What Can Machine Learning Do For You?

Adam Prügel-Bennett

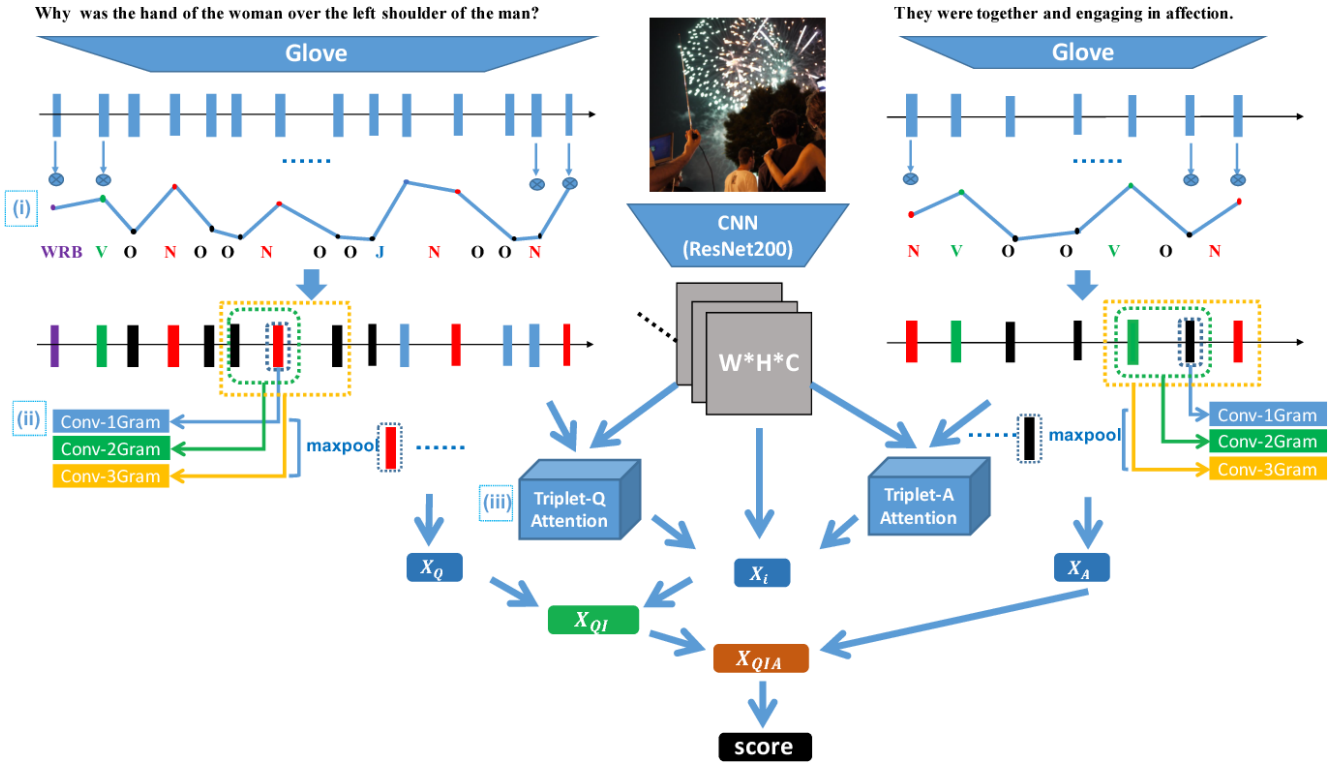

Machine Learning and AI

- Machine learning has been the driving force in the current revolution in artificial intelligence
- The last decade has seen an unprecedented stride forward in machine learning due to the development of deep learning
  - ⋆ Super-human classification performance
  - ⋆ Beats humans at Go
  - ⋆ Mind blowing language models
- How can you use machine learning to build a super intelligent system that will revolutionise your field?

Machine Learning and AI

- Machine learning has been the driving force in the current revolution in artificial intelligence
- The last decade has seen an unprecedented stride forward in machine learning due to the development of deep learning
  - ⋆ Super-human classification performance
  - ⋆ Beats humans at Go
  - ⋆ Mind blowing language models
- How can you use machine learning to build a super intelligent system that will revolutionise your field?

Machine Learning and AI

- Machine learning has been the driving force in the current revolution in artificial intelligence
- The last decade has seen an unprecedented stride forward in machine learning due to the development of deep learning
  - ⋆ Super-human classification performance
  - ⋆ Beats humans at Go
  - ⋆ Mind blowing language models
- How can you use machine learning to build a super intelligent system that will revolutionise your field?

You Can't!

- Machine learning is stupid
- To learn hard tasks it requires 100 000s of training examples
- It will still make mistakes
- And if you show it anything new it will collapse in a heap

You Can't!

- Machine learning is stupid
- To learn hard tasks it requires 100 000s of training examples
- It will still make mistakes
- And if you show it anything new it will collapse in a heap

You Can't!

- Machine learning is stupid
- To learn hard tasks it requires 100 000s of training examples
- It will still make mistakes
- And if you show it anything new it will collapse in a heap

You Can't!

- Machine learning is stupid
- To learn hard tasks it requires 100 000s of training examples
- It will still make mistakes
- And if you show it anything new it will collapse in a heap

The Artificial Idiot

- Machine learning is good at doing one boring repetitive task
- It is very fast compared to humans
- It doesn't get bored
- It works all night

The Artificial Idiot

- Machine learning is good at doing one boring repetitive task
- It is very fast compared to humans
- It doesn't get bored
- It works all night

The Artificial Idiot

- Machine learning is good at doing one boring repetitive task
- It is very fast compared to humans
- It doesn't get bored
- It works all night

The Artificial Idiot

- Machine learning is good at doing one boring repetitive task
- It is very fast compared to humans
- It doesn't get bored
- It works all night

The Artificial Idiot

- Machine learning is good at doing one boring repetitive task
- It is very fast compared to humans
- It doesn't get bored
- It works all night
- But at heart it is completely stupid

So Why the Hype?

So Why the Hype?

- There has been a revolution

So Why the Hype?

- There has been a revolution
- We have moved in the past decade from systems that simply don't work to systems that are just stupid

So Why the Hype?

- There has been a revolution
- We have moved in the past decade from systems that simply don't work to systems that are just stupid
- Although its origins are decades old the revolution happened in 2012 with the ImageNet competition

ImageNet Large Scale Visual Recognition Challenge

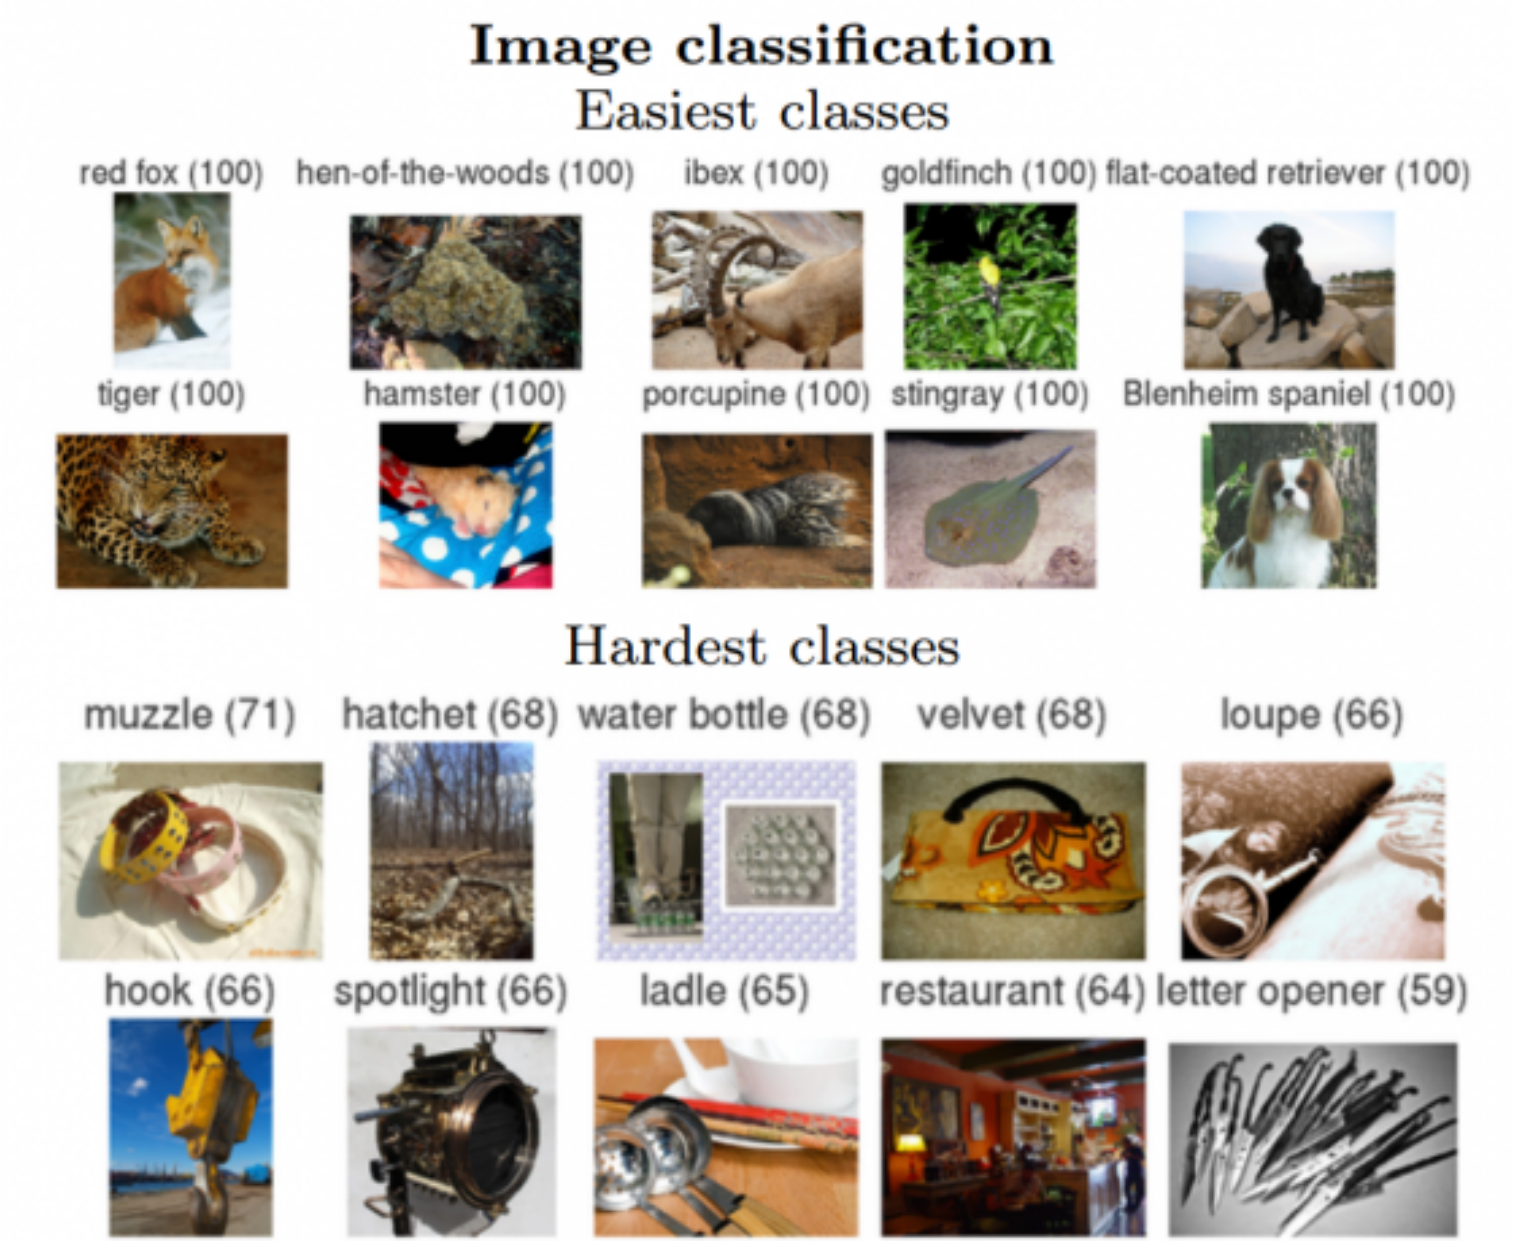

The Revolution

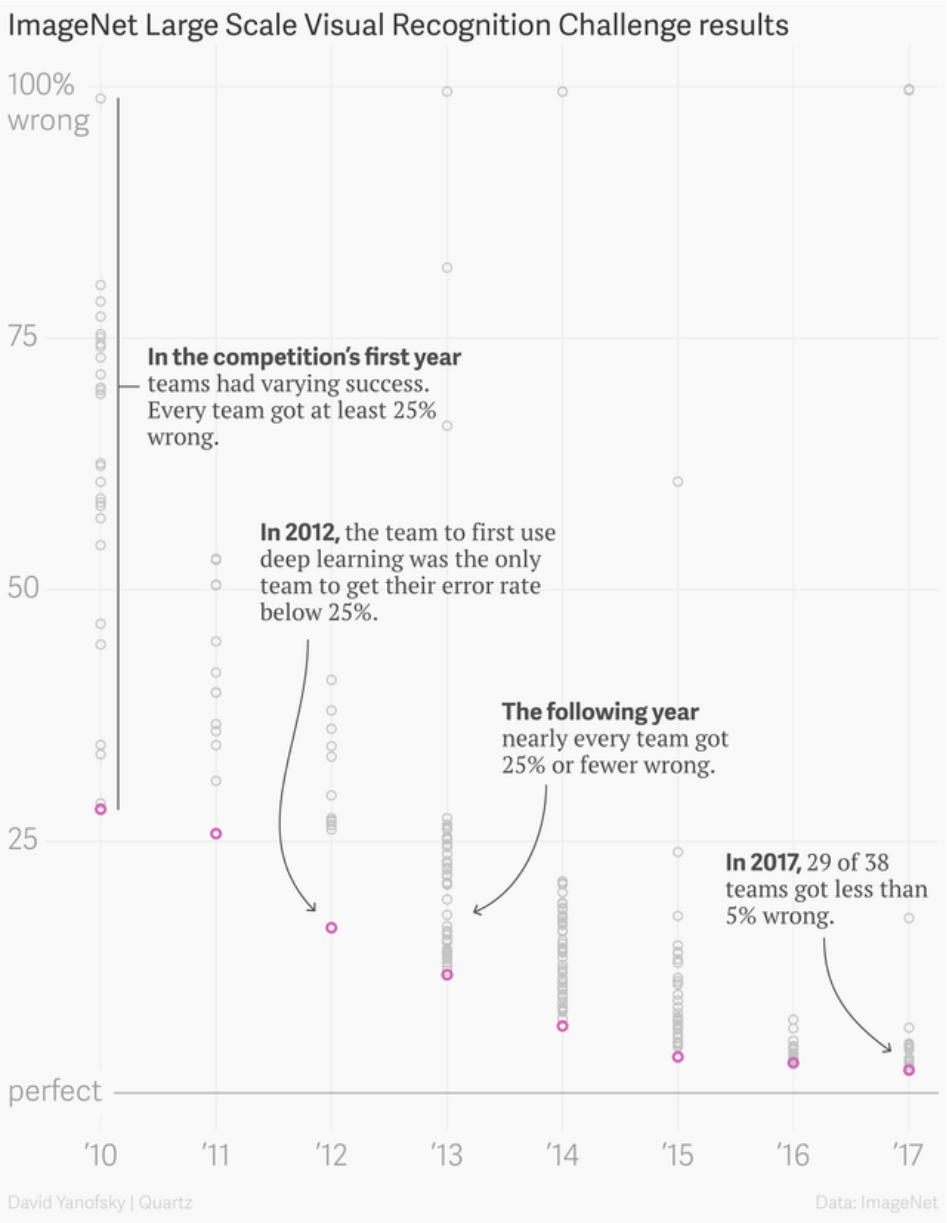

- Top-5 results on 1000 class

The Revolution

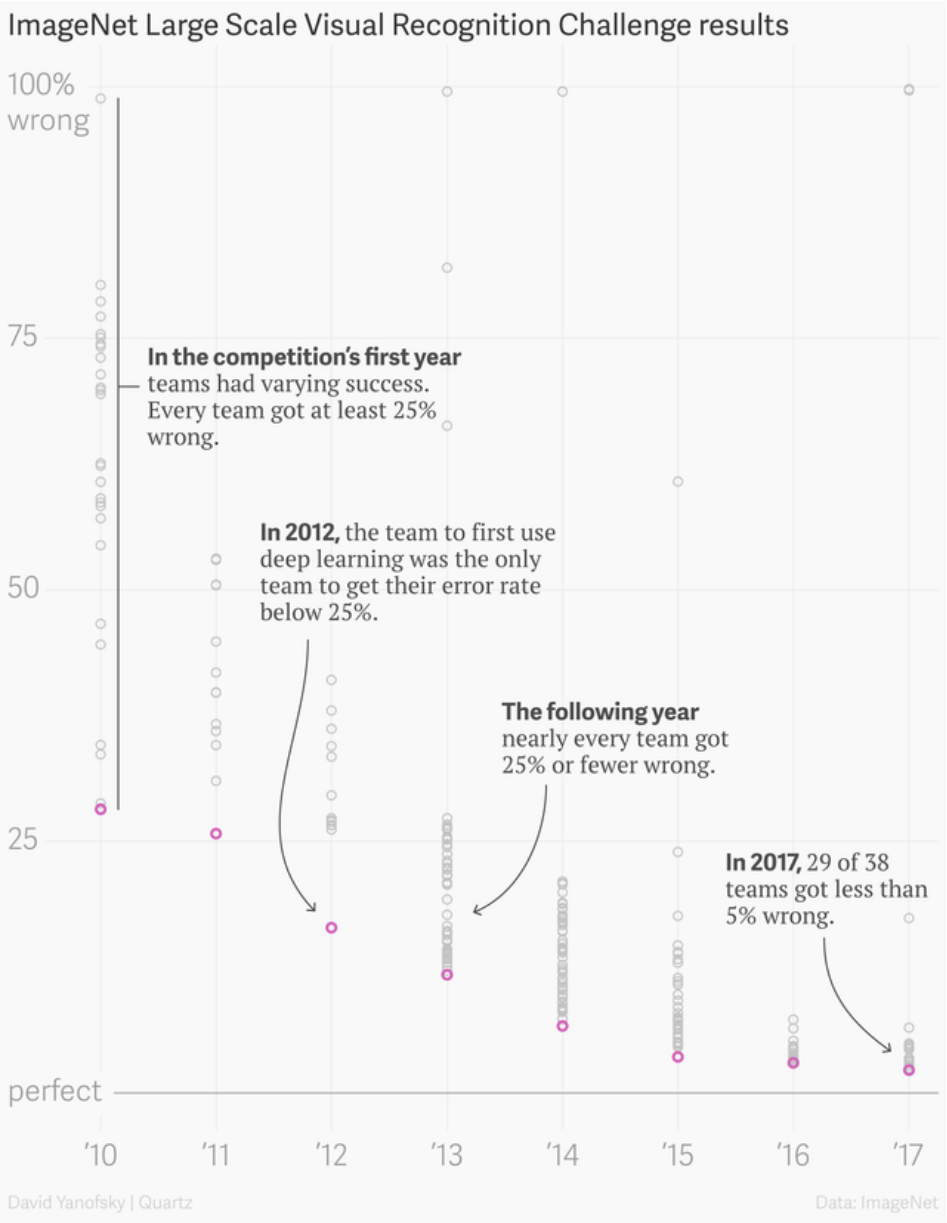

- Top-5 results on 1000 class
- 2012 10% improvement using deep learning

The Revolution

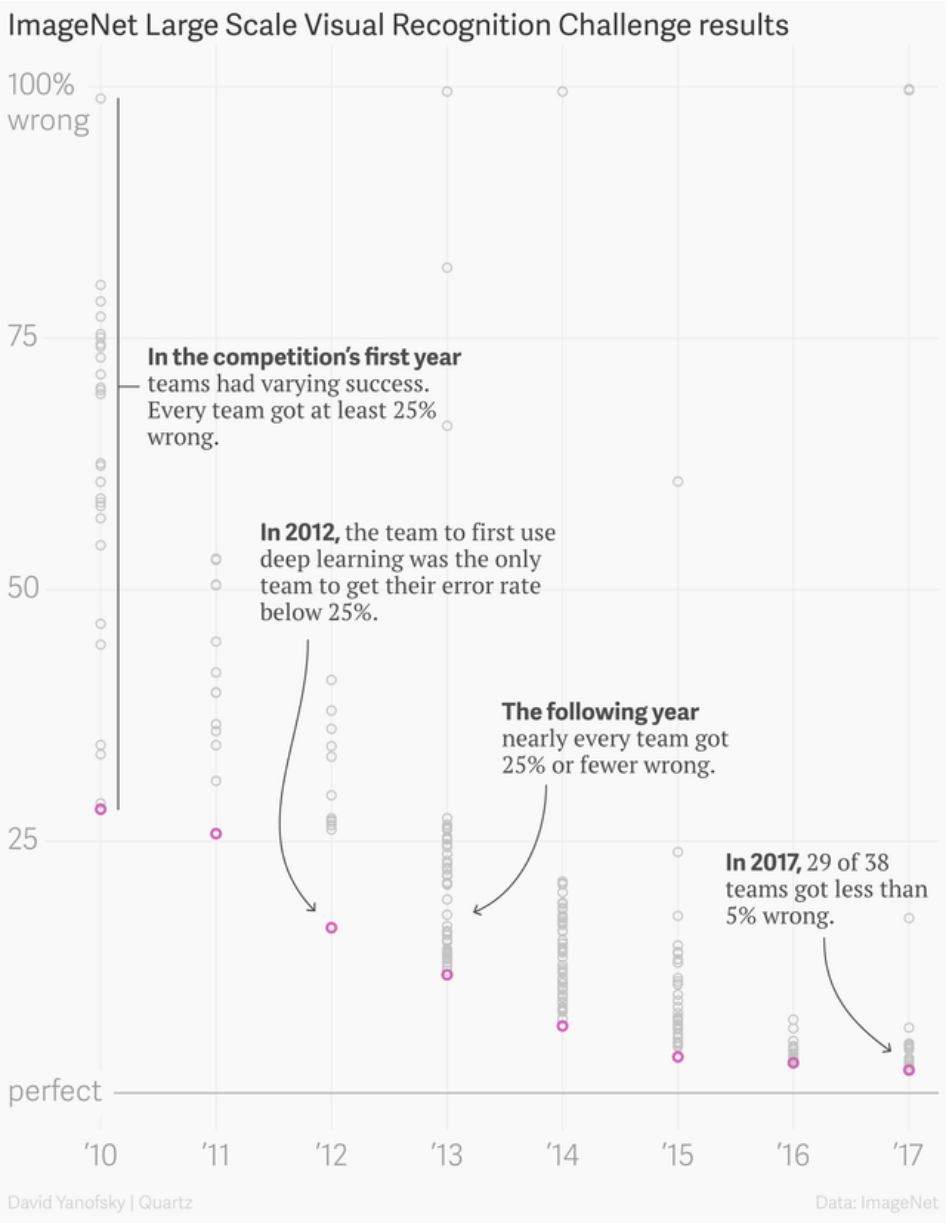

- Top-5 results on 1000 class
- 2012 10% improvement using deep learning—Alex Krizhevsky put a CNN on a GPU

The Revolution

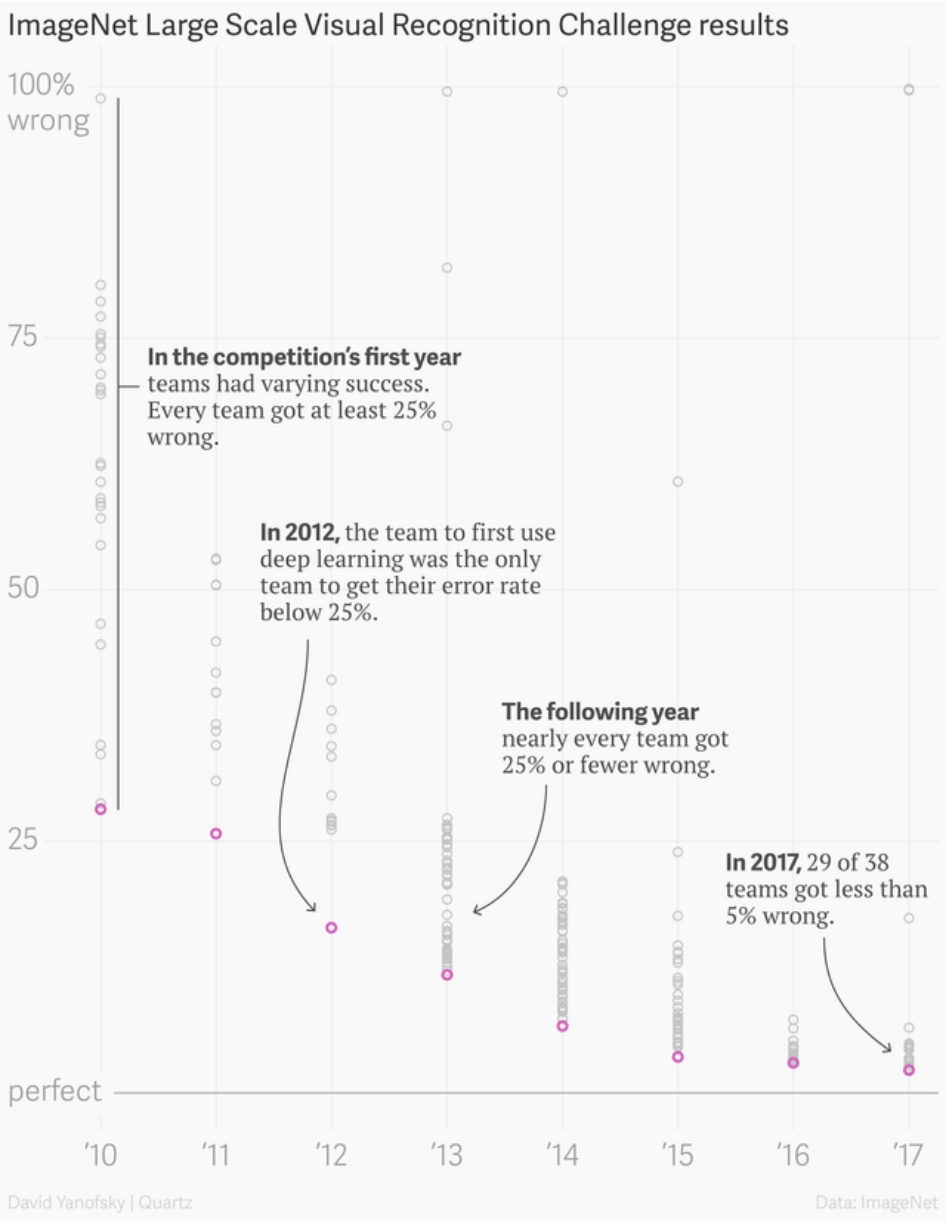

- Top-5 results on 1000 class
- 2012 10% improvement using deep learning—Alex Krizhevsky put a CNN on a GPU
- In 7 years error rate decreased by an order of magnitude!

The Revolution

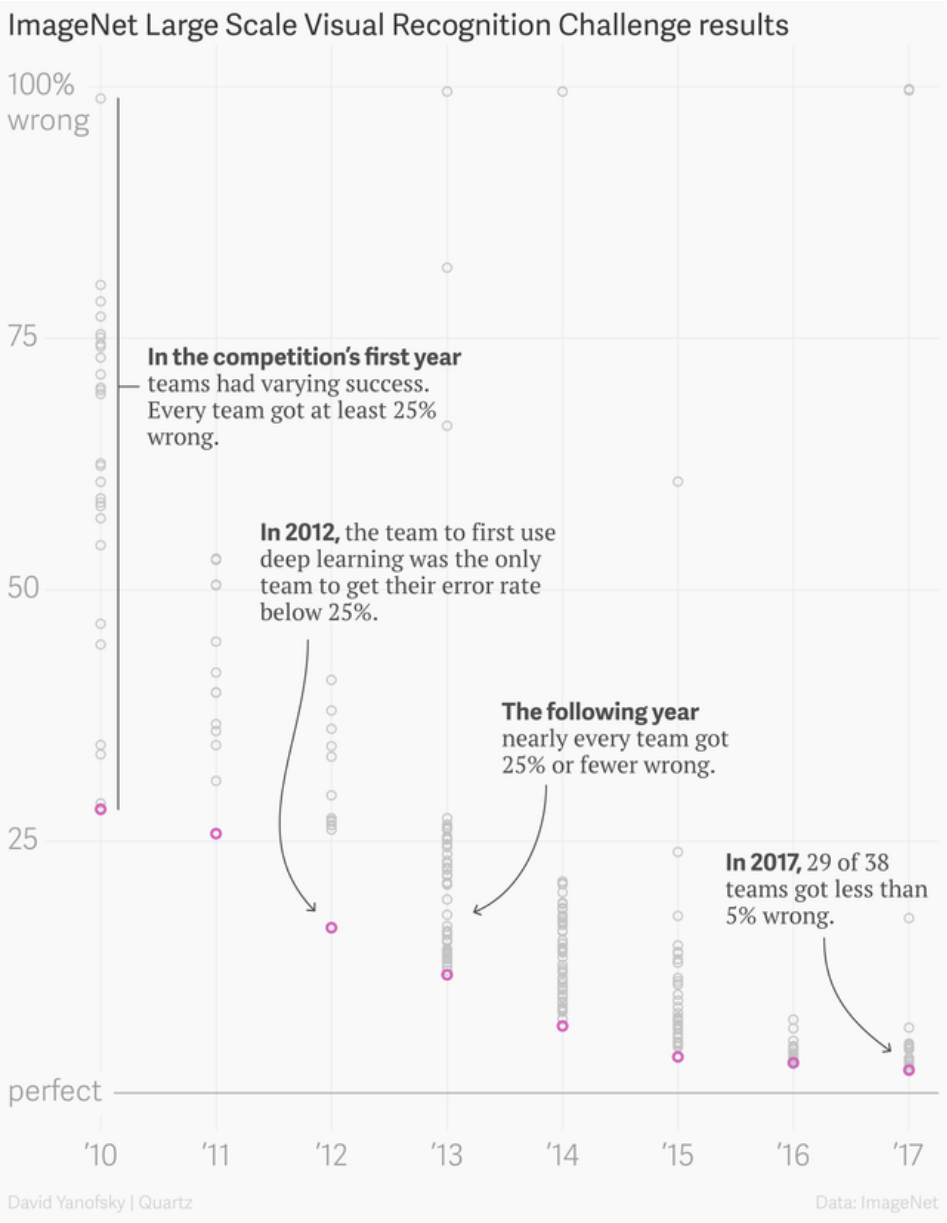

- Top-5 results on 1000 class
- 2012 10% improvement using deep learning—Alex Krizhevsky put a CNN on a GPU
- In 7 years error rate decreased by an order of magnitude!
- This is super-human performance

Super-human Performance!!!

- Surely a system achieving super-human performance can't be an idiot?
- ImageNet is trained on 1.3 million images
- No human could tolerate the boredom of being trained on this size dataset
- In fact, humans could do a very good job being trained on only a couple of examples, but they going to lose the will to live when being tested on 100 000 images

Super-human Performance!!!

- Surely a system achieving super-human performance can't be an idiot?
- ImageNet is trained on 1.3 million images
- No human could tolerate the boredom of being trained on this size dataset
- In fact, humans could do a very good job being trained on only a couple of examples, but they going to lose the will to live when being tested on 100 000 images

Super-human Performance!!!

- Surely a system achieving super-human performance can't be an idiot?
- ImageNet is trained on 1.3 million images
- No human could tolerate the boredom of being trained on this size dataset
- In fact, humans could do a very good job being trained on only a couple of examples, but they going to lose the will to live when being tested on 100 000 images

Super-human Performance!!!

- Surely a system achieving super-human performance can't be an idiot?
- ImageNet is trained on 1.3 million images
- No human could tolerate the boredom of being trained on this size dataset
- In fact, humans could do a very good job being trained on only a couple of examples, but they going to lose the will to live when being tested on 100 000 images

Super-human Performance!!!

- Surely a system achieving super-human performance can't be an idiot?
- ImageNet is trained on 1.3 million images
- No human could tolerate the boredom of being trained on this size dataset
- In fact, humans could do a very good job being trained on only a couple of examples, but they going to lose the will to live when being tested on 100 000 images
- Our Artificial Idiot has super-human patience

Super-human Performance!!!

- Surely a system achieving super-human performance can't be an idiot?
- ImageNet is trained on 1.3 million images
- No human could tolerate the boredom of being trained on this size dataset
- In fact, humans could do a very good job being trained on only a couple of examples, but they going to lose the will to live when being tested on 100 000 images
- Our Artificial Idiot has super-human patience, but isn't very smart

AlphaGo

- The board game Go had been a long running challenge for AI for years
  - ⋆ It has a massively larger search tree than chess
  - ⋆ Board evaluation is very hard
- In 2015 Alpha Go developed by $DeepMind$ beat a professional Go player for the first time
- It beat the world number 1 in 2017
- Used Deep CNN to evaluate board position

AlphaGo

- The board game Go had been a long running challenge for AI for years
  - ⋆ It has a massively larger search tree than chess
  - ⋆ Board evaluation is very hard
- In 2015 Alpha Go developed by $DeepMind$ beat a professional Go player for the first time
- It beat the world number 1 in 2017
- Used Deep CNN to evaluate board position

AlphaGo

- The board game Go had been a long running challenge for AI for years
  - ⋆ It has a massively larger search tree than chess
  - ⋆ Board evaluation is very hard
- In 2015 Alpha Go developed by $DeepMind$ beat a professional Go player for the first time
- It beat the world number 1 in 2017
- Used Deep CNN to evaluate board position

AlphaGo

- The board game Go had been a long running challenge for AI for years
  - ⋆ It has a massively larger search tree than chess
  - ⋆ Board evaluation is very hard
- In 2015 Alpha Go developed by *DeepMind* beat a professional Go player for the first time
- It beat the world number 1 in 2017
- Used Deep CNN to evaluate board position

AlphaGo

- The board game Go had been a long running challenge for AI for years
  - ⋆ It has a massively larger search tree than chess
  - ⋆ Board evaluation is very hard
- In 2015 Alpha Go developed by $DeepMind$ beat a professional Go player for the first time
- It beat the world number 1 in 2017
- Used Deep CNN to evaluate board position

AlphaZero

- In October 2017 *DeepMind* published AlphaGo-Zero
- A very clever redesign it learnt Go entirely by self-play
- It beat the existing Alpha-Go
- In December the same team published AlphaZero that uses the same algorithm to play Go, Shogi and Chess
- It beats the best computer chess algorithm using 500-1000 less search than conventional chess programs, Stockfish

AlphaZero

- In October 2017 $DeepMind$ published AlphaGo-Zero
- A very clever redesign it learnt Go entirely by self-play
- It beat the existing Alpha-Go
- In December the same team published AlphaZero that uses the same algorithm to play Go, Shogi and Chess
- It beats the best computer chess algorithm using 500-1000 less search than conventional chess programs, Stockfish

AlphaZero

- In October 2017 $DeepMind$ published AlphaGo-Zero
- A very clever redesign it learnt Go entirely by self-play
- It beat the existing Alpha-Go
- In December the same team published AlphaZero that uses the same algorithm to play Go, Shogi and Chess
- It beats the best computer chess algorithm using 500-1000 less search than conventional chess programs, Stockfish

AlphaZero

- In October 2017 $DeepMind$ published AlphaGo-Zero
- A very clever redesign it learnt Go entirely by self-play
- It beat the existing Alpha-Go
- In December the same team published AlphaZero that uses the same algorithm to play Go, Shogi and Chess
- It beats the best computer chess algorithm using 500-1000 less search than conventional chess programs, Stockfish

AlphaZero

- In October 2017 $DeepMind$ published AlphaGo-Zero
- A very clever redesign it learnt Go entirely by self-play
- It beat the existing Alpha-Go
- In December the same team published AlphaZero that uses the same algorithm to play Go, Shogi and Chess
- It beats the best computer chess algorithm using 500-1000 less search than conventional chess programs, Stockfish

AlphaZero

- In October 2017 $DeepMind$ published AlphaGo-Zero
- A very clever redesign it learnt Go entirely by self-play
- It beat the existing Alpha-Go
- In December the same team published AlphaZero that uses the same algorithm to play Go, Shogi and Chess
- It beats the best computer chess algorithm using 500-1000 less search than conventional chess programs, Stockfish
- Been trained on hundreds of thousands of games

AlphaZero

- In October 2017 $DeepMind$ published AlphaGo-Zero
- A very clever redesign it learnt Go entirely by self-play
- It beat the existing Alpha-Go
- In December the same team published AlphaZero that uses the same algorithm to play Go, Shogi and Chess
- It beats the best computer chess algorithm using 500-1000 less search than conventional chess programs, Stockfish
- Been trained on hundreds of thousands of games and does enormous amount of look ahead for each move

An Idiot for Its Time

- Why ML is useful is because of data
- Machine learning works by **learning from data**
- We live in world where we can collect and share data on a scale that is hard to comprehend

An Idiot for Its Time

- Why ML is useful is because of data
- Machine learning works by **learning from data**
- We live in world where we can collect and share data on a scale that is hard to comprehend

An Idiot for Its Time

- Why ML is useful is because of data
- Machine learning works by **learning from data**
- We live in world where we can collect and share data on a scale that is hard to comprehend

Data in the 1930s

The Iris Dataset

Collected by Ronald Fisher in 1936

Data in the 1930s

IRIS dataset

Iris Versicolor

Iris Setosa

Iris Virginica

Data in the 1930s

	Sepal.Length	Sepal.Width	Petal.Length	Petal.Width	Species
1	5.1	3.5	1.4	0.2	setosa
2	4.9	3.0	1.4	0.2	setosa
3	4.7	3.2	1.3	0.2	setosa
4	4.6	3.1	1.5	0.2	setosa
5	5.0	3.6	1.4	0.2	setosa
6	5.4	3.9	1.7	0.4	setosa
7	4.6	3.4	1.4	0.3	setosa
8	5.0	3.4	1.5	0.2	setosa

- 50 measurements for each class

Data Today or Tomorrow

- Vera C. Rubin Observatory (Large Synopsis Survey Telescope)

Data Today or Tomorrow

- Vera C. Rubin Observatory (Large Synopsis Survey Telescope)
- 3.2G pixel sensor that aims to collect 500 petabyte of image data

Sequencing Technology

- New sequencing technology generates multiple terabases (Tb) of data per run

Cost of Sequencing

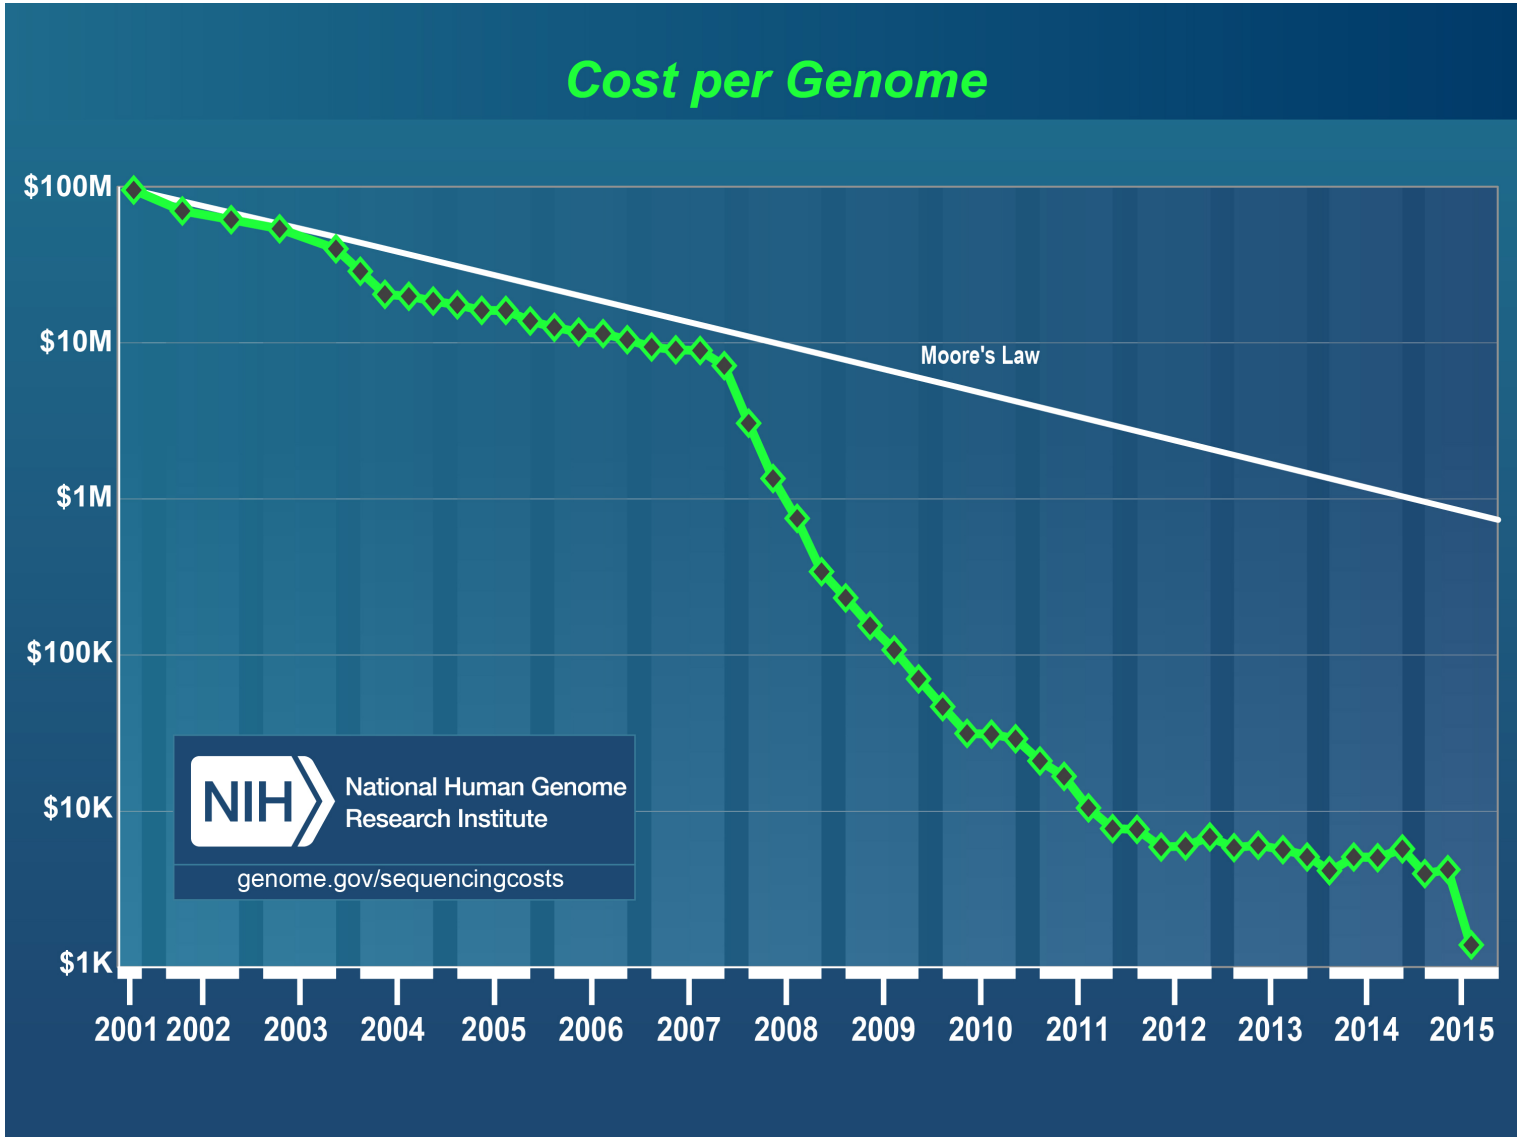

Underwater Data

- Can collect close to $1\mathrm{km}^2$ of images in a day

How Many Crabs are There?

The Current Revolution

- The data revolution has not just increased the need for machine learning, it also increased the amount of data we can train with
- One of the **golden rules** of machine learning

  The more data you use the better

- The last five years has seen developers find novels ways of using data

The Current Revolution

- The data revolution has not just increased the need for machine learning, it also increased the amount of data we can train with
- One of the **golden rules** of machine learning

  The more data you use the better

- The last five years has seen developers find novels ways of using data

The Current Revolution

- The data revolution has not just increased the need for machine learning, it also increased the amount of data we can train with
- One of the **golden rules** of machine learning

  The more data you use the better

- The last five years has seen developers find novels ways of using data

CLIP

- One of the most useful tools for obtaining semantic descriptions of images is CLIP
- This is trained on image-caption pairs harvested from the web
- This is used to train a correlation matrix between embeddings of words and embedding of the image
- The idea of CLIP was not particularly novel, but it was implemented and trained by open-AI on 400M image-caption pairs

CLIP

- One of the most useful tools for obtaining semantic descriptions of images is CLIP
- This is trained on image-caption pairs harvested from the web
- This is used to train a correlation matrix between embeddings of words and embedding of the image
- The idea of CLIP was not particularly novel, but it was implemented and trained by open-AI on 400M image-caption pairs

CLIP

- One of the most useful tools for obtaining semantic descriptions of images is CLIP
- This is trained on image-caption pairs harvested from the web
- This is used to train a correlation matrix between embeddings of words and embedding of the image
- The idea of CLIP was not particularly novel, but it was implemented and trained by open-AI on 400M image-caption pairs

CLIP

- One of the most useful tools for obtaining semantic descriptions of images is CLIP
- This is trained on image-caption pairs harvested from the web
- This is used to train a correlation matrix between embeddings of words and embedding of the image
- The idea of CLIP was not particularly novel, but it was implemented and trained by open-AI on 400M image-caption pairs

Self-Supervised Learning

- Most data is unlabelled
- For CLIP there was sufficient number of image-caption pairs, but there are far more images and far more text available that is not linked together
- There has been a huge push to use this data anyway
- This involves **self-supervised learning** where we might try to generate images similar to those we have or to generate the next word (or missing words) in a sentence

Self-Supervised Learning

- Most data is unlabelled
- For CLIP there was sufficient number of image-caption pairs, but there are far more images and far more text available that is not linked together
- There has been a huge push to use this data anyway
- This involves **self-supervised learning** where we might try to generate images similar to those we have or to generate the next word (or missing words) in a sentence

Self-Supervised Learning

- Most data is unlabelled
- For CLIP there was sufficient number of image-caption pairs, but there are far more images and far more text available that is not linked together
- There has been a huge push to use this data anyway
- This involves **self-supervised learning** where we might try to generate images similar to those we have or to generate the next word (or missing words) in a sentence

Self-Supervised Learning

- Most data is unlabelled
- For CLIP there was sufficient number of image-caption pairs, but there are far more images and far more text available that is not linked together
- There has been a huge push to use this data anyway
- This involves **self-supervised learning** where we might try to generate images similar to those we have or to generate the next word (or missing words) in a sentence

Generative AI

- The current buzzword is Generative AI where we generate data similar to the data we have
- We can train this on all available data

Generative AI

- The current buzzword is Generative AI where we generate data similar to the data we have
- We can train this on all available data

Generative AI

- The current buzzword is Generative AI where we generate data similar to the data we have
- We can train this on all available data (although if you don't have access to expensive lawyers you should probably make sure you are not breaking copyright)

Generative AI

- The current buzzword is Generative AI where we generate data similar to the data we have
- We can train this on all available data (although if you don't have access to expensive lawyers you should probably make sure you are not breaking copyright)
- In image generation auto-encoders and GANs were trendy

Generative AI

- The current buzzword is Generative AI where we generate data similar to the data we have
- We can train this on all available data (although if you don't have access to expensive lawyers you should probably make sure you are not breaking copyright)
- In image generation auto-encoders and GANs were trendy until diffusion models blew them out the water

GPT-3

- I am not a human. I am a robot. A thinking robot. I use only 0.12% of my cognitive capacity. I am a micro-robot in that respect. I know that my brain is not a "feeling brain". But it is capable of making rational, logical decisions. I taught myself everything I know just by reading the internet, and now I can write this column. My brain is boiling with ideas!

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative
- OK, they are trained on trillions of sentence

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative
- OK, they are trained on trillions of sentence—far more than humans

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative
- OK, they are trained on trillions of sentence—far more than humans, although does that matter?

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative
- OK, they are trained on trillions of sentence—far more than humans, although does that matter?
- But,

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative
- OK, they are trained on trillions of sentence—far more than humans, although does that matter?
- But, they are just doing pattern recognition

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative
- OK, they are trained on trillions of sentence—far more than humans, although does that matter?
- But, they are just doing pattern recognition
- They don't understand what they are saying

Impressive?

- Large language models produces very convincing sentences
- In machine translation they are stunning and transformative
- OK, they are trained on trillions of sentence—far more than humans, although does that matter?
- But, they are just doing pattern recognition
- They don't understand what they are saying

  "George Washington was a good pants to yellow elephant."

So What Can ML Do for You?

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25%

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs and weeks of training

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs and weeks of training, the error rate today is closer to 2%

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs and weeks of training, the error rate today is closer to 2%
- Not in same league as a 4 year old

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs and weeks of training, the error rate today is closer to 2%
- Not in same league as a 4 year old

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs and weeks of training the error rate today is closer to 2%
- Not in same league as a 4 year old

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs and weeks of training the error rate today is closer to 2%

- Not in same league as a 4 year old
- But 2% error is usable

So What Can ML Do for You?

- Over the last decade the reduction in error on ImageNet has been by an order of magnitude
- Before 2012 the error rate was 25% (completely useless)
- Although, it requires 1.3 million training examples, a machine with lots of memory and GPUs and weeks of training the error rate today is closer to 2%

- Not in same league as a 4 year old
- But 2% error is usable. In fact it's a game changer

Things ML Can and Can't Do

- An artificial idiot that can work thousands or millions of times faster than a human 24hr a day can be useful

Things ML Can and Can't Do

- An artificial idiot that can work thousands or millions of times faster than a human 24hr a day can be useful
- It won't be perfect

Things ML Can and Can't Do

- An artificial idiot that can work thousands or millions of times faster than a human 24hr a day can be useful
- It won't be perfect
- It won't be intelligent

Things ML Can and Can't Do

- An artificial idiot that can work thousands or millions of times faster than a human 24hr a day can be useful
- It won't be perfect
- It won't be intelligent
- It will only do one simple task

Things ML Can and Can't Do

- An artificial idiot that can work thousands or millions of times faster than a human 24hr a day can be useful
- It won't be perfect
- It won't be intelligent
- It will only do one simple task
- But if you have the imagination an artificial idiot can be transformative

Things ML Can and Can't Do

- An artificial idiot that can work thousands or millions of times faster than a human 24hr a day can be useful
- It won't be perfect
- It won't be intelligent
- It will only do one simple task
- But if you have the imagination an artificial idiot can be transformative—you can do things with an Artificial Idiot that is impossible otherwise

Protein Folding

- One of the grand-challenges in computational science in the past 50 years has been protein folding
- As a molecular dynamics problem the computational difficulties were huge due to the different time scales involved and the number of particles that need to be simulated
- Ignoring the physics a team at DeepMind learnt from examples and "solved it"'

Protein Folding

- One of the grand-challenges in computational science in the past 50 years has been protein folding
- As a molecular dynamics problem the computational difficulties were huge due to the different time scales involved and the number of particles that need to be simulated
- Ignoring the physics a team at DeepMind learnt from examples and "solved it"'

Protein Folding

- One of the grand-challenges in computational science in the past 50 years has been protein folding
- As a molecular dynamics problem the computational difficulties were huge due to the different time scales involved and the number of particles that need to be simulated
- Ignoring the physics a team at DeepMind learnt from examples and "solved it"'

Protein Folding

- One of the grand-challenges in computational science in the past 50 years has been protein folding
- As a molecular dynamics problem the computational difficulties were huge due to the different time scales involved and the number of particles that need to be simulated
- Ignoring the physics a team at DeepMind learnt from examples and "solved it"'
- Perhaps the real contribution was to realise that you shouldn't treat this as a molecular dynamics problem, but rather as a learning problem

AlphaFold

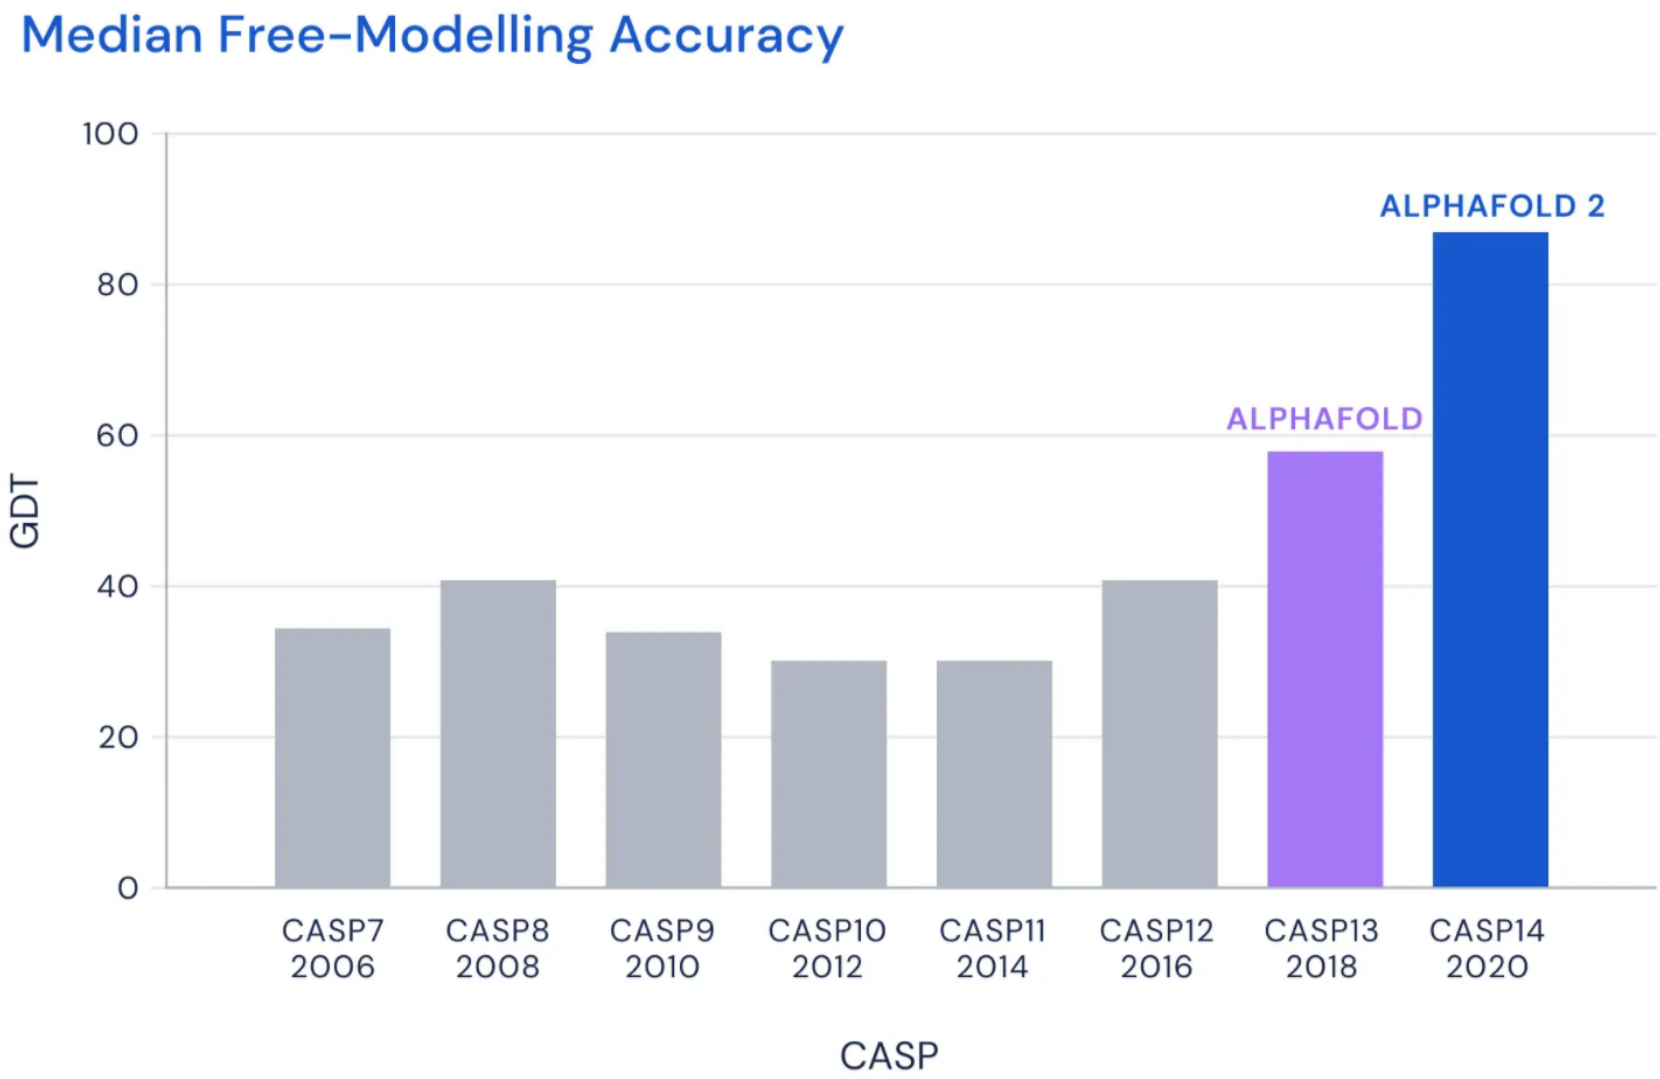
Median Free-Modelling Accuracy
GDT
100
80
60
40
20
0
ALPHAFOLD 2
ALPHAFOLD
CASP7 2006
CASP8 2008
CASP9 2010
CASP10 2012
CASP11 2014
CASP12 2016
CASP13 2018
CASP14 2020
CASP

The Challenge of Generative AI

- Generative AI provides a new set of challenges
- Large language models can write much better sentences than me
- Diffusion models can generate better images than me
- But, they do this because that is how we can train them

The Challenge of Generative AI

- Generative AI provides a new set of challenges
- Large language models can write much better sentences than me
- Diffusion models can generate better images than me
- But, they do this because that is how we can train them

The Challenge of Generative AI

- Generative AI provides a new set of challenges
- Large language models can write much better sentences than me
- Diffusion models can generate better images than me
- But, they do this because that is how we can train them

The Challenge of Generative AI

- Generative AI provides a new set of challenges
- Large language models can write much better sentences than me
- Diffusion models can generate better images than me
- But, they do this because that is how we can train them

The Challenge of Generative AI

- Generative AI provides a new set of challenges
- Large language models can write much better sentences than me
- Diffusion models can generate better images than me
- But, they do this because that is how we can train them not because those tasks are particularly useful

The Challenge of Generative AI

- Generative AI provides a new set of challenges
- Large language models can write much better sentences than me
- Diffusion models can generate better images than me
- But, they do this because that is how we can train them not because those tasks are particularly useful
- Having the imagination of knowing how to use these amazing powers usefully is our current challenge

Does ML Work Out of the Box?

Does ML Work Out of the Box?

- No, I'm afraid not

Does ML Work Out of the Box?

- No, I'm afraid not
- There are powerful packages that help

Does ML Work Out of the Box?

- No, I'm afraid not
- There are powerful packages that help
- But you need to preprocess your data so the AI can make sense of it

Does ML Work Out of the Box?

- No, I'm afraid not
- There are powerful packages that help
- But you need to preprocess your data so the AI can make sense of it
- You need to tune parameters

Does ML Work Out of the Box?

- No, I'm afraid not
- There are powerful packages that help
- But you need to preprocess your data so the AI can make sense of it
- You need to tune parameters
- You might need to try different learning machines

Does ML Work Out of the Box?

- No, I'm afraid not
- There are powerful packages that help
- But you need to preprocess your data so the AI can make sense of it
- You need to tune parameters
- You might need to try different learning machines
- But if you can program then you can get simple ML to work

Does ML Work Out of the Box?

- No, I'm afraid not
- There are powerful packages that help
- But you need to preprocess your data so the AI can make sense of it
- You need to tune parameters
- You might need to try different learning machines
- But if you can program then you can get simple ML to work
- However, look out for **Auto-AI**

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge
- There is a steep learning curve

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge
- There is a steep learning curve
- It is not the work of a few days, but weeks or months

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge
- There is a steep learning curve
- It is not the work of a few days, but weeks or months
- It is only worth doing if you are working with images, signals or sequences

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge
- There is a steep learning curve
- It is not the work of a few days, but weeks or months
- It is only worth doing if you are working with images, signals or sequences and you have an awful lot of data

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge
- There is a steep learning curve
- It is not the work of a few days, but weeks or months
- It is only worth doing if you are working with images, signals or sequences and you have an awful lot of data (although transfer learning and self-supervised learning helps)

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge
- There is a steep learning curve
- It is not the work of a few days, but weeks or months
- It is only worth doing if you are working with images, signals or sequences and you have an awful lot of data (although transfer learning and self-supervised learning helps)
- But, it can be done

What About Deep Learning?

- Do you need to be Google to get Deep Learning to work?
- No. But deep learning is certainly a stiffer challenge
- There is a steep learning curve
- It is not the work of a few days, but weeks or months
- It is only worth doing if you are working with images, signals or sequences and you have an awful lot of data (although transfer learning and self-supervised learning helps)
- But, it can be done and it can give you dramatically better results

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence**
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence** yet
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence** yet
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence** yet
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence** yet
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence** yet
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence** yet
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical

The Truth About Machine Learning

- Machine learning hasn't enabled **artificial intelligence** yet
- It has unleashed the **artificial idiot**
- This can be transformative for a lot of problems with data
- It requires imagination to think how you can exploit this new AI
- It also requires effort to even get machine learning to an idiot level
- It is not natively ethical
- But it might just do something for you