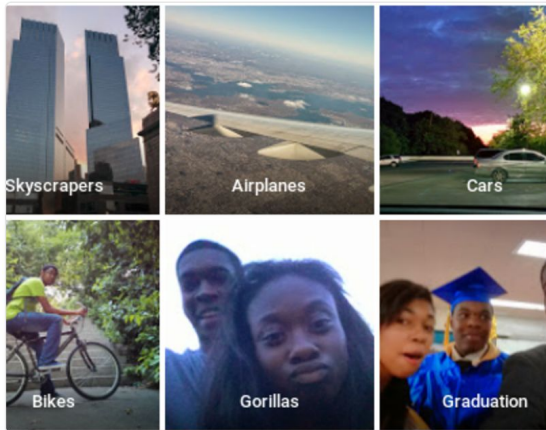
olufemi nzame

@jackyalcine

Follow



Google Photos, y'all fucked up. My friend's not a gorilla.



RETWEETS

3,181

LIKES

1,996



6:22 PM - 28 Jun 2015



224




3.2K



2.0K

Yes: Medical Imaging



December 13, 2016


Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs

Varun Gulshan, PhD¹; Lily Peng, MD, PhD¹; Marc Coram, PhD¹; [et al](#)


» [Author Affiliations](#)

JAMA. 2016;316(22):2402-2410. doi:10.1001/jama.2016.17216


Estimated
Malignancy:
93.7%




Malignant
92% similarity
28.7mm diameter
Medium spiculation




Malignant
95% similarity
32.8mm diameter
Medium spiculation



Malignant
93% similarity
26.4mm diameter
High spiculation




Malignant
93% similarity
36.1mm diameter
Low spiculation



Malignant
91% similarity
33.4mm diameter
Very High spiculation

Yes: Semantic Segmentation / 3D



Not Yet: Image Captioning




Human: A young girl asleep on the sofa cuddling a stuffed bear.

Model: A close up of a child holding a stuffed animal.

Model: A baby is asleep next to a teddy bear.

Not Yet: Question Answering



What is in the sheep's ear?

```
(describe[what]
 (and find[sheep]
       find[ear]))
```

tag



What color is she wearing?

```
(describe[color]
 find[wear])
```

white



What is the man dragging?

```
(describe[what]
 find[man])
```

boat (board)

No: Research



Andrej Karpathy ✓

@karpathy

 Follow



Coworker on RL research: "We were supposed to make AI do all the work and we play games but we do all the work and the AI is playing games!"

RETWEETS

838

LIKES

1,392



7:02 PM - 6 Oct 2016



12



838



1.4K

No: Research



Andrej Karpathy ✓
@karpathy

Follow

Coworker on RL research: "We were supposed to make AI do all the work and we play games but we do all the work and the AI is playing games!"

RETWEETS
838

LIKES
1,392



7:02 PM - 6 Oct 2016



12



838



1.4K



kamalikac @kamalikac · 6 Oct 2016

Replying to @karpathy

that's how you know the AI is really intelligent :-)



1



22



Otso Rajala @otjura · 7 Oct 2016

Replying to @karpathy

The greatest irony is that when AI systems replace labour little by little, the last ones to stop working are AI researchers.



3



3

Maybe: “Boring” Parts of Research

NEURAL ARCHITECTURE SEARCH WITH REINFORCEMENT LEARNING

Barret Zoph*, Quoc V. Le

Google Brain

{barretzoph, qvl}@google.com



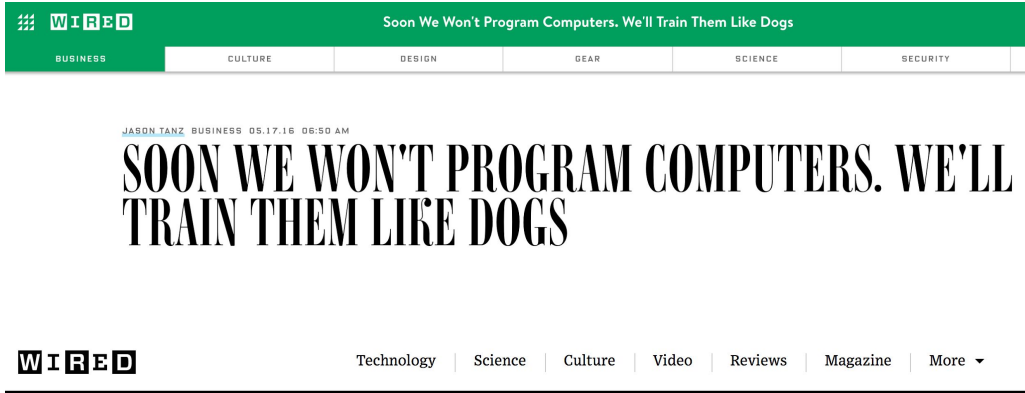


Figure 7: Convolutional architecture discovered by our method, when the search space does not have strides or pooling layers. FH is filter height, FW is filter width and N is number of filters.

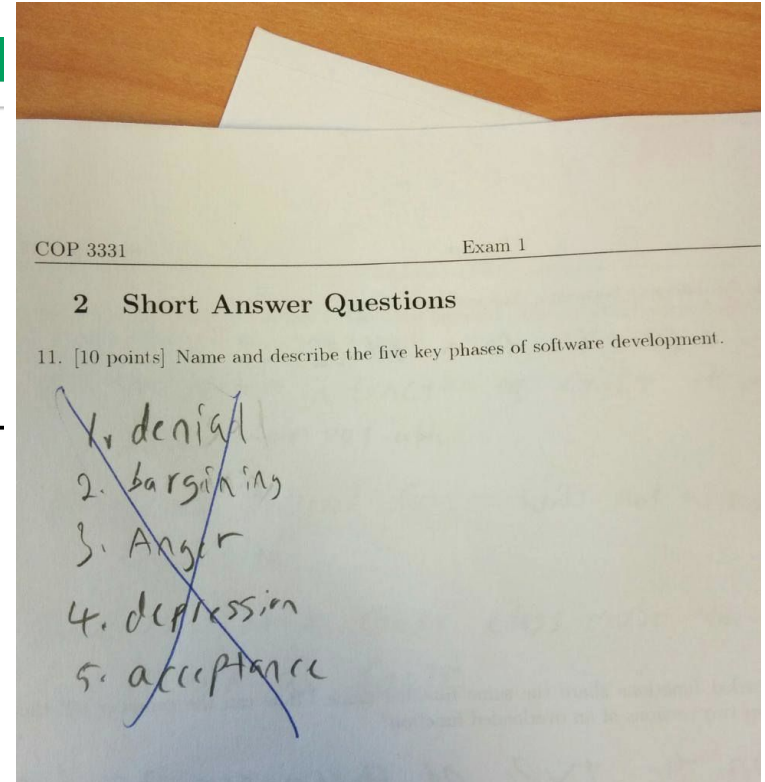
No: Software Development



Artificial Intelligence

Microsoft's AI writes code by looting other software

Developed by Microsoft and the University of Cambridge, DeepCoder takes lines of code from existing software



Very Yes: Memes



ML Hipster

@ML_Hipster

[Follow](#)

It's pretty clear what he meant to say was
"Despite the constant negative press $\text{Cov}(\phi, \phi)$
is always positive semi-definite." [#covfefe](#)

RETWEETS

403

LIKES

820



9:43 PM - 30 May 2017



5



403



820

Search or Article ID inside arXiv

All papers



Broaden your search using

[\(Help | Advanced search\)](#)

Computer Science > Learning

CONVolutional Feature Filtering Engine (CONVFefe): a neural approach to large-scale feature selection

Timotei Erwin, Andrei Keefe, Alois Bhaskara, Pejo Jitender

(Submitted on 30 May 2017)

Download

- PDF
- Other formats

(license)

Current version

cs.LG

< [prev](#) |

No: Product Definition, Not Yet: General Chatbots




David J Bland

@davidjbland

 Follow



I've analyzed the Silicon Valley Bot Startup trend and created a handy venn diagram to help explain it.



RETWEETS

5,880

LIKES

6,621



5:27 PM - 26 Apr 2016



85



5.9K



6.6K

Yes: Specific Chatbots

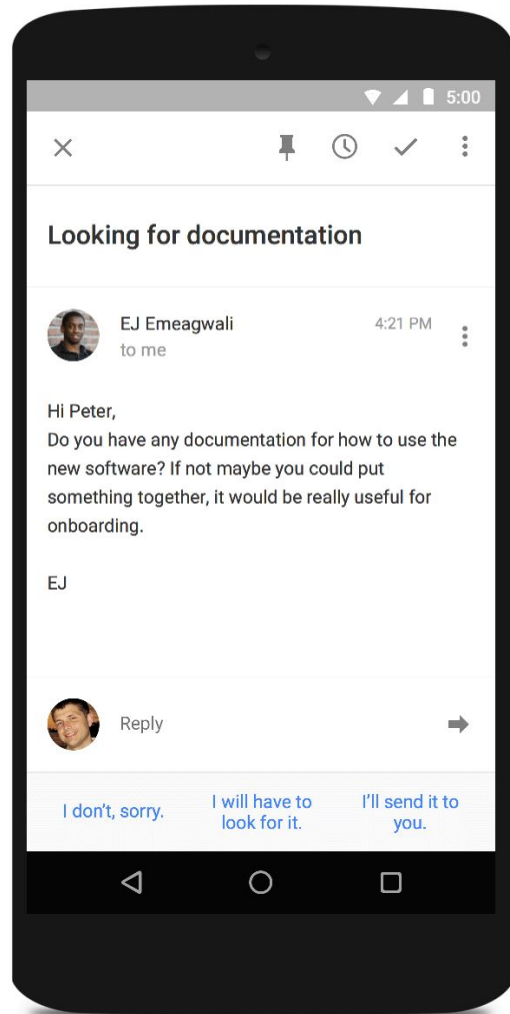


Smart Reply

April 1, 2009: April Fool's Day joke

Nov 5, 2015: Launched Real Product

Feb 1, 2016: >10% of mobile Inbox replies



Not Yet: Coherent Text



Smerity
@Smerity

 **Follow**



Today's AI isn't vaguely close to writing multiple sentences coherently, let alone articles with dependencies..?

Keith Rabois @rabois

Wait until journalists learn that today's AI is close to writing news stories on par or better than they do.

RETWEETS

42

LIKES

145



2:57 PM - 12 Nov 2016 from **San Francisco, CA**



16



42




145

Maybe: Robots



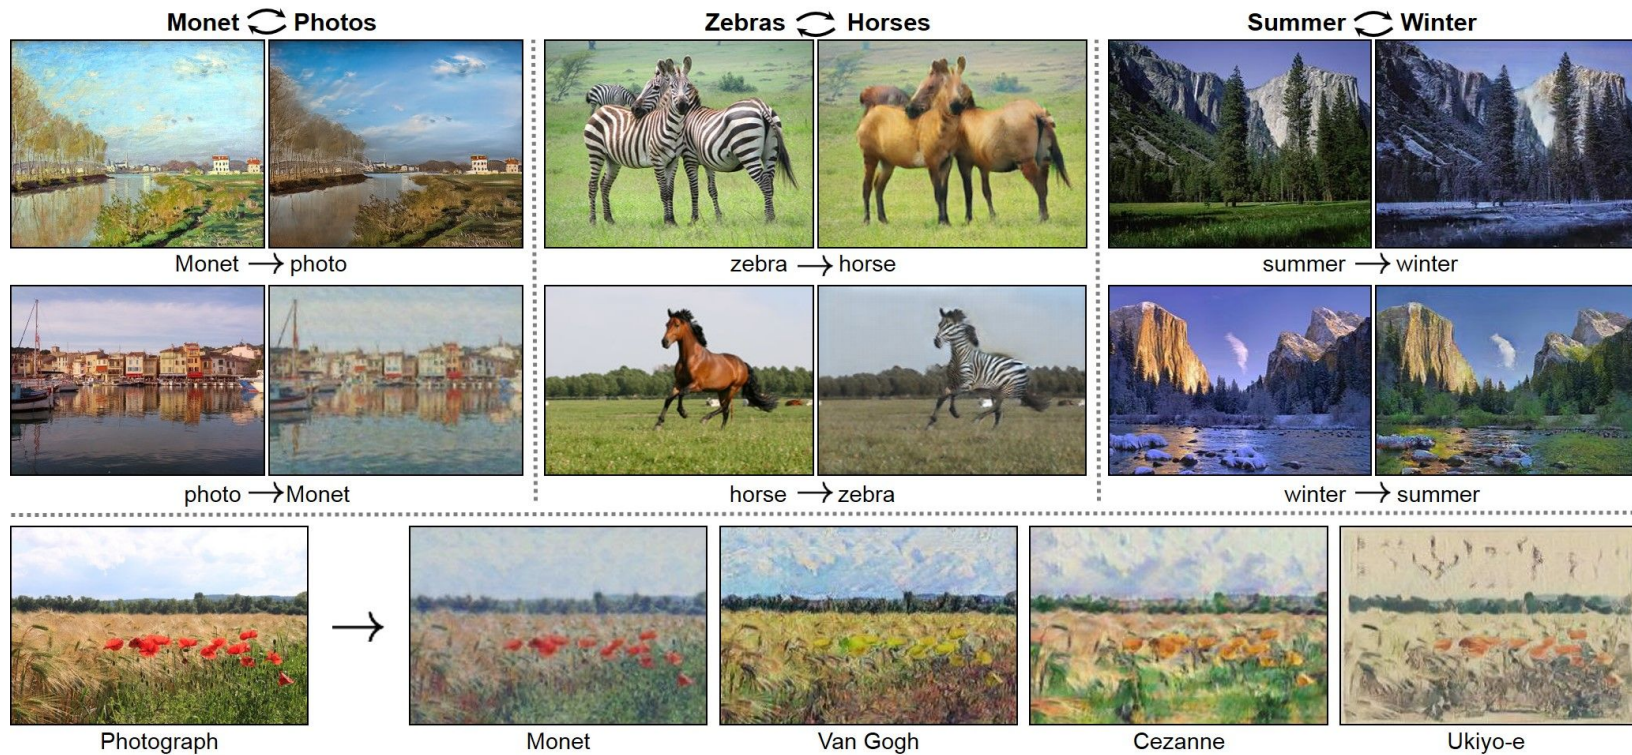
Not Yet: Game Playing



Not Yet: Game Playing



Yes: Image Transformation



Yes: Image Transformation



Not Yet: Reliable Image Transformation





Yes: Face Transformation





Not Yet: Unbiased Face Transformation



Yes: Art From Scratch




Yes: Sketching






Yes: Conditional Art

Synthesized Image

#NeuralDoodle




Yes: Style Transfer



Very Yes: Pokemon Style Transfer



+



=



Yes: All Sorts of Other Things

What Machine Learning Can Do

A simple way to think about supervised learning.

INPUT A	RESPONSE B	APPLICATION
Picture	Are there human faces? (0 or 1)	Photo tagging
Loan application	Will they repay the loan? (0 or 1)	Loan approvals
Ad plus user information	Will user click on ad? (0 or 1)	Targeted online ads
Audio clip	Transcript of audio clip	Speech recognition
English sentence	French sentence	Language translation
Sensors from hard disk, plane engine, etc.	Is it about to fail?	Preventive maintenance
Car camera and other sensors	Position of other cars	Self-driving cars

Back To The Big Questions

How Will The World Change?



Andrew Ng ✓

@AndrewYNg

 Follow



If you're trying to understand AI's near-term impact, don't think "sentience." Instead think "automation on steroids."

RETWEETS

1,287

LIKES

1,888



11:04 AM - 1 May 2017 from [Stanford, CA](#)



42



1.3K



1.9K

How Will The World Change?

David Autor:

Will automation take away all our jobs?

TEDxCambridge · 18:37 · Filmed Sep 2016

18 subtitle languages ?

View interactive transcript



Share this idea



Facebook



LinkedIn



Twitter



Link



Email



Embed

1,149,161

Total views

What Can Deep Learning Do?



Andrew Ng ✓

@AndrewYNg

 **Follow**



Cool thing about being in AI: You don't just work on AI. You get to do healthcare, robotics, logistics, finance, media,

RETWEETS

561

LIKES

1,311



9:57 AM - 27 Jan 2017 from [Sunnyvale, CA](#)



25



561



1.3K

What Can Deep Learning Do?

In a relationship
Engaged
Married
✓ It's complicated
Divorced

How Do I Take Advantage Of These Trends?



Join the Team

Full Time Roles

We're looking for talented research scientists and software engineers enthusiastic about deep learning to join us.

[VIEW JOBS](#)

Brain Residency

This 12-month program is designed to jumpstart your career in deep learning, working with our scientists and engineers from the Google Brain Team.

[VIEW RESIDENCY](#)

Visiting Faculty

Visiting Faculty work closely with our scientists and engineers, and have the opportunity to explore projects at industrial scale with state-of-the-art technology.

[VIEW VISITING FACULTY](#)

Interns

Our interns work on projects utilizing the latest techniques in deep learning. In your application, indicate your research interests in the 'Cover letter/other notes' section, so it can be routed to the appropriate recruiter.

[VIEW INTERNSHIPS](#)

Thanks!