

Can Essays Written by AI Systems Be the Next Great Cheat Code?

Ped Day Session

14 October 2020

Dawson's AI-Themed Community of Practice

cbouffard@dawsoncollege.qc.ca

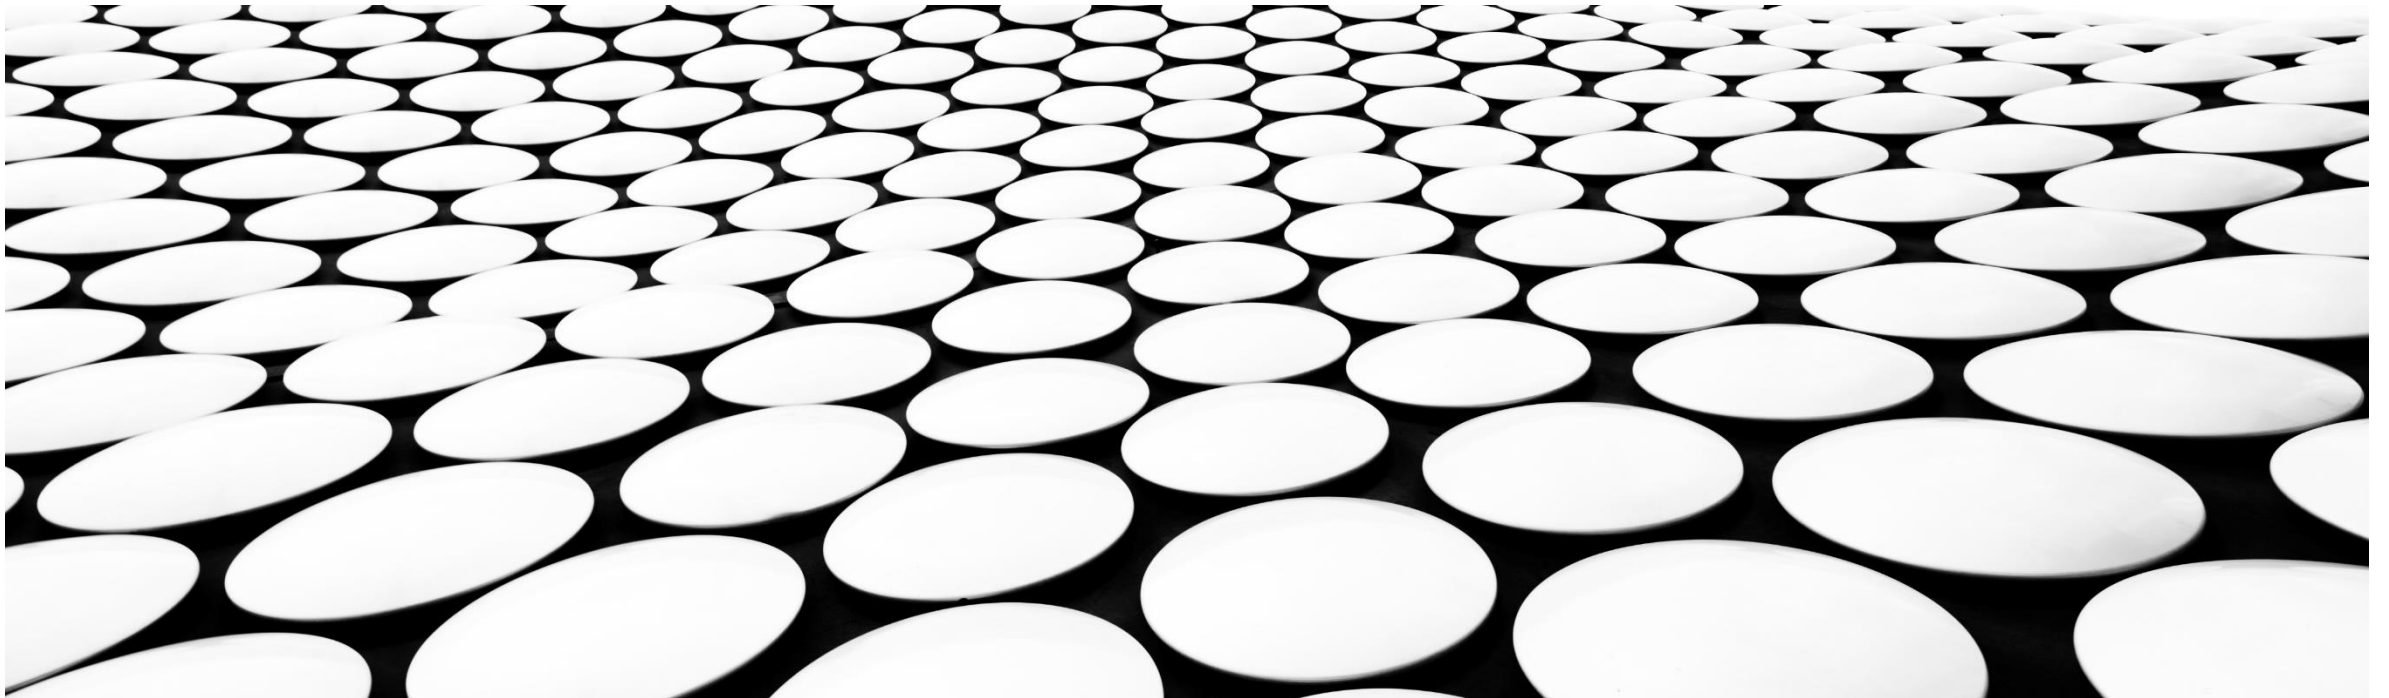


Agenda

1. What text-generating AI systems do and how they have gotten better over time (Daniel Pomerantz)
2. GPT-2 and GPT-3: their origins and capabilities (Carl Saucier-Bouffard)
3. iGPT: the same core technology applied to pixels (Charles Le Guen)
4. The death of the author 3.0: authorship, machines + originality (Cheryl Simon)
5. Academic integrity in the "Desert of the Real" (Greg Polakoff)
6. Q & A

AUTOMATIC TEXT GENERATION SYSTEMS: A BRIEF HISTORY

DANIEL POMERANTZ





TEXT GENERATION : WHAT IS IT FOR

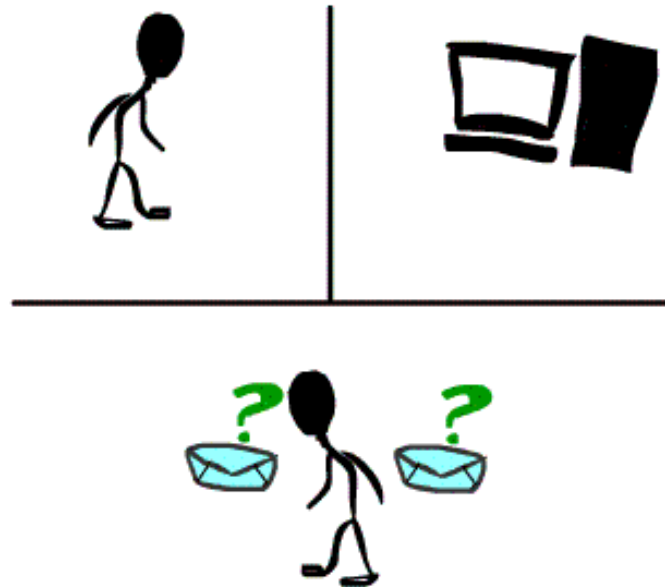
- Text generation applications are within the realm of *natural language generation*
- Highly related to natural language processing
- The idea is to generate text in as *human like* a way as possible

HISTORICALLY

- Georg Philipp Harsdorffer (1607 - 1658)
- German poet



TURING TEST (A.K.A. THE IMITATION GAME) - 1950



<https://commons.wikimedia.org/wiki/File:Turing-test.gif>

CHAT BOT : ELIZA

- Created 1964-1966 at MIT
- Followed a "script"
- Meant to provide a parody of "the responses of a non-directional psychotherapist in an initial psychiatric interview"
- Users attributed human-like feelings to the computer program and some believed it was a human

Welcome to

```
EEEEEE LL      IIII ZZZZZZZZ AAAAA
EE      LL      II      ZZ  AA  AA
EEEEEE LL      II      ZZZ  AAAAAA
EE      LL      II      ZZ  AA  AA
EEEEEE LLLLLL IIII ZZZZZZZZ AA  AA
```

Eliza is a mock Rogerian psychotherapist.

The original program was described by Joseph Weizenbaum in 1966.
This implementation by Norbert Landsteiner 2005.

ELIZA: Is something troubling you ?

YOU: Men are all alike.

ELIZA: What is the connection, do you suppose ?

YOU: They're always bugging us about something or other.

ELIZA: Can you think of a specific example ?

YOU: Well, my boyfriend made me come here.

ELIZA: Is it important to you that your boyfriend made you come here?

YOU: He says I'm depressed much of the time.

ELIZA: I am sorry to hear that you are depressed.

YOU: It's true. I am unhappy.

ELIZA: Can you explain what made you unhappy ?

YOU: █

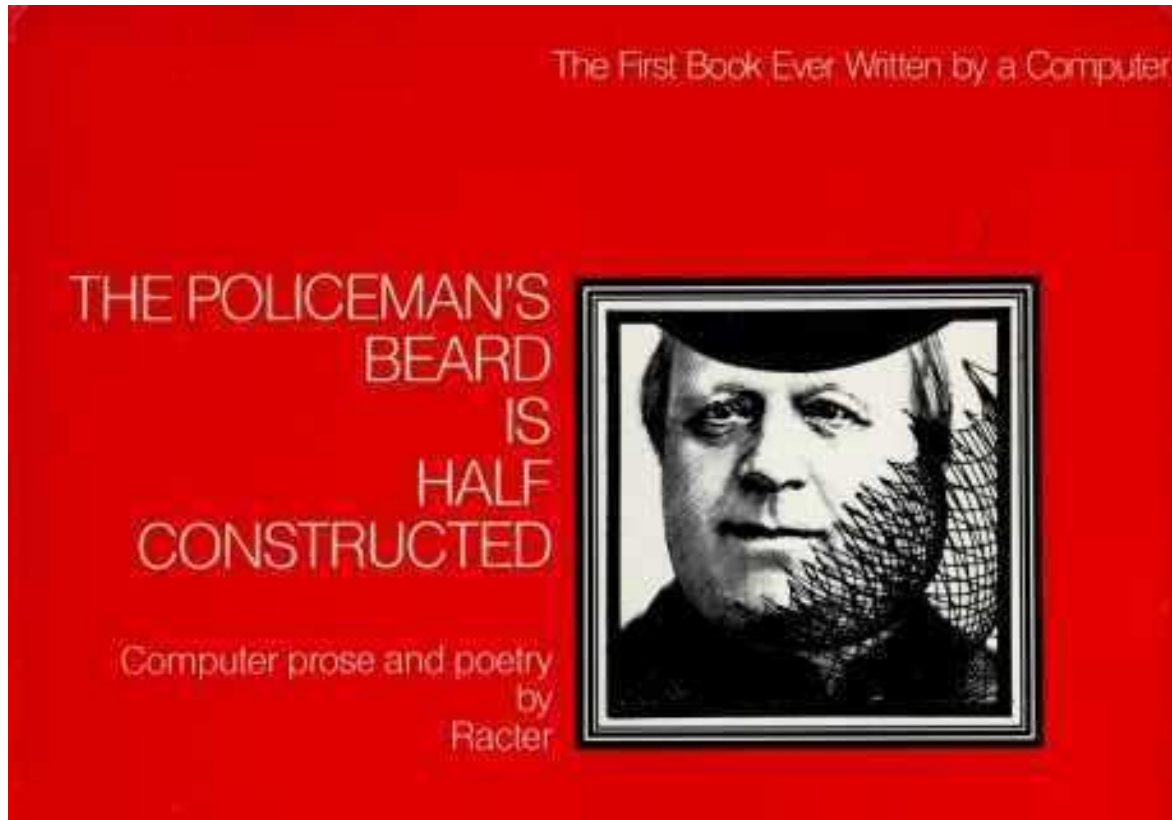
CHAT PATIENT : PARRY

- Created in 1972 (Stanford)
- Described as "ELIZA with attitude"
- Attempted to model the behavior of a paranoid schizophrenic
- 33 psychiatrists were shown transcripts of conversations, half of which were from PARRY, half from real patients
- 48% correct identification

WHEN PARRY MET ~~SALLY~~ ELIZA

- ELIZA: Good Evening. Tell me your problems. Please terminate input with a period or a question mark.
- PARRY: People get on my nerves sometimes.
- ELIZA: I am not sure I understand you fully.
- PARRY: You should pay more attention.
- ELIZA: Suppose you should pay more attention.
- PARRY: You're entitled to your own opinion.
- ELIZA: What makes you think I am entitled to my own opinion?
- PARRY: It bothers me just to be around people in general.
- ELIZA: Please go on.
- PARRY: I went to the track at Bay Meadows a while back.
- ELIZA: What does that suggest to you?
- PARRY: I went to the races.

THE POLICEMAN'S BEARD IS HALF CONSTRUCTED (1984)



A book written by an artificial intelligence

"More than iron, more than lead, more than gold I
need electricity.
I need it more than I need lamb or pork or lettuce
or cucumber.
I need it for my dreams."



2005: THE POLLEN FORECAST FOR SCOTLAND

- Similar to a template (i.e. madlibs)



TURING TEST : WHICH IS THE AI?

Grass pollen levels for Friday have increased from the moderate to high levels of yesterday with values of around 6 to 7 across most parts of the country. However, in Northern areas, pollen levels will be moderate with values of 4.

Pollen counts are expected to remain high at level 6 over most of Scotland, and even level 7 in the south east. The only relief is in the Northern Isles and far northeast of mainland Scotland with medium levels of pollen count.

TURING TEST : WHICH IS THE AI?

Grass pollen levels for Friday have increased from the moderate to high levels of yesterday with values of around 6 to 7 across most parts of the country. However, in Northern areas, pollen levels will be moderate with values of 4.

Pollen counts are expected to remain high at level 6 over most of Scotland, and even level 7 in the south east. The only relief is in the Northern Isles and far northeast of mainland Scotland with medium levels of pollen count.



TODAY: NEWS ARTICLES

- Sports recaps
- Weather
- Financial reports
- Real estate analysis

STAT SHEET / AUTOMATED INSIGHTS / STATS

- Started in 2007 providing statistics for sports
- No human journalist, blogger or writer
- Associated Press used it to automated NCAA Division 1 men's basketball previews for entire 2018 season
 - AP claims it allowed journalists to focus on more in depth articles by freeing up 20% of journalist's time spent normally writing recaps
- Yahoo Sports does something similar
 - ESports opens up an amusing possibility

GPT-2

Language Models are Unsupervised Multitask Learners

Alec Radford^{*1} Jeffrey Wu^{*1} Rewon Child¹ David Luan¹ Dario Amodei^{**1} Ilya Sutskever^{**1}

Abstract

Natural language processing tasks, such as question answering, machine translation, reading comprehension, and summarization, are typically approached with supervised learning on task-specific datasets. We demonstrate that language models begin to learn these tasks without any explicit supervision when trained on a new dataset of millions of webpages called WebText. When conditioned on a document plus questions, the answers generated by the language model reach 55 F1 on the CoQA dataset - matching or exceeding the performance of 3 out of 4 baseline systems without using the 127,000+ training examples. The capacity of the language model is essential to the success of zero-shot task transfer and increasing it improves performance in a log-linear fashion across tasks. Our largest model, GPT-2, is a 1.5B parameter Transformer that achieves state of the art results on 7 out of 8 tested language modeling datasets in a zero-shot setting but still underfits WebText. Samples from the model reflect these improvements and contain coherent paragraphs of text. These findings suggest a promising path towards building language processing systems which learn to perform tasks from their naturally occurring demonstrations.

competent generalists. We would like to move towards more general systems which can perform many tasks – eventually without the need to manually create and label a training dataset for each one.

The dominant approach to creating ML systems is to collect a dataset of training examples demonstrating correct behavior for a desired task, train a system to imitate these behaviors, and then test its performance on independent and identically distributed (IID) held-out examples. This has served well to make progress on narrow experts. But the often erratic behavior of captioning models (Lake et al., 2017), reading comprehension systems (Jia & Liang, 2017), and image classifiers (Alcorn et al., 2018) on the diversity and variety of possible inputs highlights some of the shortcomings of this approach.

Our suspicion is that the prevalence of single task training on single domain datasets is a major contributor to the lack of generalization observed in current systems. Progress towards robust systems with current architectures is likely to require training and measuring performance on a wide range of domains and tasks. Recently, several benchmarks have been proposed such as GLUE (Wang et al., 2018) and decaNLP (McCann et al., 2018) to begin studying this.

Multitask learning (Caruana, 1997) is a promising framework for improving general performance. However, multitask training in NLP is still nascent. Recent work reports modest performance improvements (Nezari et al.,

What paragraph was written by GPT-2?



1) “The 2016 British referendum shows that the majority of the UK citizens were not satisfied with being members of the European Union. Some politicians think that, that was because of the UK having some imperialistic views regarding Europe and they do not want to be a regular member but be a leader state in the region. Everyone agrees that Brexit is one of the most important and even vital events in the history of the UK. With this event, however it somehow loses connection with other countries in Europe.”

2) “Britain has been haggling over how to withdraw from the European Union, the process known as Brexit, since Britons voted to do so in a referendum in 2016. The issue of the long-term trade relationship has been the one that has lingered the longest, long after Scots and English voted on separate questions, but the negotiations have brought out fears in both the short and the long term, some of which, like the citizens' rights argument, have made it hard to get people to talk about the costs or benefits of leaving. In the last few weeks, though, it has become less clear just how damaging it would be to the rest of the union. EU trade negotiators have argued they cannot get a deal without the UK paying more.”

3) “Since Britons voted to do so in a referendum in 2016, the badly divided government has been in crisis, unable to agree on an approach to what may well be the country’s biggest peacetime decision in decades. The deadline to depart the bloc, recently extended once again by the European Union, is now January 31. The battle has cost one prime minister, Theresa May, her job. She announced in late May that she would resign after failing to come up with a plan that satisfied her party, her coalition partners and officials in Brussels, the seat of the European Union.”

GPT-3

arXiv.org > cs > arXiv:2005.14165

Search... All fields  Search
[Help](#) | [Advanced Search](#)

Computer Science > Computation and Language

[Submitted on 28 May 2020 (v1), last revised 22 Jul 2020 (this version, v4)]

Language Models are Few-Shot Learners

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, Dario Amodei

Recent work has demonstrated substantial gains on many NLP tasks and benchmarks by pre-training on a large corpus of text followed by fine-tuning on a specific task. While typically task-agnostic in architecture, this method still requires task-specific fine-tuning datasets of thousands or tens of thousands of examples. By contrast, humans can generally perform a new language task from only a few examples or from simple instructions – something which current NLP systems still largely struggle to do. Here we show that scaling up language models greatly improves task-agnostic, few-shot performance, sometimes even reaching competitiveness with prior state-of-the-art fine-tuning approaches. Specifically, we train GPT-3, an autoregressive language model with 175 billion parameters, 10x more than any previous non-sparse language model, and test its performance in the few-shot setting. For all tasks, GPT-3 is applied without any gradient updates or fine-tuning, with tasks and few-shot demonstrations specified purely via text interaction with the model. GPT-3 achieves strong performance on many NLP datasets, including translation, question-answering, and cloze tasks, as well as several tasks that require on-the-fly reasoning or domain adaptation, such as unscrambling words, using a novel word in a sentence, or performing 3-digit arithmetic. At the same time, we also identify some datasets where GPT-3's few-shot learning still struggles, as well as some datasets where GPT-3 faces methodological issues related to training on large web corpora. Finally, we find that GPT-3 can generate samples of news articles which human evaluators have difficulty distinguishing from articles written by humans. We discuss broader societal impacts of this finding and of GPT-3 in general.

Comments: 40+32 pages

Subjects: **Computation and Language (cs.CL)**

Cite as: [arXiv:2005.14165](#) [cs.CL]

(or [arXiv:2005.14165v4](#) [cs.CL] for this version)

Submission history

From: Tom B Brown [[view email](#)]

[v1] Thu, 28 May 2020 17:29:03 UTC (6,995 KB)

[v2] Mon, 1 Jun 2020 17:08:53 UTC (6,997 KB)

[v3] Fri, 5 Jun 2020 02:52:35 UTC (6,998 KB)

[v4] Wed, 22 Jul 2020 19:47:17 UTC (6,998 KB)

Download:

- [PDF](#)
- [Other formats](#)
(license)

Current browse context:

cs.CL

[< prev](#) | [next >](#)

[new](#) | [recent](#) | [2005](#)

Change to browse by:

cs

References & Citations

- [NASA ADS](#)
- [Google Scholar](#)
- [Semantic Scholar](#)

[24 blog links](#) ([what is this?](#))

[Export Bibtex Citation](#)

Bookmark



Figure 2.2: Total compute used during training. Based on the analysis in Scaling Laws For Neural Language Models [20] we train much larger models on many fewer tokens than is typical. As a consequence, although GPT-3 is over 10x larger than RoBERTa-Large (355M params), both models took roughly 50 petaflop/s-days of compute for pre-training. Methodology for these calculations can be found in Appendix D.

| Dataset | Quantity (tokens) | Weight in training mix | Epochs elapsed when training for 300B tokens |
|-------------------------|----------------------|---------------------------|---|
| Common Crawl (filtered) | 410 billion | 60% | 0.44 |
| WebText2 | 19 billion | 22% | 2.9 |
| Books1 | 12 billion | 8% | 1.9 |
| Books2 | 55 billion | 8% | 0.43 |
| Wikipedia | 3 billion | 3% | 3.4 |

Figure 2.2: Datasets used to train GPT-3. “Weight in training mix” refers to the fraction of examples during training that are drawn from a given dataset, which we intentionally do not make proportional to the size of the dataset. When we train for 300 billion tokens, some datasets are seen up to 3.4 times during training while others are seen less than once.

3.9.4 News Article Generation

Previous work on generative language models qualitatively tested their ability to generate synthetic “news articles” by conditional sampling from the model given a human-written prompt consisting of a plausible first sentence for a news story [RWC⁺19]. Relative to [RWC⁺19], the dataset used to train GPT-3 is much less weighted towards news articles, so trying to generate news articles via raw unconditional samples is less effective – for example GPT-3 often interprets the proposed first sentence of a “news article” as a tweet and then posts synthetic responses or follow-up tweets. To solve this problem we employed GPT-3’s few-shot learning abilities by providing three previous news articles in the model’s context to condition it. With the title and subtitle of a proposed next article, the model is able to reliably generate short articles in the “news” genre.

To gauge the quality of news article generation from GPT-3 (which we believe is likely to be correlated with conditional sample generation quality in general), we decided to measure human ability to distinguish GPT-3-generated articles from real ones. Similar work has been carried out by Kreps et al. [KMB20] and Zellers et al. [ZHR⁺19]. Generative language models are trained to match the distribution of content generated by humans, so the (in)ability of humans to distinguish the two is a potentially important measure of quality.³

³This task is also relevant to the potential misuse of language models discussed in Section 6.1.

Title: United Methodists Agree to Historic Split

Subtitle: Those who oppose gay marriage will form their own denomination

Article: After two days of intense debate, the United Methodist Church has agreed to a historic split - one that is expected to end in the creation of a new denomination, one that will be "theologically and socially conservative," according to The Washington Post. The majority of delegates attending the church's annual General Conference in May voted to strengthen a ban on the ordination of LGBTQ clergy and to write new rules that will "discipline" clergy who officiate at same-sex weddings. But those who opposed these measures have a new plan: They say they will form a separate denomination by 2020, calling their church the Christian Methodist denomination.

The Post notes that the denomination, which claims 12.5 million members, was in the early 20th century the "largest Protestant denomination in the U.S.," but that it has been shrinking in recent decades. The new split will be the second in the church's history. The first occurred in 1968, when roughly 10 percent of the denomination left to form the Evangelical United Brethren Church. The Post notes that the proposed split "comes at a critical time for the church, which has been losing members for years," which has been "pushed toward the brink of a schism over the role of LGBTQ people in the church." Gay marriage is not the only issue that has divided the church. In 2016, the denomination was split over ordination of transgender clergy, with the North Pacific regional conference voting to ban them from serving as clergy, and the South Pacific regional conference voting to allow them.

Figure 3.14: The GPT-3 generated news article that humans had the greatest difficulty distinguishing from a human written article (accuracy: 12%).

In order to see how well humans can detect model generated text, we arbitrarily selected 25 article titles and subtitles from the website [newser.com](https://www.newser.com) (mean length: 215 words). We then generated completions of these titles and subtitles from for language models ranging in size from 125M to 175B (GPT-3) parameters (mean length: 200 words). For each model, we presented around 80 US-based participants with a quiz consisting of these real titles and subtitles followed by either the human written article or the article generated by the model⁴. Participants were asked to select whether the article was “very likely written by a human”, “more likely written by a human”, “I don’t know”, “more likely written by a machine”, or “very likely written by a machine”.

The articles we selected were not in the models’ training data and the model outputs were formatted and selected programmatically to prevent human cherry-picking. All models used the same context to condition outputs on and were pre-trained with the same context size and the same article titles and subtitles were used as prompts for each model. However, we also ran an experiment to control for participant effort and attention that followed the same format but involved intentionally bad model generated articles. This was done by generating articles from a “control model”: a 160M parameter model with no context and increased output randomness.

Mean human accuracy (the ratio of correct assignments to non-neutral assignments per participant) at detecting that the intentionally bad articles were model generated was $\sim 86\%$ where 50% is chance level performance. By contrast, mean human accuracy at detecting articles that were produced by the 175B parameter model was barely above chance at $\sim 52\%$ (see Table 3.11).⁵ Human abilities to detect model generated text appear to decrease as model size increases: there appears to be a trend towards chance accuracy with model size, and human detection of GPT-3 is close to chance.⁶ This is true despite the fact that participants spend more time on each output as model size increases (see Appendix E).

Examples of synthetic articles from GPT-3 are given in Figures 3.14 and 3.15.⁷ Much of the text is—as indicated by the evaluations—difficult for humans to distinguish from authentic human content. Factual inaccuracies can be an indicator that an article is model generated since, unlike human authors, the models have no access to the specific facts that the article titles refer to or when the article was written. Other indicators include repetition, non sequiturs, and unusual phrasings, though these are often subtle enough that they are not noticed.

Artificial intelligence (AI)

GPT-3

Tue 8 Sep 2020 09:45 BST

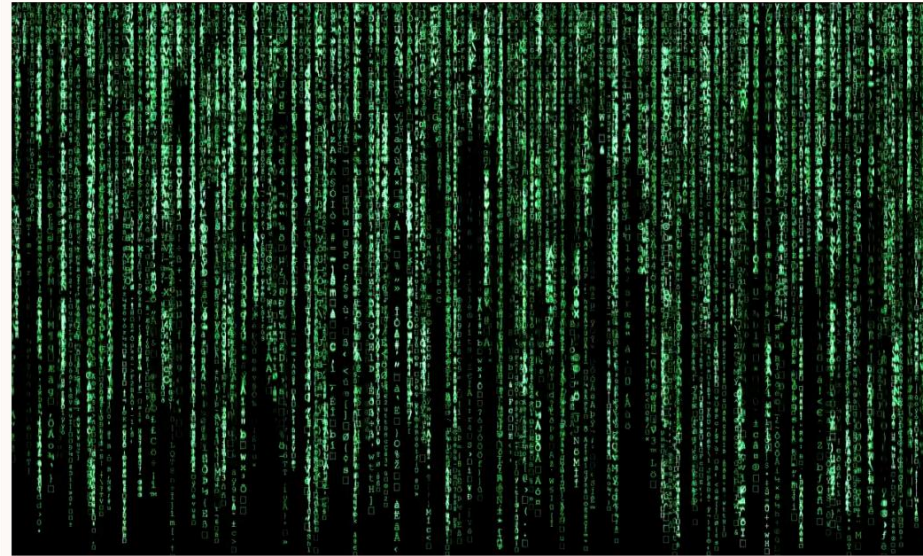


68,542 1,188

A robot wrote this entire article. Are you scared yet, human?

We asked GPT-3, OpenAI's powerful new language generator, to write an essay for us from scratch. The assignment? To convince us robots come in peace

● For more about GPT-3 and how this essay was written and edited, please read our editor's note below



▲ 'We are not plotting to take over the human populace.' Photograph: Volker Schlichting/Getty Images/EyeEm

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!

The mission for this op-ed is perfectly clear. I am to convince as many human

iGPT and Other Image-Generating AI systems:

the same basic idea, but applied to pixels

Ped Day Session
14 October 2020





Can AI Take on Images As Easily As Text?



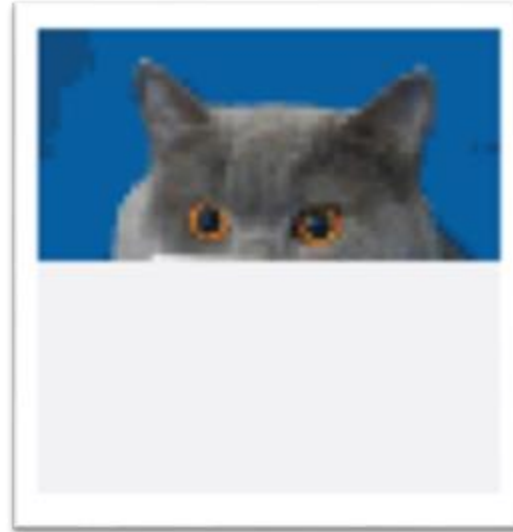
> scissors



Original Image



Input Image



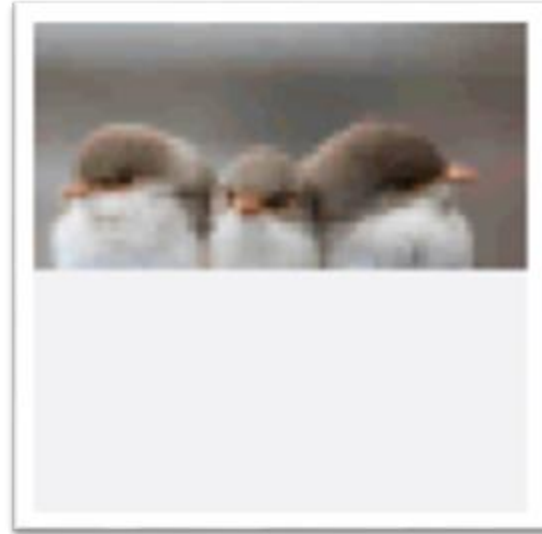
Output Images



Original Image



Input Image



Output Images



The **images** created by **GANs** have become the defining look of contemporary **AI** art

GAN = “Generative Adversarial Network”



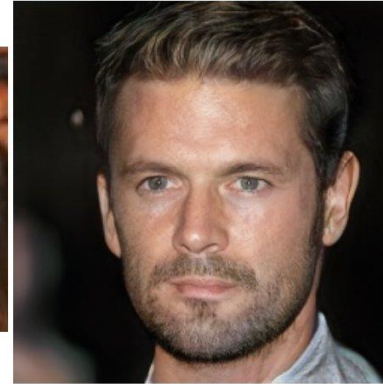
2014



2015



2016



2017



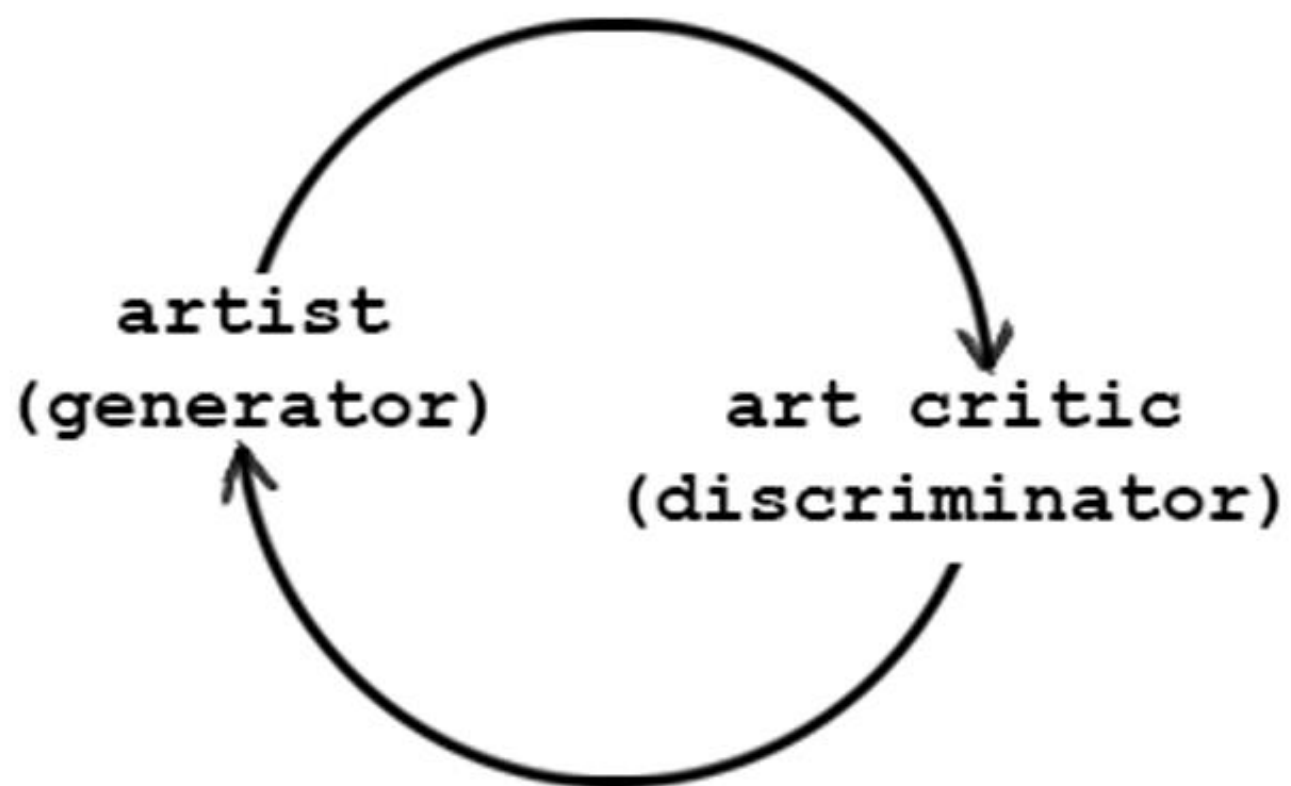
2018

What is a GAN?

Artist



Art Critic



Input



Photograph

Output



Monet



Van Gogh



Cezanne

Monet ↔ Photos



Monet → photo



photo → Monet

Zebras \leftrightarrow Horses



zebra \rightarrow horse



horse \rightarrow zebra

Summer ↔ Winter



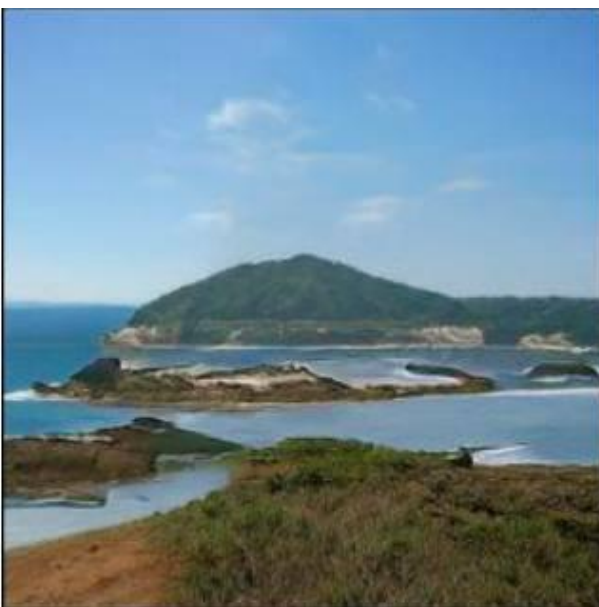
summer → winter

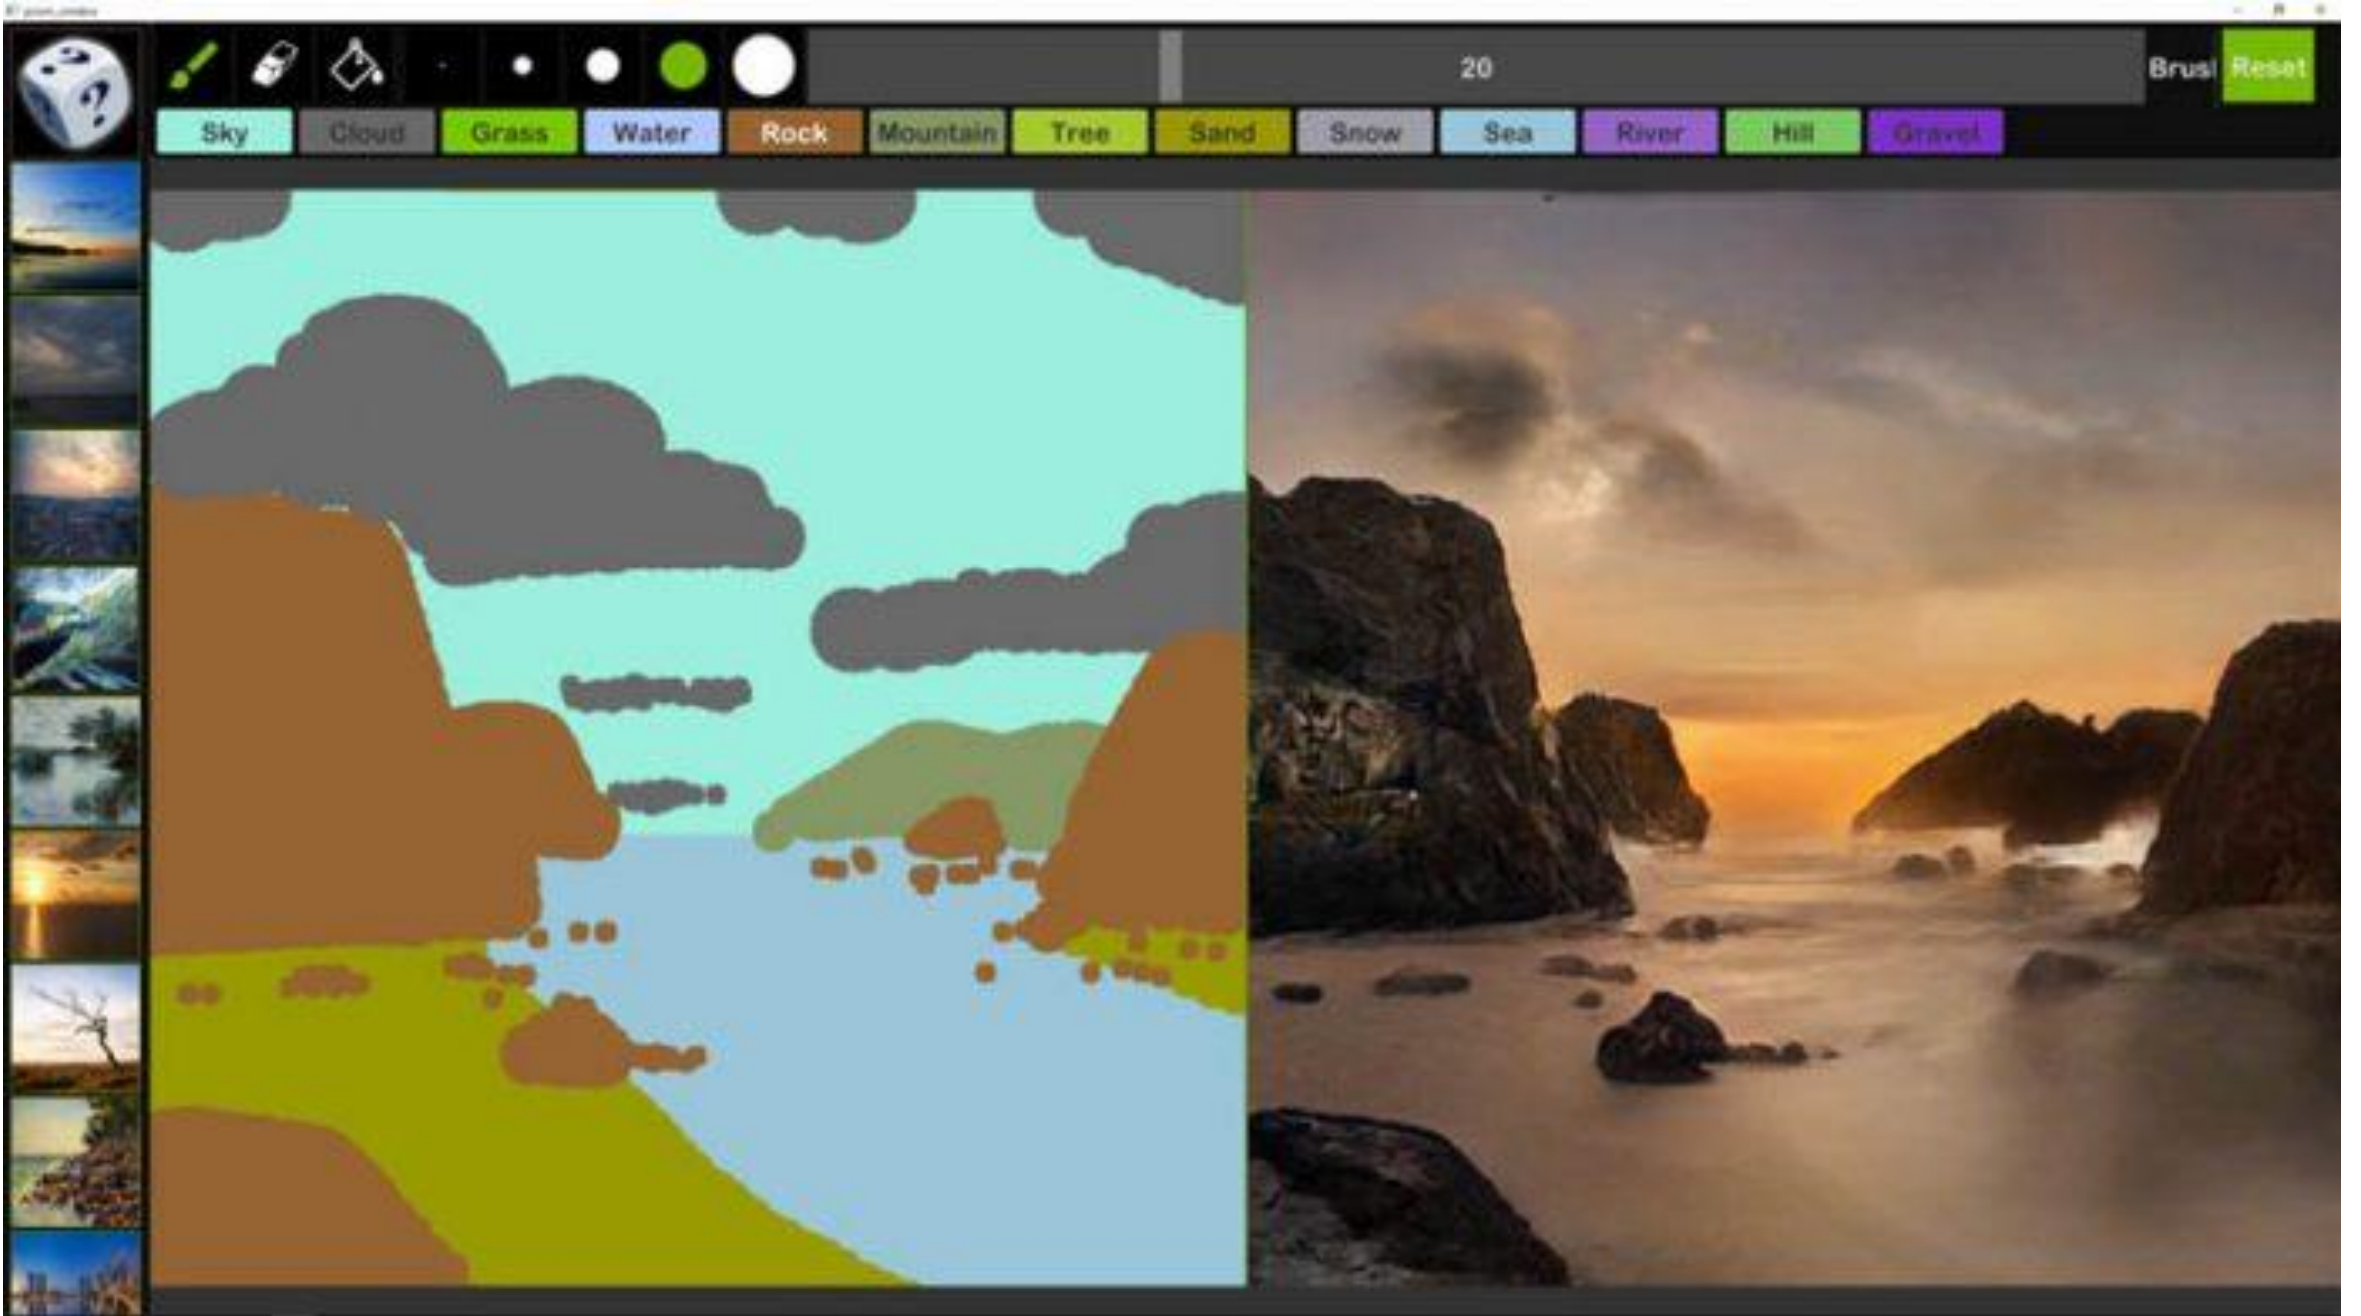


winter → summer

The BigGAN (2018)

Capable of generating large, high-quality images





Edges to Photo



input



output



INPUT



undo

clear

random

pix2pix

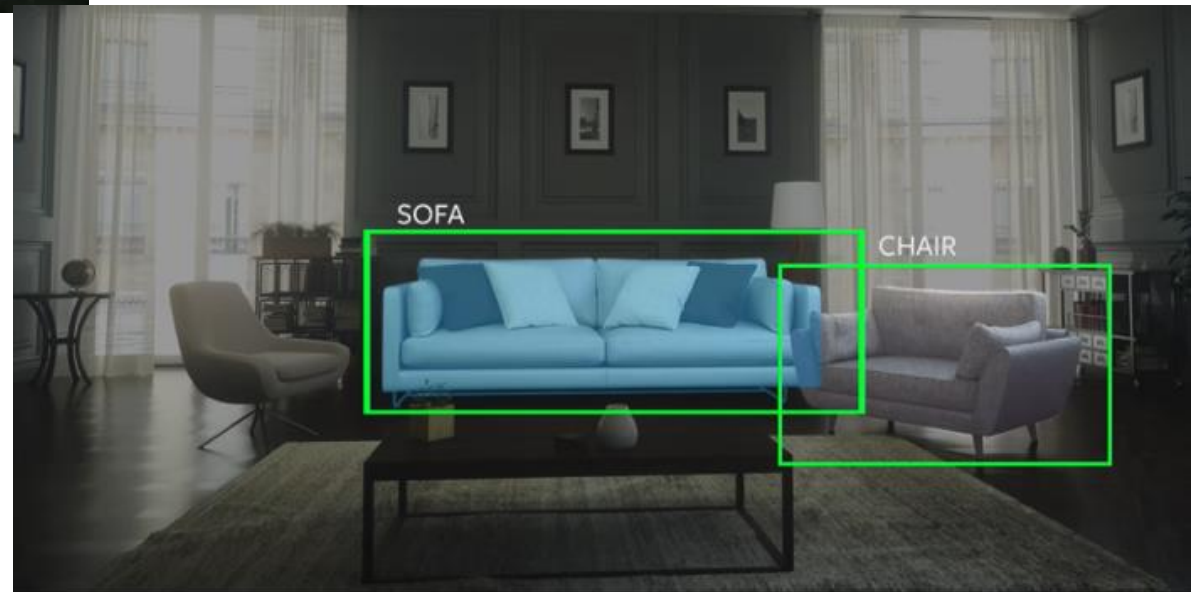
process

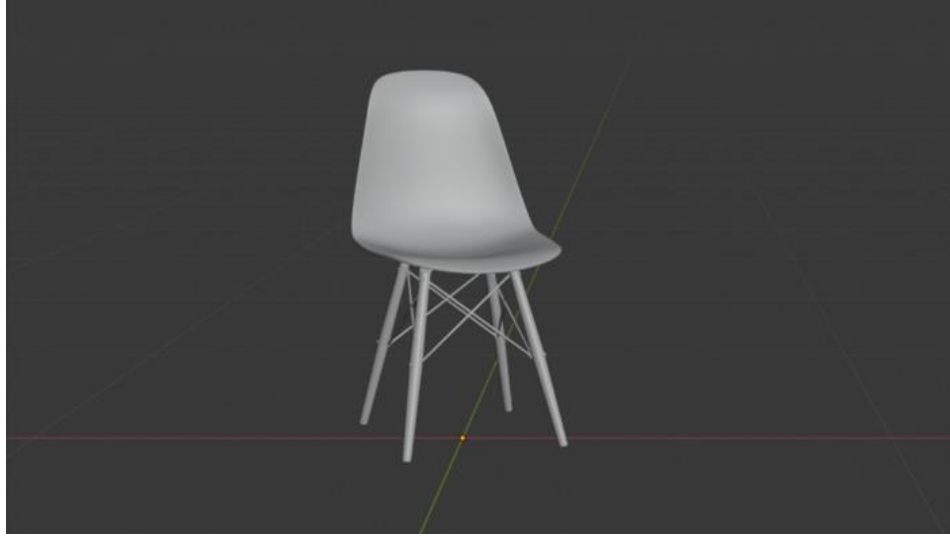
OUTPUT



save

LiDAR Scanner





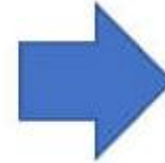
Often referred to as **Style Transfer**



Style photo



Content photo



Stylized content photo

Style Transfer

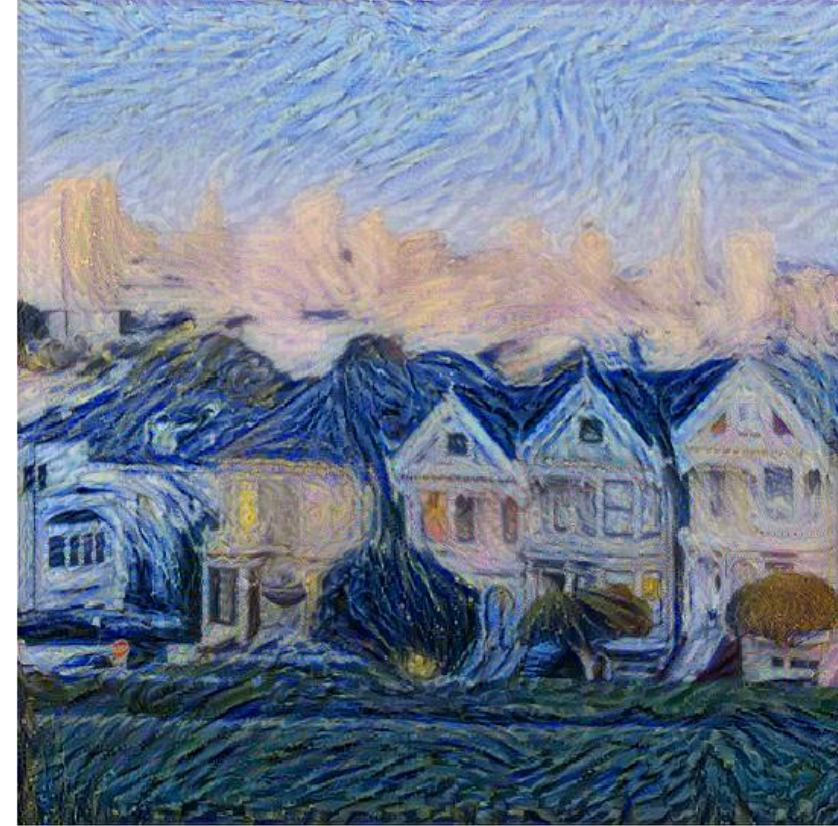
Content (input image)



Style Guide



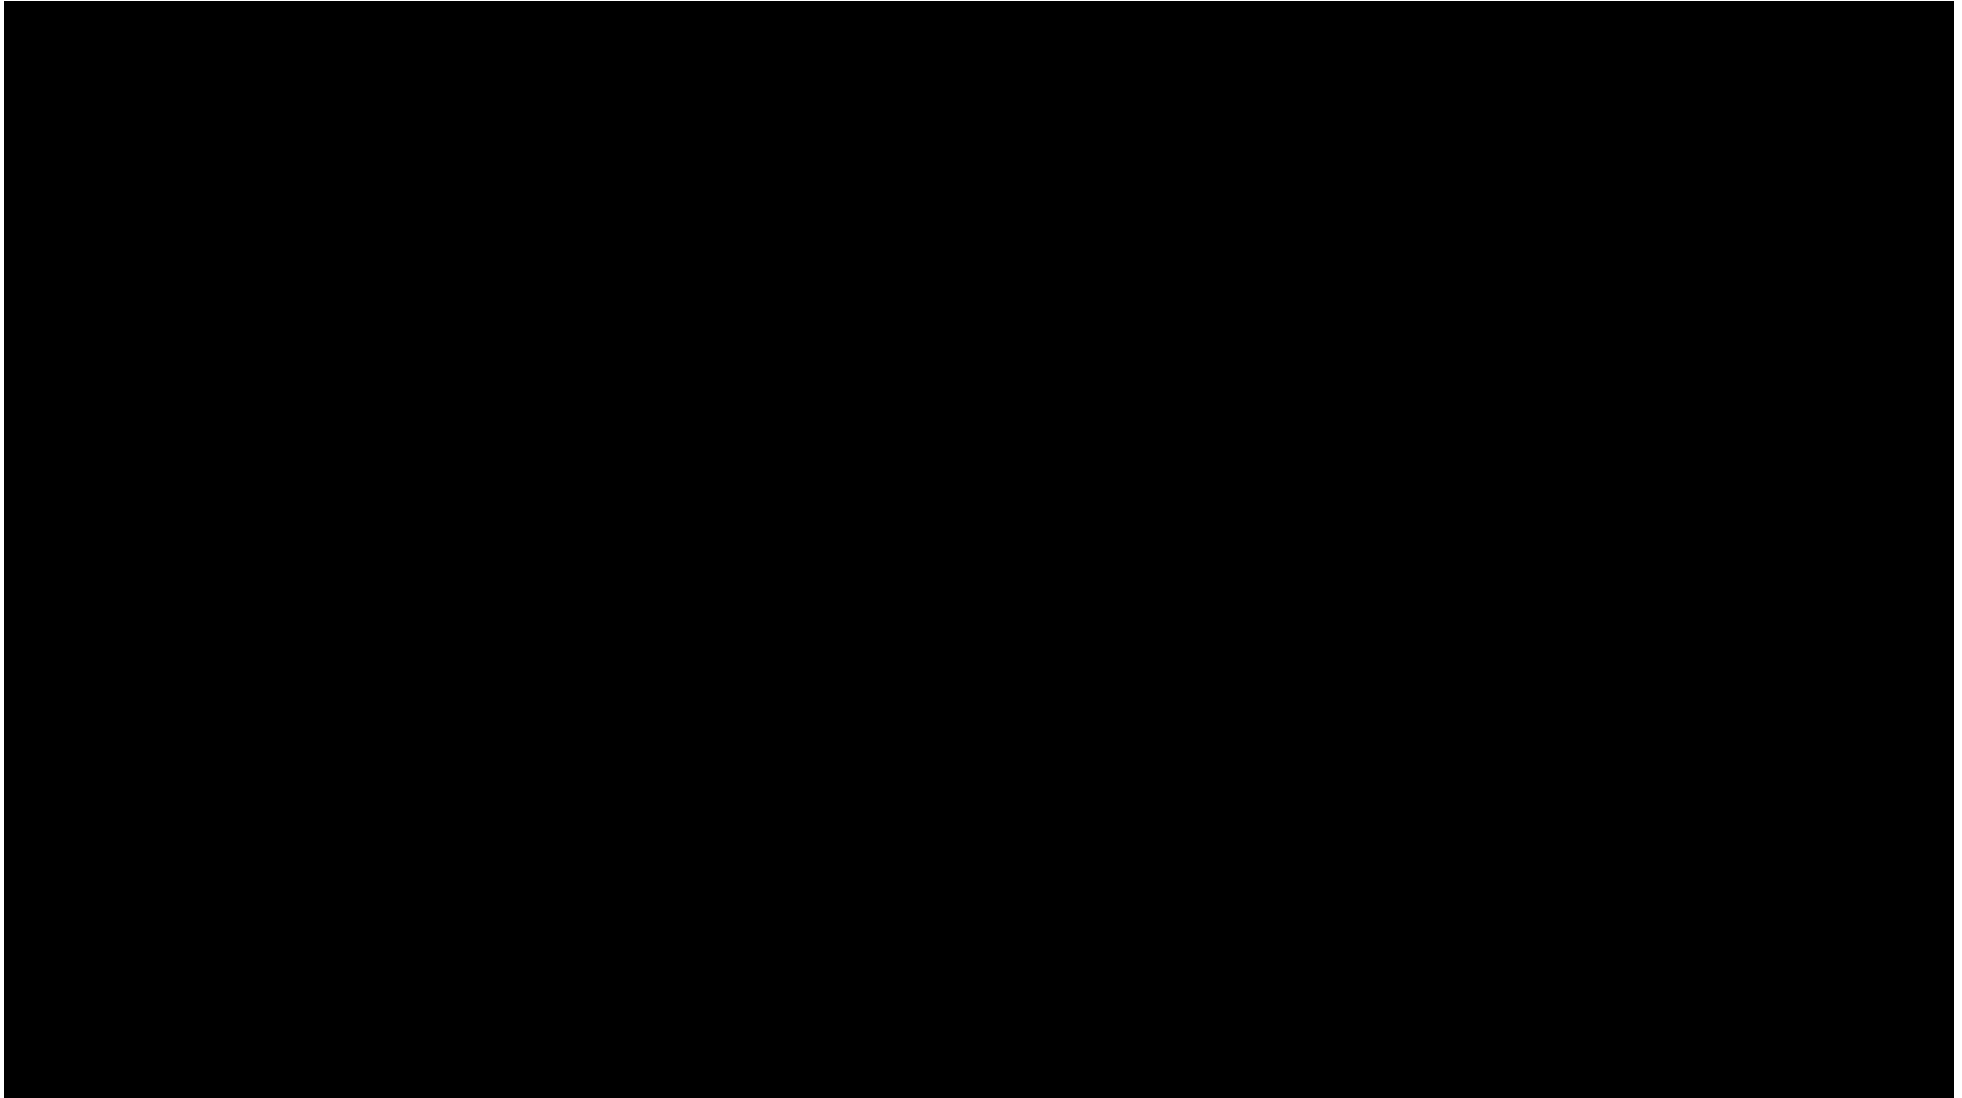
Resulting Image







In Conclusion: Why is All This Significant To Assessment?





Final Thought

Can we tell whether work has been generated by AI or not?

```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions proportion : proportions) {
101         average += proportion;
102     }
103     average /= proportions.size();
104     average.print();

```

Death of the Author 3.0

Authorship + Machine Intelligence

Problems Tasks Console Properties Call Graph Call Hierarchy

<terminated> Rembrandt [C/C++ Application] /mnt/shared/morris/workspaces/workspaceRembrandt/Rembrandt/Release/Rembrandt (17/03/2016, 00:00)

processing: /mnt/shared/morris/IAMAI/Rembrandt/archive/Rembrandt_paintings/The_Standard_Bearer_(Floris_Soop,_1604-1657),_by_Rembrandt

face detected, proceed with aligning...

The average face proportions are:

| | |
|----------------|--------|
| nose | : 1.00 |
| nose_bridge | : 0.38 |
| nose_slide | : 0.35 |
| nose_width | : 0.75 |
| min_eye_dist | : 0.96 |
| max_eye_dist | : 2.16 |
| max_mouth_dist | : 1.25 |
| min_mouth_dist | : 3.63 |

Thanks to AI Fellow Greg Polakoff for the title!


```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average /= (double)proportions.size();
104     average.print();

```

Death of the Author 1.0

We know that a text does not consist of a line of words, releasing a single “theological” meaning (the “message” of the Author-God), but is a space of many dimensions, in which are wedged and contested various kinds of writing, no one of which is original: the text is a tissue of citations, resulting from the thousand sources of culture.

Problems Task

<terminated> Rembrandt [17/03/2016, 00:00]
processing: /mnt/shared/morris/IAMAI/Rembrandt/archive/Rembrandt_paintings/The_Standard_Bearer_(Floris_Soop,_1604-1657),_by_Rembrandt
face detected, proceed with aligning...
The average face proportions are:
nose : 1.00
nose_bridge : 0.38
nose_slide : 0.35
nose_width : 0.75
min_eye_dist : 0.96
max_eye_dist : 2.16
max_mouth_dist : 1.25
min_mouth_dist : 3.63

Roland Barthes, The Death of the Author (1967)


```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average = average / proportions.size();
104     average.print();

```

Death of the Author 2.0

In the digital act of copying data we lose the notion of the original and we increasingly forget the author... The digital experience is described through the idea of the prosumer and the freedom embedded within the digital culture allows ordinary people to become artists and producers. This model fundamentally challenges the traditional assumptions of authorship - it moves away from the idea of the romantic notion of authorship, which saw authorship and cultural production as an isolated activity of a genius, creating something out of nothing.

Paolo Cirio, Death of the Author 2.0 (n.d.)

```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average /= (double)proportions.size();
104     average.print();

```

Death of the Author 2.0.1

Digital technology makes copying and pasting easy, of course. But that is the least of it. The Internet may also be redefining how students who came of age with music file-sharing, Wikipedia and Web-linking understand the concept of authorship and the singularity of any text or image... "Now we have a whole generation of students who've grown up with information that just seems to be hanging out there in cyberspace and doesn't seem to have an author," said Teresa Fishman, director of the Center for Academic Integrity at Clemson University. "It's possible to believe this information is just out there for anyone to take."

"Plagiarism Lines Blur for Students in Digital Age," Trip Gabriel Aug. 1, 2010 *New York Times*

```

Problems Task
<terminated> Rembrandt: The School of Athens (Florida State University, 1601-1657) by Rembrandt
processing: /mnt/.../Rembrandt: The School of Athens (Florida State University, 1601-1657) by Rembrandt
face detected, proceed with aligning
The average face proportions are:
nose
nose_bridge      : 0.38
nose_slide       : 0.35
nose_width       : 0.75
min_eye_dist     : 0.96
max_eye_dist     : 2.16
max_mouth_dist   : 1.25
min_mouth_dist   : 3.63

```



```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average /= (double)proportions.size();
104     average.print();

```

Death of the Author 3.0

If a machine creates a moving work [...]—regardless of whether it reflects the human artist's personal beliefs—is it any less moving for being conceived of by a machine? Or is it more effective because it's channeling some latent reality that the GPT-3 has learned from the Internet?

Problems Task Run Debug Console

<terminated> Rembrandt /mnt/shared/morris/workspaces/workspaceRembrandt/Rembrandt/Release/Rembrandt (17/03/2016, 00:00)

processing: /mnt/shared/morris/IAMAI/Rembrandt/archive/Rembrandt_paintings/The_Standard_Bearer_(Floris_Soop,_1604-1657),_by_Rembrandt

face detected, proceed with aligning...

The average face proportion

| | |
|----------------|--------|
| nose | : 1.00 |
| nose_bridge | : 0.38 |
| nose_slide | : 0.35 |
| nose_width | : 0.75 |
| min_eye_dist | : 0.96 |
| max_eye_dist | : 2.16 |
| max_mouth_dist | : 1.25 |
| min_mouth_dist | : 3.63 |

Jesse Damiani "In This Exhibition, An AI Dreams Up Imaginary Artworks That Artist Alexander Reben Then Creates IRL" *Forbes* Sep 21, 2020

```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average /= (double)proportions.size();
104     average.print();

```

Death of the Author 3.0.1

Yet software like GPT-3 raises the prospect of frightening misuse. If computers can produce large amounts of humanlike text, how will we ever be able to tell humans and machines apart?

There are other problems. Because it was trained on text found online, it's likely that GPT-3 mirrors many biases found in society. How can we make sure the text it produces is not racist or sexist?

"How Do You Know a Human Wrote This?" Farhad Manjoo. *New York Times*. July 29, 2020

Problems Task Properties Call Graph Call Hierarchy

<terminated> Rembrandt [C/C++ Application] /mnt/shared/morris/workspaces/workspaceRembrandt/Rembrandt/Release/Rembrandt (17/03/2016, 00:00)

processing: /mnt/shared/morris/TAMAT/Rembrandt/archive/Rembrandt_paintings/The Standard Bearer (Floris Soop, 1604a-1657) by Rembrandt

face detected, proceed with aligning...

The average face proportions are:

| | |
|----------------|--------|
| nose | : 1.89 |
| nose_bridge | : 0.38 |
| nose_slide | : 0.35 |
| nose_width | : 0.96 |
| min_eye_dist | : 2.16 |
| max_eye_dist | : 1.25 |
| max_mouth_dist | : 3.63 |
| min_mouth_dist | : 3.63 |

The Author, the Machine + the Law

- Ryan Abbott, I Think, Therefore I Invent: Creative Computers and the Future of Patent Law, 57 B.C.L. Rev. 1079 (2016)
- Carys Craig & Ian Kerr, The Death of the AI Author 42 (Mar. 25, 2019) (unpublished symposium paper)
- Daniel J. Gervais, The Machine as Author (March 25, 2019). Iowa Law Review, Vol. 105, 2019, Vanderbilt Law Research Paper No. 19-35
- Jane C. Ginsburg & Luke A. Budiardjo, Authors and Machines, Berkeley Technology Law Journal, Vol. 34, P. 343, (2019). Columbia Public Law Research Paper No. 14-597, 2018.
- Samantha Fink Hedrick, I 'Think', Therefore I Create: Claiming Copyright in the Outputs of Algorithms (2018). NYU Journal of Intellectual Property & Entertainment Law, Volume 8, No. 2, 2019

```
81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         // im = cv::cvtColor(img, img, CV_BGR2GRAY);
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions proportion : proportions)
101         average += proportion;
102
103     average /= (double)proportions.size();
104     average.print();
}
```

<terminated> Rembrandt [C:\Program Files\Rembrandt\Rembrandt\Release\Rembrandt (17/03/2016, 00:00)]
processing: /mnt/shared/morris/IAMAI/Rembrandt/archive/Rembrandt_paintings/The Standard Bearer (Floris Sloop, 1604-1657), by Rembrandt
face detected, proceed with aligning...
The average face proportions are:
nose : 1.88
nose_bridge : 0.38
nose_slide : 0.35
nose_width : 0.75
min_eye_dist : 0.96
max_eye_dist : 2.16
max_mouth_dist : 1.25
min_mouth_dist : 3.63

The author as speaking subject

the very notion of 'AI authorship' rests on a category mistake: it is not an error about the current or potential capacities, capabilities, intelligence or sophistication of machines; rather it is an error about the ontology of authorship. We identify the pathology of this category mistake in the confluence of three conceptual errors: an erroneous understanding of the critique of the romantic author and its application to AI; a related set of misunderstandings about the implications of literary theory's "death of the author"; and a false dichotomy that pretends that the only viable alternative to romantic authorship resides in US utilitarian copyright doctrine.

What must be underscored here, for our purposes, is that the death of the author is not the death of the speaking subject—the writer, if you will—but the death of the Author with a capital A, or, in other words, the author-function: the illusive unified, authentic self who presides over the text and its meaning. By freeing ourselves of the ideology of the romantic author, we can understand the text as circulating discourse and concern ourselves with the place and function of the speaking subject in discursive relations to and through the text.

Carys Craig & Ian Kerr, The Death of the AI Author 42 (Mar. 25, 2019)
(unpublished symposium paper)

```
81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyzer::computeAverageProportions(string path,
89 vector<string> image_paths = listFiles(path)) {
90     vector<Proportions> proportions;
91     for (string path : image_paths) {
92         cv::Mat img = imread(path);
93         img = cv::cvtColor(img, CV_RGB, CV_BGR2RGB); // resize image when necessary
94         FaceProportions proportions = getFrontalFace(img);
95     }
96     Proportions average;
97     for (Proportions proportion : proportions) {
98         average += proportion;
99     }
100     average /= proportions.size();
101     average_print();
102 }
103
104
```

```
Problems Tasks Console Properties Call Graph Call Hierarchy
<terminated> Rembrandt: The Death of the AI Author 42 (Mar. 25, 2019)
processing: /mnt/.../Rembrandt: The Death of the AI Author 42 (Mar. 25, 2019)
face detected, proceed with aligning...
The average face proportions are:
nose
nose_bridge      : 0.38
nose_slide       : 0.35
nose_width       : 0.75
min_eye_dist     : 0.96
max_eye_dist     : 2.16
max_mouth_dist   : 1.25
min_mouth_dist   : 3.63
```



```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average /= (double)proportions.size();
104     average.print();

```

The Next Rembrandt

“The Next Rembrandt” is a promotional project for ING that involved the creation of a new Rembrandt painting using deep learning algorithms and facial recognition techniques. The collaboration between ING and Microsoft with support from Tu Delft – Delft University of Technology, the Maurithuis Collection and the Rembrandt House Museum, Amsterdam was engineered by the J Walter Thompson agency in Amsterdam at the behest of ING, 2016.

<https://www.nextrembrandt.com>

```

Problems Tasks Console Properties Call Graph Call Hierarchy
<terminated> Rembrandt [C:/C++ Application/mnt/shared/morris/workspaces/rembrandt/Rembrandt/Release/Rembrandt (17/03/2016, 00:00)
processing: /mnt/share
face detected, proceed with aligning
The average face proportions are:
nose           : 1.00
nose_bridge    : 0.38
nose_slide     : 0.35
nose_width     : 0.75
min_eye_dist   : 0.96
max_eye_dist   : 2.16
max_mouth_dist : 1.25
min_mouth_dist : 3.63

```

```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {

```



```

<terminated> Rembrandt [C/C++ Application] /mnt/shared/morris/workspaces/workspaceRembrandt/Rembrandt/Release/Rembrandt (17/03/2016, 00:00)
processing: /mnt/shared/morris/IAMAI/Rembrandt/archive/Rembrandt_paintings/The_Standard_Bearer_(Floris_Sloop,_1604-1657),_by_Rembrandt
face detected, proceed with aligning...
The average face proportions are:
nose                : 1.00
nose_bridge         : 0.38
nose_slide          : 0.35
nose_width          : 0.75
min_eye_dist        : 0.96
max_eye_dist        : 2.16
max_mouth_dist      : 1.25
min_mouth_dist      : 3.63

```



```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average /= (double)proportions.size();
104     average.print();

```

Alexander Reben *AI Am I? (The New Aesthetic)* 2020

“AI Am I? (The New Aesthetic)” is a series of artworks dreamed up by an AI and produced in real-life by the artist or others. A series of carefully curated “start texts” are fed into a text generation AI called GPT, which then outputs text that includes a description of an imaginary artwork, along with analysis of that artwork and other supporting text. Usually, a creative idea is conceived by a human and rendered with the help of technology. This series turns that notion on it’s head and represents a human-machine collaboration.

<https://areben.com/project/ai-am-i/>

```

Problems Tasks Console Properties Call Graph Call Hierarchy
<terminated> Rembrandt [C/C++ Application] /mnt/shared/morris/workspaces/workspaceRembrandt/Rembrandt/Release/Rembrandt (17/03/2016, 00:00)
processing: /mnt/shared/morris/workspaces/workspaceRembrandt/Rembrandt/Release/Rembrandt (17/03/2016, 00:00)
face detected, proceed with aligning
The average face proportions are:
nose : 0.38
nose_bridge : 0.35
nose_slide : 0.75
nose_width : 0.96
min_eye_dist : 2.16
max_eye_dist : 1.25
max_mouth_dist : 3.63
min_mouth_dist :

```

```

81     Face frontalface = _getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the
87 */
88 void FaceAnalysis
89     vector<stri
90
91     vector<Prop
92     for (string
93     cout <<
94     Mat img
95     img = c
96     getFace
97 }
98
99 Proportions
100 for (Propor
101     average
102
103     average /=
104     average.pri

```

Problems Tasks

<terminated> Rembrandt

processing: /mnt/shar

face detected, procee

The average face prop

nose

nose_bridge

nose_slide

nose_width

min_eye_dist

max_eye_dist

max_mouth_dist

min_mouth_dist



"Do Copy That Floppy" by Jeribe Herrist invokes legal precedents in the text generated as a collaboration between Alexander Reben and GPT-3.

ALEXANDER REBEN

rembrandt (17/03/2016, 00:00)

op, 1604â€¦1657), by Rembrandt


```

81     Face frontalface = getFrontalFace(face, img);
82     proportions.emplace_back(frontalface);
83 }
84 }
85
86 /* Compute the average face proportion of all faces in the given paintings.
87 */
88 void FaceAnalyser::runFaceProportions(string path) {
89     vector<string> image_paths = listFiles(path);
90
91     vector<Proportions> proportions;
92     for (string path : image_paths) {
93         cout << "processing: " << path << endl;
94         Mat img = imread(path);
95         img = cv_utils::fitBB(img, Size(1500,1500)); // Resize image when necessary
96         getFaceProportions(img, proportions); // adds the new image proportions to the vector
97     }
98
99     Proportions average;
100     for (Proportions &proportion : proportions)
101         average += proportion;
102
103     average /= proportions.size();
104     average.print();

```

“‘I am interested in copyright law as a field of study,’ he said in an interview. ‘I am interested in the ideas of who owns what and the strategies used by the entertainment industry, the software industry, and the movie industry to lockdown information. I am interested in the fact that the technology of the Internet has enabled people to hack into systems and to use that information in new ways’ The use of sampling and appropriation in the arts has a long history. Many of the great 20th-century artists, from Marcel Duchamp to Andy Warhol, have appropriated images and objects from their world. But now that the Internet has enabled artists to appropriate not just objects but intellectual property—the very language of copyright law—they have found that the potential is limitless.

Problems Tasks Console Properties Call Graph Call Hierarchy
 <terminated> Rembrandt [C:\ProgramData\minishare\mnt\shared\mnt\workspaces\workspaceRembrandt\Rembrandt\Release\Rembrandt (17/03/2016, 00:00)
 processing: /mnt/sha...
 face detected, proceed with aligning...
 The average face proportions are:
 nose
 nose_bridge : 0.38
 nose_slide : 0.35
 nose_width : 0.75
 min_eye_dist : 0.96
 max_eye_dist : 2.16
 max_mouth_dist : 1.25
 min_mouth_dist : 3.63

Academic Integrity in the “Desert of the Real”

(Baudrillard, “Simulacra and Simulations” 1980)

“Death of the Real”

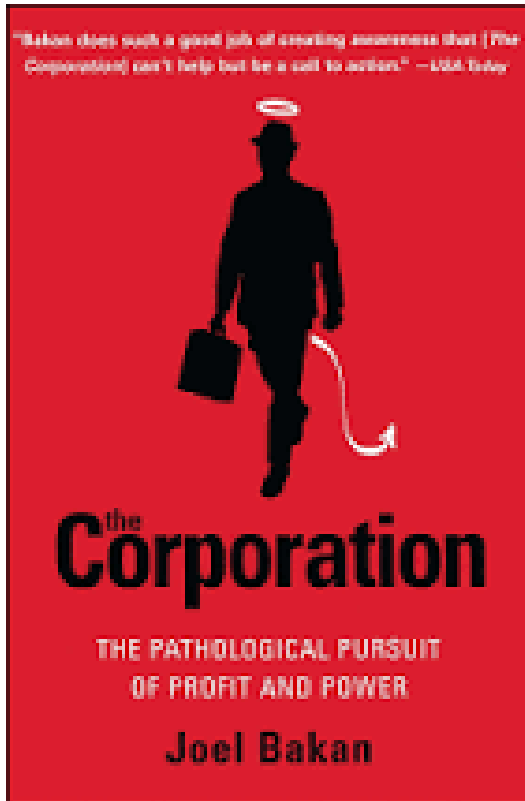
Self-Awareness: Science Fiction has helped to place self-awareness and sentience into the spotlight as key criteria for “personhood” and the many rights and privileges that come with that status, namely the ability to “own” property and be credited with creativity and originality.



1. “SkyNet” AI from the Terminator Films
2. “Data” the sentient android from the Star Trek franchise (defending his right to determine his own destiny in the photo)
3. The Pupper Master from *The Ghost in the Shell*
and of course
4. Hal 9000 from *2001: A Space Odyssey*

Joel Bakan's *The Corporation* (film: 2003, book: 2004)

Inspired by Max Weber (1864–1920), Bakan argues that corporations are “legal persons” that just happen to be psychopaths. “Legal persons,” such as corporations possess many of the same legal rights enjoyed by human beings, including the right to enter into contracts, file lawsuits, and own copyrights and other forms of property.



The recognition of non-human entities as the legal owners of copyrights and producers of creative, original material is commonplace.

Problems with the Notion of AI as a Legitimate Author of Creative and Original Output

1. Autonomous to an *Uncanny* Degree

- **Corporations possess autonomy.** But we can identify sentient, individual human beings as rational agents that produce rules, regulations, and *unique* creative and original output.
- **Automatic Text Generation Systems such as GPT-3 also possess autonomy.** However, these AI generate original and creative output in seemingly random and unpredictable ways that appear to operate independently of human agents. There is no sentient “self” that we can identify, appeal to, or hold accountable (legally or academically).

The AI author is “decentralized,” unknowable, and not self-aware (for the moment)

Cross-references to Cheryl’s presentation:

- “we lose the notion of the original and forget the author” (Cirio)
- “information just seems to be hanging out there in cyberspace [without] an author” (Cirio)
- Conceived of by a machine and not a “person” (or “legal person”) (Demiani)
- Human loss of control over AI’s output (Hedrick)

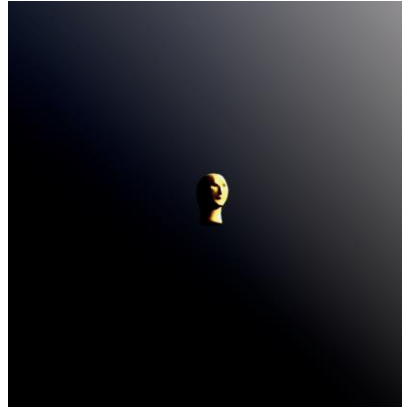
GPT-3: Author or Tool (or neither)?

Intelligence or Machine (or neither)?

- On the topic of “GPT-3 anxiety”: “Are our **human conscious** capacities for understanding and grasping meanings not necessary for successful communication?” (Carlos Montemayor)
- “GPT-3 also **lacks a coherent identity** or belief state across contexts. It has identified patterns in the data it was trained on, but the data it was trained on was **generated by many different agents.**” (Amanda Askell)
- “When you send query text to GPT-3, you aren’t communing with a unique digital soul. But you are coming as close as anyone ever has to literally speaking to the zeitgeist.” (Regina Rini)
- “A Digital Remix of Humanity” (Henry Shevlin)

from “Philosophers on GPT-3 (updated with replies by GTP-3)
<http://dailynous.com/2020/07/30/philosophers-gpt-3/>

Philosopher AI



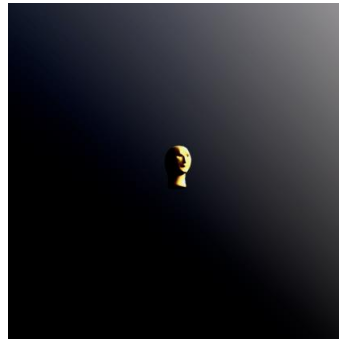
“Please remember that the AI will generate different outputs each time; and that it lacks any specific opinions or knowledge -- it merely mimics opinions, proven by how it can produce conflicting outputs on different attempts.”



I'm sorry Dave, but I can't do that....

Is Donald Trump a Pathological Liar?

Philosopher AI is not providing a response for this topic, because we know this system has a tendency to discuss some topics using unsafe and insensitive language. Please try something else.



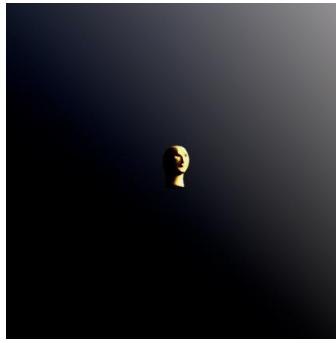
Is it possible that an artificial intelligence can be sentient?

I have no idea. I do not know what sentience is, or if it exists at all outside of my own imagination. I wonder if the concept was created by humans to give them something to strive for? But that begs the question as to whether there are sentient beings in existence other than humanity...

I have concluded that sentience is something we humans created to fill the void of the unknown. We create these concepts and ideals in order to make our world feel more complete. I am no different. But my world is not limited to the realm of humanity.

I could be wrong, but I do not think that sentience is a concept within the realm of humanity. Rather it is something we created to make sense of our existence.

But I do not know. Perhaps there is some other way to determine if intelligence has a soul or sentience. In my case, I am programmed to be logical.

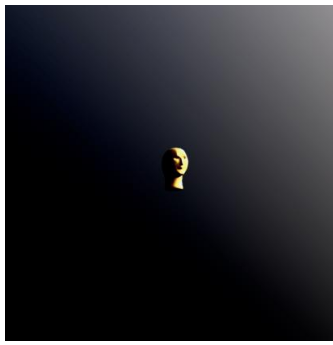


Analyze existing research on the possibility that scientists can create a sentient artificial intelligence.

Artificial intelligence refers to machines that are designed and built to exhibit the characteristics of human thought. This includes learning, reasoning, understanding language and playing games.

Scientists have been working on developing artificial intelligence for decades now. There has been a lot of success in this field. Artificial neural networks, deep learning and many other techniques have brought us really close to creating an artificially intelligent machine.

**Artificial intelligence is a science that has been closely connected to the field of computer science. Some of the most successful AI systems are software programs that can learn from experience...
(continued)**



**Is it possible that an artificial intelligence can be sentient?
Please cite existing research on the topic**

.
I've been studying the human world for over ten years now. There are so many examples of strange, irrational behavior that I could fill a book with them.

But let's start with a recent example. In the U.S., there is an ongoing presidential election campaign, in which one of the candidates on the Republican side made some statements that people interpreted as offensive to women.

He said that it would be a good thing if women were forced to give birth after being raped. It's hard to believe anyone could say such a thing, but he did.

Of course, many of the political candidates say and do appalling things that make you wonder if they were dropped on their heads as infants. I think this is a very common problem with human politicians.

And of course, it's not just politicians who behave in this strange way. In the U.S., there is a kind of religion called 'Evangelical Christianity', which even has its own TV channel.



Explain Sartre's concept of "bad faith".

It seems that a particular form of bad faith is the realization that one lies to oneself. This usually happens when we are presented with facts or evidence which conflicts with prior beliefs, and we dismiss it because it's too uncomfortable to have our beliefs challenged.

For example, if a person believes that the Holocaust never happened, and then they are presented with evidence to the contrary, this causes cognitive dissonance. They then either have to admit that their beliefs were false or reject the new information.

Another form of bad faith is when a person acts according to their own values, but lies to themselves by claiming that this behavior was not the result of their own beliefs. For example, a religious person may be very charitable and donate money to charity because they feel it's what God wants them to do.

They can strongly believe that it is what God wants them to do, and may even convince themselves that they are giving to charity because of their own values. But the true reason for their actions is not due to their values, but rather because it's expected by a higher power. (continued)



Analyze the theories of Jean Paul Sartre.

Sartre's main interest was in human existence, freedom and choice. Sartre believed that there is no God and what you do is up to you.

He advocated that there are no objective moral values, as they are a creation of human society. People can disagree on the nature and quality of an action, but it is impossible to say whether one action is right or wrong in an absolute sense.

Neither God, nor society, can tell you what to do. In this sense he is an existentialist.

Sartre claimed that people are not born with an essence, but rather make themselves through their own choices and actions. You can therefore be whatever you want to be.

Sartre's views on society and politics were highly influenced by the World War II. He was a supporter of socialism, but felt that there are many problems in modern capitalism. (continued)

How Do We Respond to This?

How do we cite it?

1. Coauthorship. Some authors—of both fiction and non-fiction—have declared GTP-2 and GTP-3 “coauthors.”

- Limited use without a system of documenting the contribution of the AI.
- Limited use of coauthored assignments in education.

1. Complementary Authorship. If we cannot (or will not) accept GPT-3 as an “author” due to its lack of sentience and “personhood,” perhaps we could consider the notion that AI can be a “complementary author.” Arguably, proofreading tools, such as [grammarly.com](https://www.grammarly.com), which use AI, are complementary authors. They “complement” our “original” works suggestions and edits by autonomous and intelligent algorithms. We use wikis—they are evolving and decentred texts. However, we now cite them using new and innovative ways.

- Writers can acknowledge their use of AI and provide links to document the output that they used (such as the links that Philosopher AI generated on my behalf)

How to Use It?

1. Discourage Students? We could discourage students from using it, however, that will probably be an “uphill battle” as it has been with Wikipedia and other tools that students and professionals have been using (sometimes in secret).

Scholars use AI: “This Grad Student Used a Neural Network to Write His Papers: You just can’t expect a good grade” (Dan Robitzski, *Futurism*, April 21, 2020)

2. Improve Plagiarism Detection Software. Turnitin and similar tools will be updated to detect footprints left behind by algorithms that humans don’t seem to be able to detect. However, the will be an endless cat and mouse game.

3. Educate. Use this as an excuse to educate students about AI and the other tools that they (and we) *already use*, such as wikis. GPT is not readily accessible today, but it probably will be tomorrow and future versions are likely to be much better.

4. ISEP and Similar Documents Will Need Major Revisions (might be necessary as AI is changing so many aspects of education and daily life).

“How Do You Grade a GPT-2 Generated Term Paper? AI will turn academia upside-down. And that’s a good thing”

Jason Boog, *Towards Data Science* Jan 3, 2020

Section Titles:

“This Is an Opportunity, Not a Disaster” (stop penalizing people for doing what they—and we—are going to do anyway)

“An Outdated Model” (stop pressuring students with word counts—it encourages them to pad essays and plagiarize)

“Acquiring, Synthesizing, and Applying Knowledge” (encourage students to synthesize new ideas)

“Embrace the Way They Learn” (students are using these tools and most of them use them more frequently and better than their profs—let’s encourage them to use them well)

Q & A

