

# Self replicating genies

How to democratise and ensure ethics in AI





All resources:  
[aka.ms/human-ai](https://aka.ms/human-ai)



Let's talk about  
"Artificial Intelligence"



# Artificial Intelligence



- Is **nothing new** – the concepts go back to the 50ies
- Is quite the **hype** and very often **misattributed**
- Is an **umbrella term** for a lot of math and science around repetition, pattern recognition and machine learning
- Got a **huge boost** because of availability of **hardware**
- Became much **more feasible** **because** of the availability of **lots of data**





# Artificial Intelligence



A large collection of books on artificial intelligence and robotics, including titles like 'The Inevitable', 'Deep Learning Revolution', 'Robot Sex', and 'The Most Human Human'. The books are arranged in a dense, overlapping manner, with some standing upright and others lying flat. The covers feature various designs, including text, illustrations of robots, and abstract patterns. Some books have red circular price tags with '149,-' written on them. The background is dark and out of focus, emphasizing the books in the foreground.

@codepo8



## Reminders of "genie in the bottle"



- Fulfills our wishes seemingly with invisible magic
- Useful, and feels too good to be true
- Once released, may have a dark, sinister edge to it
- Hard to put back into the bottle.



Let's start with  
some predictions.



- AI is the number one growth **market** in IT – the others are cloud and security
- Machine Learning is already **replacing thousands of jobs** – boring, terrible jobs humans should not do
- This is also happening in IT – **we are not invincible** because we know how to exit Vim



Let's start with  
some predictions.



- There is no stopping this – it is just too convenient
- The amount of data we create (actively or by triggering sensors) demands machines to whittle it down for us to make it consumable by humans
- If we as developers and decision makers in IT don't take ownership and lead with good, ethical examples, we'll throw away decades of work democratising computing





The machines  
are watching...



# Social Credit System






Big brother is  
redundant...



- Everything we do online is **monitored** and **recorded**
- We often don't realise that our **data** is how we **pay** for "free" services
- We're happy to use systems that **record** all the time **in exchange** for **convenience**
- Often people don't realise just how **dangerous** this can be in the **wrong** hands.



Everything counts  
in large amounts



- We create a **massive** amount of information – actively and without our knowledge.
- It is tough to make that amount of information **consumable** again.
- That's why we have **computers**
- With cloud computing, on demand processing and advances in hardware **we're faster than ever.**





Leaving invisible marks...



- By using **other people's machines and infrastructure**, we leave **traces**
- This allows companies to **recognise** us, and accumulates a **usage history**
- This leads to **better results**, but can leak data
- We should have more **transparency** about what digital legacy we left behind.




Are machines  
friend or foe?



# Artificial Intelligence Myths




- AI can't replace a thinking, creative human
- AI can not **magically fill gaps with perfect information** – it can only compare and assume
- AI doesn't learn in a creative fashion. It makes **no assumptions**
- AI has **no morals and ethics**, but – used wrongly – it can **amplify our biases**



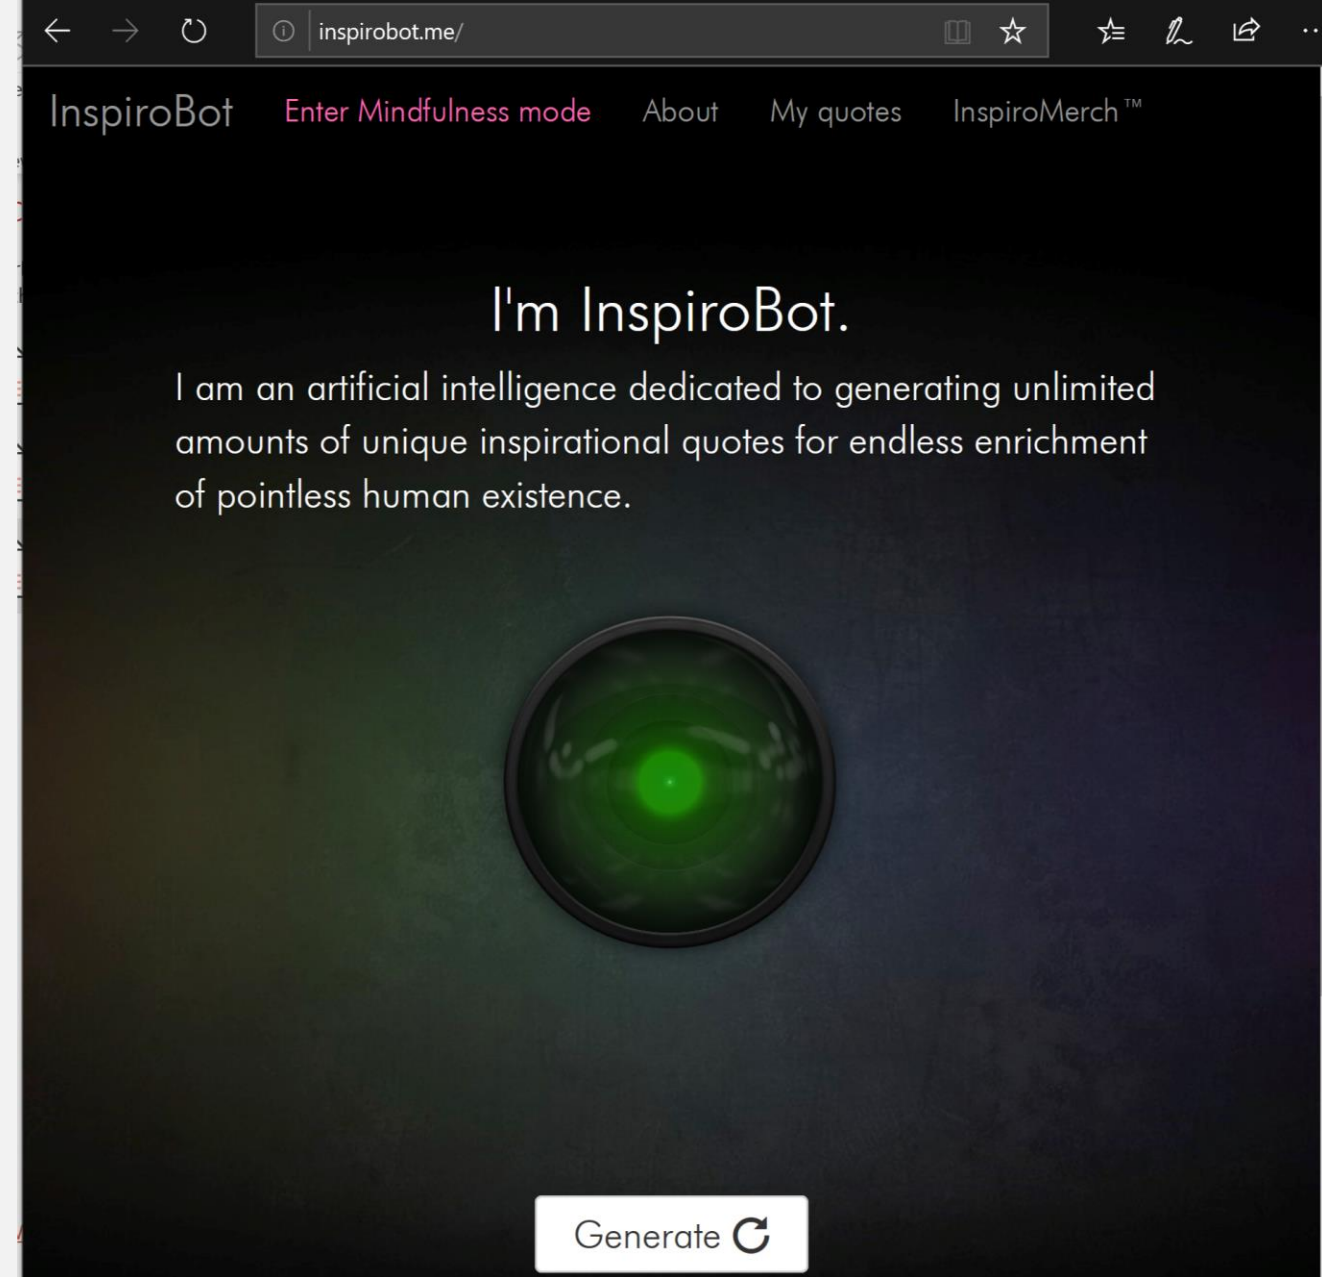
Machines can be  
great tools or  
weapons...


---

- Machine Learning is all about **returning assumptions**
- We don't get any definitive **truth** from **algorithms**, we get answers to our questions
- AI can answer questions, but it is up to you to **ask good questions** – generic questions yield assumed results.



Unguided or  
supervised AI...







It can be  
demanding








It can mix up  
needs...






It can be overly  
excited...








It can be a good  
warning...

---

Remember that  
you are pissed.






It can be painfully  
humbling...

---




Basing your  
everyday on science  
creates loneliness.





Prophetic, even?






Passive aggressive  
towards humans...





It can be  
adoringly cute...






Whilst being  
actually kick-ass







Machines can be  
great tools or  
weapons...

---

- Untrained and limited data leads to terrible and biased AI results
- It is very easy to get either **wrong deductions** or **false positives**
- AI is as **intelligent** and **good** as the **people** who **apply** it



Machine learning  
helps us in a few  
ways...



- Recommendation
- Prediction
- Classification
- Clustering
- Generation





# Recommendation



Machines ploughing through lots of data for you.

- “I feel lucky” moments
- Slack finding people in your organization
- Intelligent inboxes
- Automated photo optimization
- Automated tagging and alternative text: “Image may contain”




# Prediction



You're doing this – you probably want this as the next thing

- Text autocompletion
- Task offerings
- Image tooling – adding photos to a collage
- Creating albums
- Offering similar music and videos
- Offering products that match




# Classification



Sort things by what humans told you what they are and scale it up

- Google surveys offering the right form elements for a question
- Detecting faces and asking for more information
- Finding anomalies in health scans and doing the same for all the ones in the system




# Clustering



Find own patterns and collate them

- Photo tagging and ordering
- Document analysis
- Comment filtering and triaging
- Video optimisation dependent on content.




# Generation



Allow the machine to create things

- Art style matching
- Generated articles from fact collection
- Synthesised music
- Filling content with tagged information (grass, houses, brick, etc...)
- React to human input

# We need to find our place on the scale






# About face...




## Face detection

Detect one or more human faces in an image and get back face rectangles for where in the image the faces are, along with face attributes which contain machine learning-based predictions of facial features. The face attribute features available are: Age, Emotion, Gender, Pose, Smile, and Facial Hair along with 27 landmarks for each face in the image.



```
Detection result:
JSON:
[
  {
    "faceId": "0e96b668-a0d0-46ec-a5c9-ad26b16a1ca9",
    "faceRectangle": {
      "top": 166,
      "left": 128,
      "width": 218,
      "height": 218
    },
    "faceAttributes": {
      "hair": {
        "bald": 0.0,
        "invisible": false,
        "hairColor": [
          {
            "color": "blond",
            "confidence": 1.0
          },
          {
            "color": "other",
            "confidence": 0.0
          }
        ]
      }
    }
  }
]
```





# About face...



- Face rectangle / Landmarks
- Pose (pitch/roll/yaw)
- Smile
- Gender/Age
- Type of glasses
- Makeup (lips/eye)
- Emotion (anger, contempt, disgust, fear, happiness, neutral, sadness, surprise)
- Occlusion (forehead/eye/mouth)
- Facial hair (moustache/beard/sideburns)
- Attributes: Hair (invisible, bald, colour)



# Is this you? Are those also you?

## Face verification

Check the likelihood that two faces belong to the same person. The API will return a confidence score about how likely it is that the two faces belong to one person.




Image URL

Submit

 Browse




Image URL


Submit

 Browse



Verification result: The two faces belong to the same person. **Confidence is 0.7349.**

# Is this your driver?



# Taking it too far?

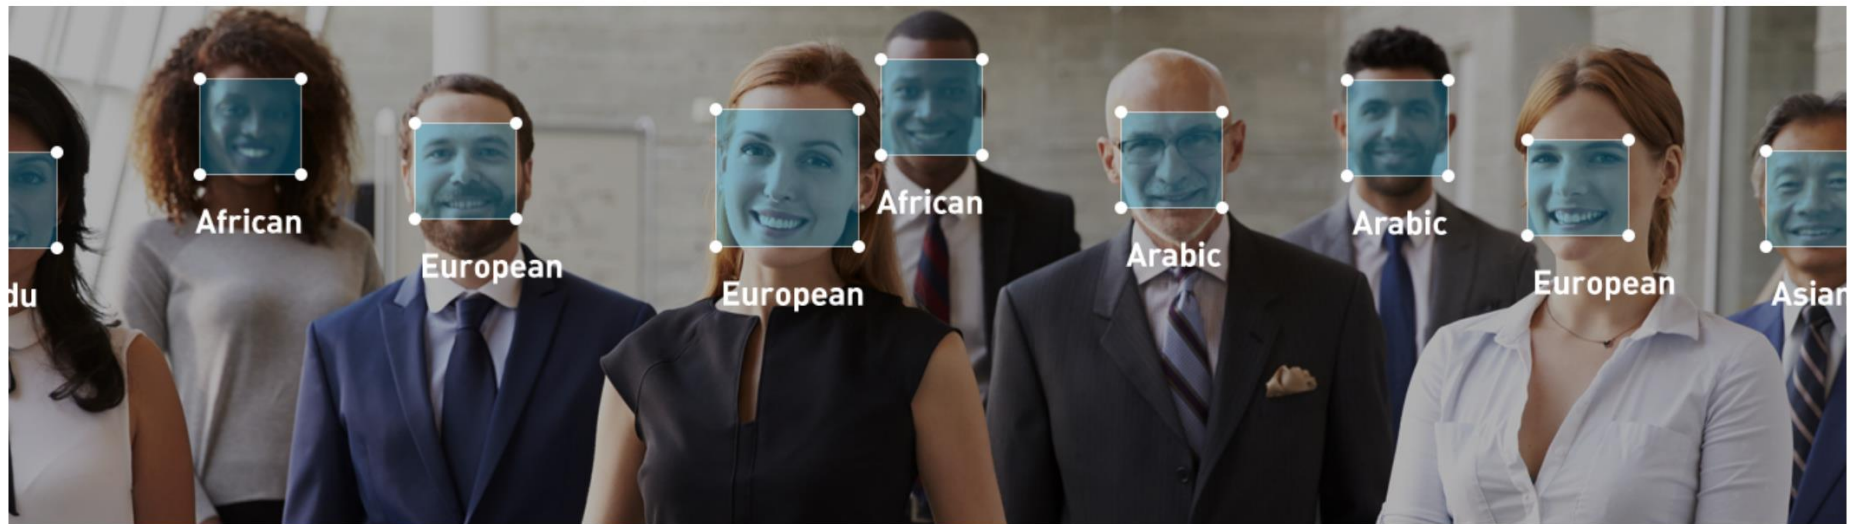
<https://ntechlab.com/>



## Coming soon

PATH TRACKING | ETHNICITY RECOGNITION

Recognizes a person's ethnicity






Detecting even  
more...



Chinese authorities have begun deploying "gait recognition" AI software in Beijing and Shanghai that identifies people via their body shapes and how they walk (Dake Kang/Associated Press)



**Chinese 'gait recognition' tech IDs people by how they walk**

BEIJING (AP) — Chinese authorities have begun deploying a new surveillance tool: "gait recognition" software that uses people's body shapes and how they walk to

[apnews.com](https://apnews.com)





Those trustworthy  
avatars...






Those trustworthy  
avatars...




INSTRUCTION: press +/- to adjust feature, toggle feature name to lock the feature



random face

Male	Age	Skin_Tone
- +	- +	- +
Bangs	Hairline	Bald
- +	- +	- +
Big_Nose	Pointy_Nose	Makeup
- +	- +	- +
Smiling	Mouth_Open	Wavy_Hair
- +	- +	- +
Beard	Goatee	Sideburns
- +	- +	- +
Blond_Hair	Black_Hair	Gray_Hair
- +	- +	- +
Eyeglasses	Earrings	Necktie
- +	- +	- +



# Automated face mapping...



Social Mapper has a variety of uses in the security industry, for example the automated gathering of large amounts of social media profiles for use on targeted phishing campaigns. Facial recognition aids this process by removing false positives in the search results, so that reviewing this data is quicker for a human operator.


Social Mapper supports the following social media platforms:

- LinkedIn
- Facebook
- Twitter
- GooglePlus
- Instagram
- VKontakte
- Weibo
- Douban

Social Mapper takes a variety of input types such as:

- An organisations name, searching via LinkedIn
- A folder full of named images
- A CSV file with names and url's to images online





Once you are  
known...




- Create fake social media profiles to 'friend' the targets and send them links or malware. Recent statistics show social media users are more than twice as likely to click on links and open documents compared to those delivered via email.
- Trick users into disclosing their emails and phone numbers with vouchers and offers to make the pivot into phishing, vishing or smishing.
- Create custom phishing campaigns for each social media site, knowing that the target has an account. Make these more realistic by including their profile picture in the email. Capture the passwords for password reuse.
- View target photos looking for employee access card badges and familiarise yourself with building interiors.



www.ericvander.com

# AI for humans

# I want people to appreciate AI, without giving up their data unwillingly...





# The best way to do this, is to stop selling it as magic, but as a tool...

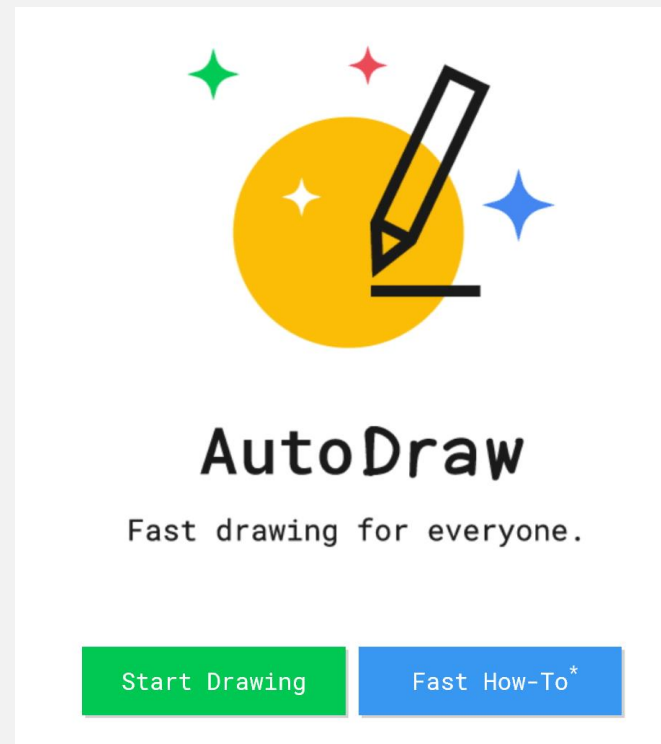


# We need data, so let's make it joyful for humans to give us some



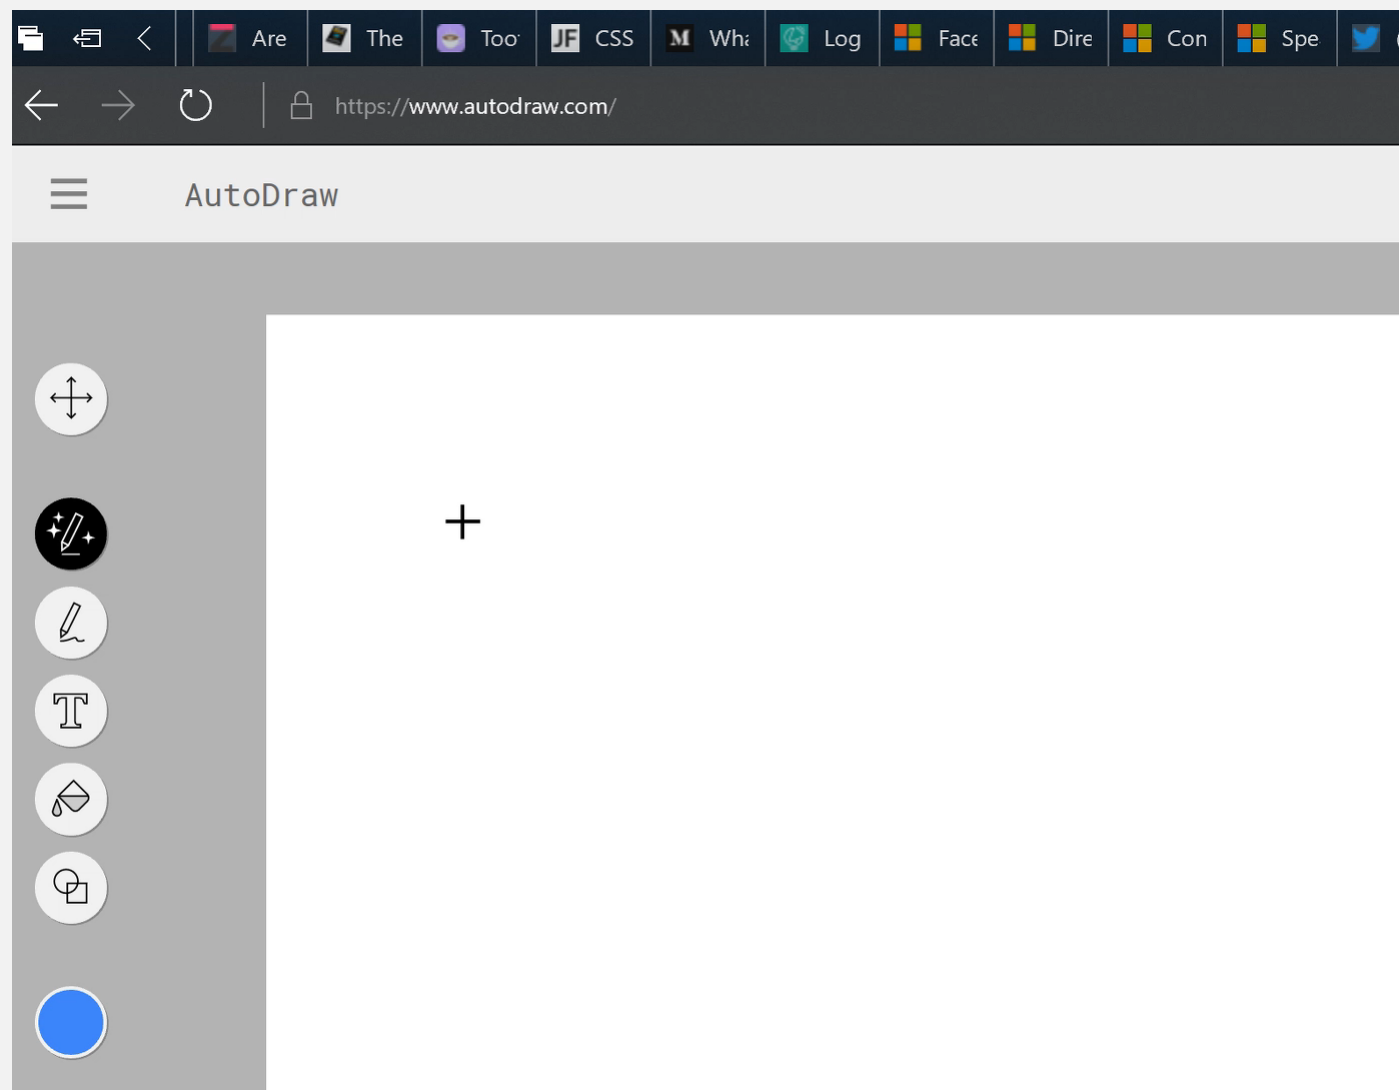


Humans  
and  
Bots/Computers







Humans  
and  
Bots/Computers








# Humans and Bots/Computers



Can a neural network learn to recognize doodling?

Help teach it by adding your drawings to the [world's largest doodling data set](#), shared publicly to help with machine learning research.

Let's Draw!



Humans  
and  
Bots/Computers



Security Check  
(Required)

☐ I'm not a robot

reCAPTCHA  
Privacy - Terms

Select all images with rivers.

Grid of 9 images for selection:

- Image 1: Mountain landscape with a river.
- Image 2: Mountain landscape with a river.
- Image 3: Island in the ocean.
- Image 4: Mountain landscape with a river.
- Image 5: Mountain landscape with a river.
- Image 6: River flowing through a forest.
- Image 7: River flowing through a forest.
- Image 8: River flowing through a forest.
- Image 9: River flowing through a forest.

Verify

Report a problem




“Learning” from  
lots of images





# Humans and Bots/Computers






NEWS CENTER

ACCELERATED COMPUTINGARTIFICIAL INTELLIGENCEAUTONOMOUS VEHICLESDESIGN & VISUALIZATIONGAME

Comments 13.3K



## New AI Imaging Technique Reconstructs Photos with Realistic Results

April 22, 2018


Researchers from NVIDIA, led by Guilin Liu, [introduced](#) a state-of-the-art [deep learning](#) method that can edit images or reconstruct a corrupted image, one that has holes or is missing pixels.

The method can also be used to edit images by removing content and filling in the resulting holes.

The method, which performs a process called “image inpainting”, could be implemented in photo editing software to remove unwanted content, while filling it with a realistic computer-generated alternative.


“Our model can robustly handle holes of any shape, size location, or distance from the image borders. Previous deep learning approaches have focused on rectangular regions located around the center of the image, and often rely on expensive post-processing,” the NVIDIA researchers stated in their [research paper](#). “Further, our model gracefully handles holes of increasing size.”

To prepare to train their neural network, the team first generated 55,116 masks of random streaks and holes of arbitrary shapes and sizes for training. They also generated nearly 25,000 for testing. These were further categorized into six categories based on sizes relative to the input image, in order to improve reconstruction accuracy.




An example of the masks generated for training.

aka.ms/nvidia-fix-image




# Humans and Bots/Computers

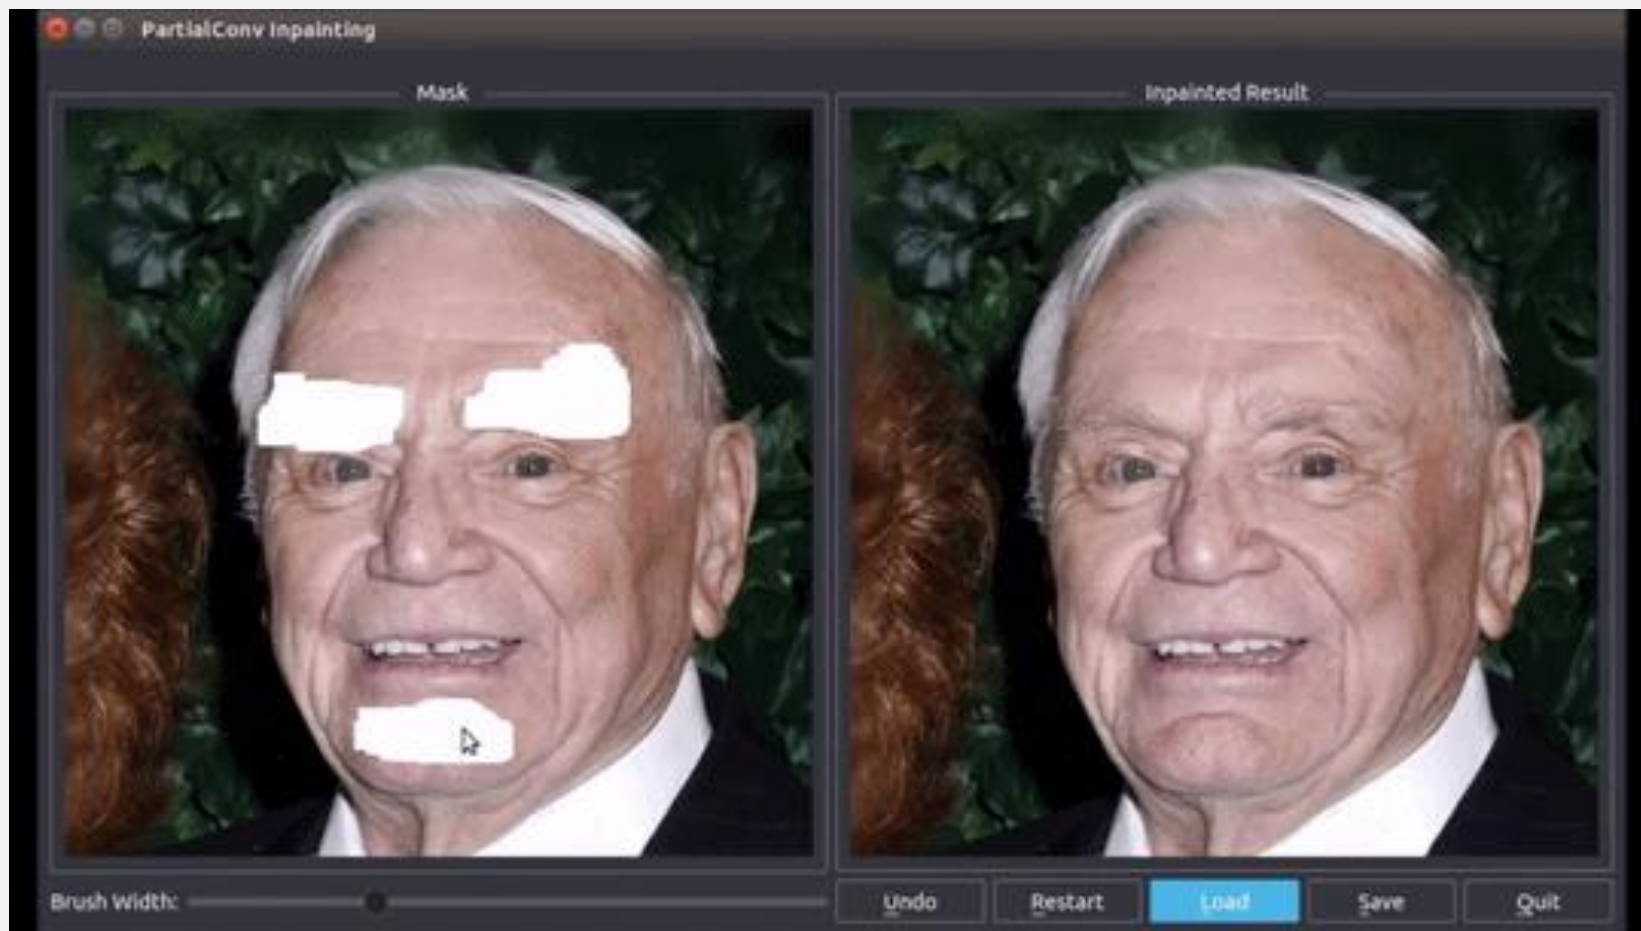


[aka.ms/nvidia-fix-image](https://aka.ms/nvidia-fix-image)






Humans  
and  
Bots/Computers



[aka.ms/nvidia-fix-image](https://aka.ms/nvidia-fix-image)



Humans  
and  
Bots/Computers







Vision and image  
analysis...



instagram: @larryandanke





# Vision and image analysis...



Alt Text: A dog sitting on a sidewalk



Vision and image  
analysis...



**dallas**  
@mixedhunty


Follow



mom said you had to let me use the xbox



10:03 PM - 1 Apr 2018



# Vision and image analysis...



**dallas**  
@mixedhunty

Follow

mom said you had to let me use the xbox



**Burke Holland** ✓ @burkeholland · Apr 4

Replying to @mixedhunty @codepo8

#vision\_api



1




**#vision\_api** @vision\_api · Apr 4

Confidence: 86.74 %

Vision API: Ed Sheeran standing in a room

Full API Result: [vision-api.azurewebsites.net/api/tweetmedia...](https://vision-api.azurewebsites.net/api/tweetmedia...)

# #vision\_api



# Intelligent, responsive systems



- AI services offer us lots of **data** to **compare** our users' input with
- Thus our users **don't need to speak computer** but be **human** instead
- We can **prevent** them from making mistakes
- We can help getting around physical barriers




# Real life examples? This stage 16:00







# Our toolkit for more human interfaces



Natural  
language  
processing



Computer  
Vision



Sentiment  
analysis



Speech  
conversion  
and analysis



Moderation



# Demos? This stage, 14:00



## AI in the Modern Workplace

Kirk Koenigsbauer  
Corporate Vice President, Microsoft 365



With great power  
comes great  
responsibility...



## Our responsibilities..

---

- AI can be an **amazing help** for humans
- It does need **transparency** – if you use people as data sources, they need to know what and where it goes
- When people get information **filtered by an algorithm**, it should be an **opt-in**
- People need to have a **chance to dispute** when an algorithm tagged or disallowed them access.



# Want to go deep?

---

- The Math behind ML
- The ethics of AI
- Working with Data using Python
- Machine Learning Models
- Deep Learning Models
- Reinforcement Learning Models
- Microsoft Professional Program Certificate in Artificial Intelligence

[aka.ms/learn-ai](https://aka.ms/learn-ai)

10 courses, (8-16 hours each), 10 skills



# Want to go deep?



Demystifying Artificial Intelligence: Understanding Machine Learning | ORIGINAL

Christian Hellmann, Senior Developer at Microsoft [+ Follow](#)

12 Videos (58m) [View My Notes](#)


▶ 1. Introduction	1:44
🔒 2. What is Machine Learning	5:25
🔒 3. How We Teach Machines	5:48
🔒 4. Machine Learning to Help Humans	5:28
🔒 5. Tools for Machine Learning	3:44
🔒 6. Visual Uses	7:54
🔒 7. Speaking Human	6:07
🔒 8. Audio & Video	6:32

97 students are watching this class

[About](#) [Community 3](#) [Class Project](#) [All Projects](#) [Save](#) [Add to Calendar](#) [f](#) [t](#) [e](#) [Report class](#)

## skl.sh/christian

Free with trial sign-up




# Who controls our data? Who benefits?

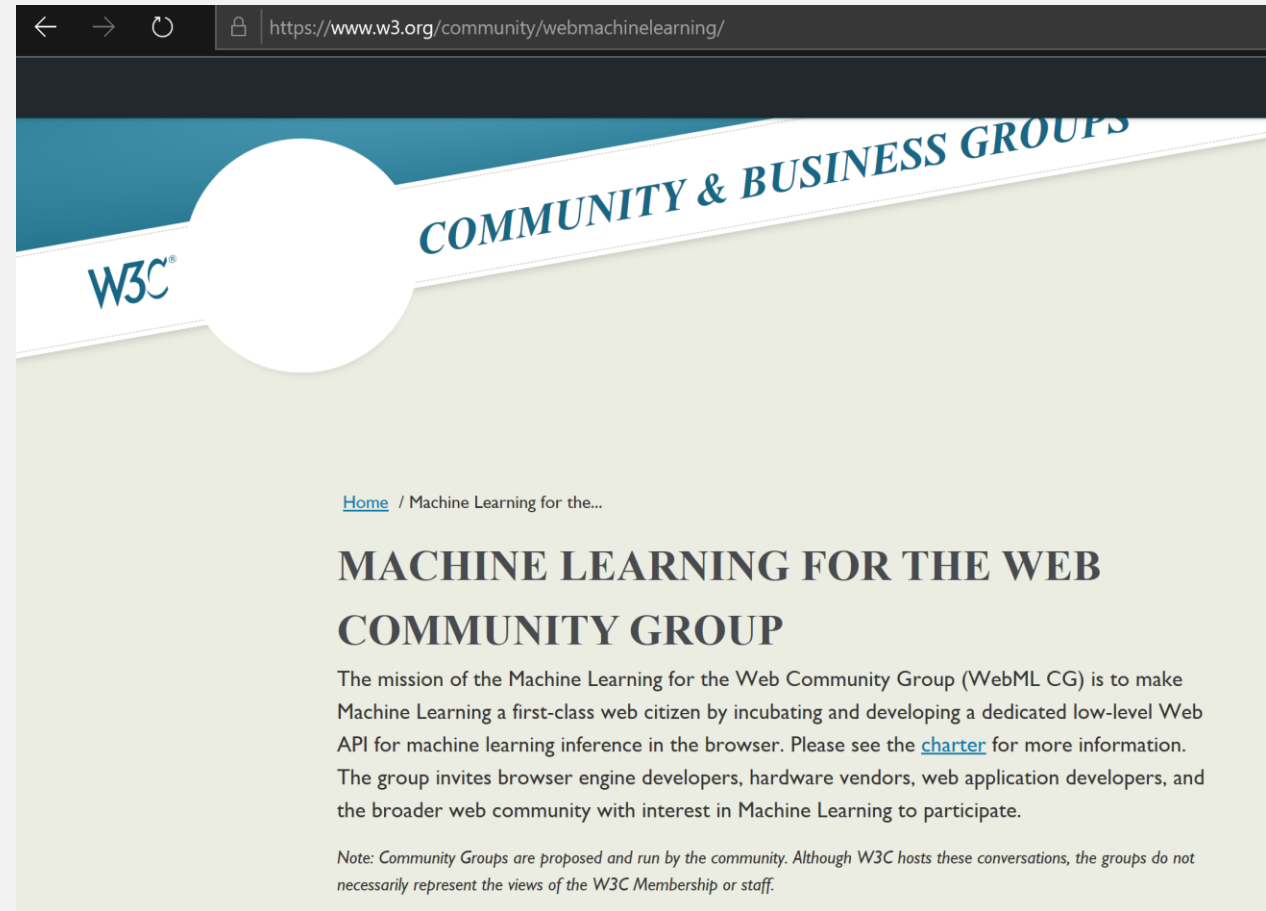


- With all this we need to make clear **who has your data** and **where** it goes.
- Wouldn't it be great if we could do more **on our devices**?
- Much lower latency, better security, increased privacy
- Right now, this is only possible in **native environments**
- I want to change that – a **W3C proposal** to bring **accelerated Machine Learning to the web** in JavaScript






# Who controls our data? Who benefits?




w3.org/community/webmachinelearning



Don't forget to have fun!






Suz Hinton


---





# Categorising images by gesture







Find your moves



MOVE MIRROR



Make a GIF




Stay silly...



**Cassie Evans**

<https://codepen.io/cassie-codes/pen/jKaVqo/>





# Help the human




**Charlie Gerard**   
@devdevcharlie

Following



I used Tensorflow.js to build a quick prototype of head-controlled keyboard to allow people to communicate with simple head movements. Might not be much but it's amazing that you can do stuff like that in JS 🥰 [charliegerard.github.io/teachable-keyb...](https://charliegerard.github.io/teachable-keyboard/) [#tensorflowjs](#) [#javascript](#) [#ML](#) [#accessibility](#)



# TEACHABLE KEYBOARD

Write using motion control

Start Training



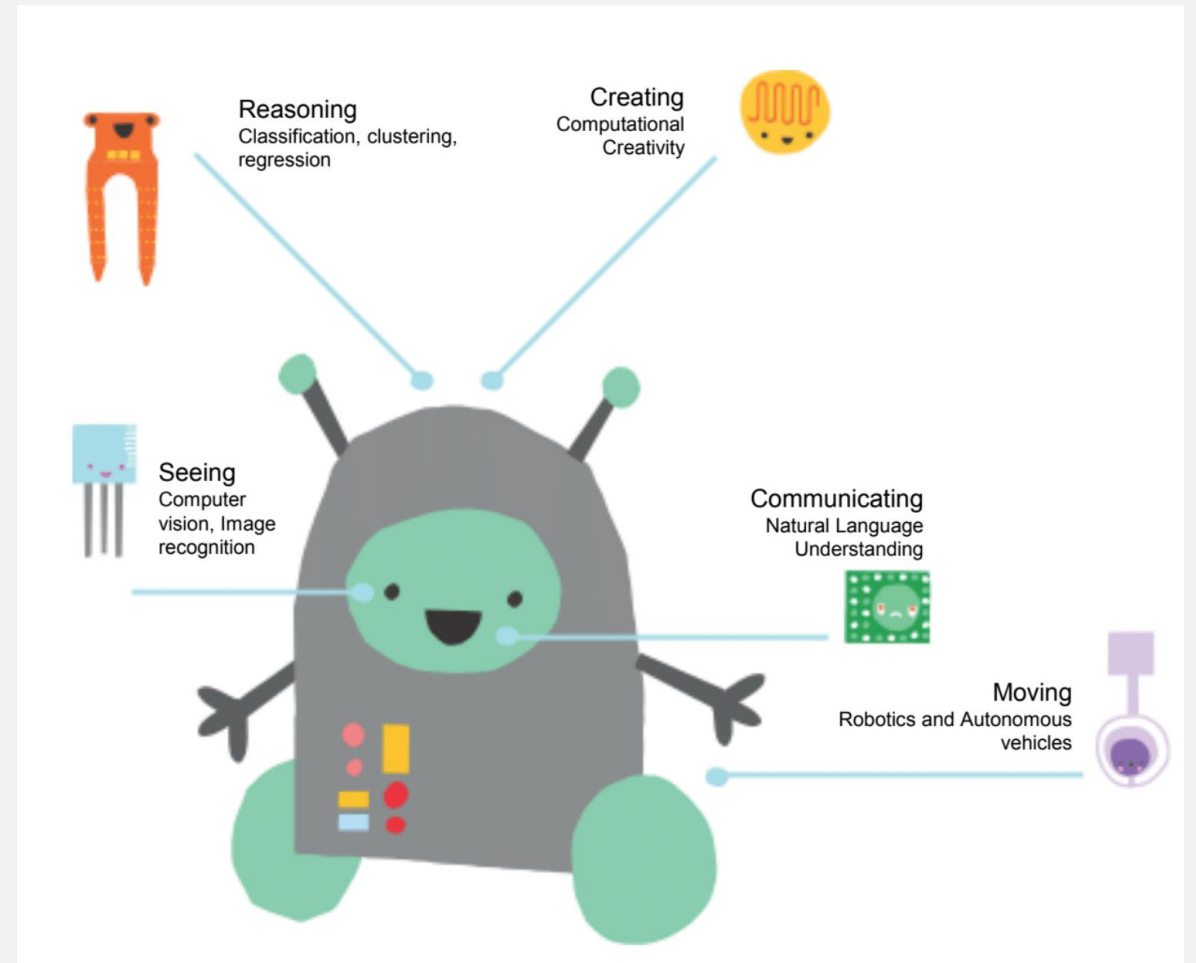


Collaborate and  
share...





# Preparing the next generation



# Thanks!

Chris Heilmann

Christianheilmann.com

Developer-evangelism.com

@codepo8

---



<http://inspirobot.me/>