

Infinite Machine Creativity

Menno van Doorn
Sander Duivestijn
Deepa Mamtani
Thijs Pepping



Infinite Machine Creativity

1	Introduction: the creative machine	3
1.1	Innov-AI-tion: a medicine in no time	4
1.2	GANs: behind the scenes	6
1.3	Pharma's AlphaGo moment	8
1.4	The structure of the report and fifteen GAN conclusions	9
2	Compression of time-to-create	12
2.1	Science: analyze, predict, synthesize	12
2.2	Design, fashion, and style	18
2.3	Design in architecture and industry	20
2.4	Synthetic media and art: fake is the new real	23
2.5	In conclusion: the short nose of innovation	32
3	HUMAN-GAN-creativity	34
3.1	Are computers actually creative?	34
3.2	HUMAN-GAN: a closer look	35
3.3	A new collaboration, to what level?	36
3.4	Finally: human creativity revealed	37
4	In conclusion: a future CIOs dream of	38
	Notes	44
	Image credits	46

We thank the Oxford Internet Institute (OII) for the insights and ideas gained from a joint workshop, especially Luciano Floridi (Director of Research, Professor of Philosophy and Ethics of Information), Mariarosaria Taddeo (Senior Researcher Fellow OII), Paul Timmers (Visiting Research Fellow Cybersecurity Policy and Digital Transformation), Matthias Holweg (American Standard Companies Professor of Operations Management at Said Business School), Vincent Wang (Research Assistant at DELab, MSc student of the foundations of Quantum Computing), Andreas Tsamados (Research Assistant OII; AI for social good), Carl Ohman (Doctoral Candidate OII; Ethics and Politics of 'Digital Human Remains'), Huw Roberts (MSc student OII; Ethics & Chinese technologies), Nikita Aggarwal (Research Associate Digital Ethics Lab, Research Fellow and Doctoral Candidate Faculty of Law), Jessica Mörley (MSc student OII, tech adviser at the UK Department of Health and Social Care), Jakob Mokander (MSc student OII; Ethical guidelines for trustworthy AI).

We would like to point out that neither the OII nor any of the persons mentioned above are responsible for the text in this report.

1 Introduction: the creative machine



Scientists say we need dreams to process our impressions. Computers, on the other hand, have nothing to process and therefore have no need to dream. They have few needs at all; they are quite indifferent. They can switch on or off, but even that doesn't matter to them. Nor how they are programmed, whether they run on Linux or Windows, whether they have a graphics processor or not, or whether they are used for deep learning or a computer game. Despite this knowledge, Google claims to have created a computer that shows us its dreams. They call it Google DeepDream.¹ Perhaps computers could dream, you speculate. If they can dream, does that mean that there is fantasy in these machines and that computers can also be creative?

In Oxford, the Creative Algorithmic Intelligence course has recently started.² One wonders out loud what it means for human creativity when computers suddenly – quite recently – compose pieces of music and write lyrics of reasonable quality. Professor Marcus Du Sautoy of Oxford University explains in his book *The Creativity Code: How AI is learning to write, paint and think* how computers make art. At the same time IBM answers the question “What's next for AI?” with “The quest for AI creativity”.³ In a recent MIT Press book by Professor Arthur Miller, we read that computers are already more creative than people in some areas and that they are going to catch up with our creativity.⁴ And if you want to get started with AI creativity right away, RunwayML allows you to “discover the power of artificial intelligence in creative projects”.

It is all to do with the new discovery of so-called Generative Adversarial Networks, known as GANs. The name was invented by Ian Goodfellow when he was still responsible for AI at Google Brain. He now has a similar position at Apple.

We go back to 2014 during a get-together where students challenged Goodfellow with the question of whether computers are able to use their fantasy. When he returned home, he spent all night building his first GAN application in 24 hours. His creation turned out to be a hit; other researchers and developers came up with variations and the principle is now widely praised. The enthusiasm of the experts is an extra stimulus for us to explore this further. Yann LeCun, Chief AI Scientist at Facebook, calls it the most interesting discovery in machine learning of the last decade. Andrew Ng, the founder of Google Brain and now Chief Scientist at Baidu, calls it promising. Geoffrey Hinton, professor at the University of Toronto and often mentioned in the same breath as the other two “AI godfathers”, says GAN is a breakthrough. The number of GAN applications and GAN networks is growing day by day and we are seeing continuous improvements in the basic model and specializations for niche applications. We now count more than five hundred, brought together in the so-called GAN Zoo.⁵

To understand where that enthusiasm comes from, it's good to first see exactly what we're talking about. We can best illustrate this with an example, in this case of medicine inventor Insilico Medicine. This start-up from Johns Hopkins University, which moved their headquarters to Hongkong in 2019, has discovered a method to find and create a medicine for a disease in no time. At the basis we find this new application of artificial intelligence: the Generative Adversarial Networks.




From left to right: *The Artist in the Machine*, in which Professor Miller states that computer creativity will surpass human creativity; RunwayML, which offers AI tools to become more creative; and Marcus du Sautoy's book, which explains how it all works.

1.1 Innov-AI-tion: a medicine in no time

Developing a new drug is like looking for a needle in a haystack. It takes an average of \$2.6 billion and 10 years of hard work and experimentation. At the end of the ride, you also need to make sure the drug works without any unpleasant side effects or toxic effects. It's getting harder and harder to invent a new drug. Whereas the return on R&D in the pharmaceutical industry used to be 10 cents on every dollar invested, now it's 2 cents and the ten largest pharmaceutical companies invest around \$80 billion in R&D.⁶ The low hanging fruit in terms of drug development has already been harvested; as it becomes increasingly difficult to discover a new drug, any contribution to speed up this process is welcome.

If such a pin is found in the haystack in just 46 days, you know that something special is going on. We're talking about the announcement of the biotechnology company Insilico Medicine in the *Nature Biotechnology* publication on 2 September 2019. Thanks to generative AI, this company has succeeded in designing a molecule that is medicine for the prevention of fibrosis and a number of related conditions. The cost was about \$150,000. The announcement was not presented without a sense of drama: "It's the AlphaGo moment of the pharmaceutical industry". We'll come back to AlphaGo in a moment, first we'll go into how it went.


The key word here is generative AI, artificial intelligence that "generates" something. Generative AI is not new, we already know the phenomenon of generating probabilities. For example, in a picture of a cat or a dog, artificial intelligence gives an estimate, such as 80 percent probability that this is a dog and 95 percent that this is a cat, as in the picture below. But this new form of generative AI delivers something completely different, it generates a molecule.




Two forms of generative AI: AI that generates probabilities (above) and the molecule generated by AI (below).



The biotechnology company Insilico Medicine uses AI, including GANs, to find and create new medicines.



Deep learning enables rapid identification of potent DDR1 kinase inhibitors



Over a period of 46 days, the AI application was able to generate 30,000 different molecules, one of which was eventually tested on a mouse.

The discovery starts with a computer model (in silico) representing molecules that can act as medication. This process takes 21 days. The synthetic medication is then produced and analyzed in a test tube (in vitro) and in 46 days the medication is ready to be tested on a mouse (in vivo). Deepknowledge Ventures, which finances Insilico, had to wait patiently for two years until the theoretical framework was ready to test the technique (GENTRL). They certainly don't hide their enthusiasm for the breakthrough. For example, they speak of a discovery after 21 days instead of 46, because the molecule had actually already been discovered. The fact that the name of the company is identical to the computer phase of the medication development is not a coincidence. Insilico is a technology company that specializes in AI and wants to speed up the production of medicines, as Henry Ford once did with the production of cars. In this context, we are constantly looking for new words. Artificial intelligence has long been classed as "General Purpose Technologies" because of its general deployability. A trio of scientists from MIT, Harvard and Stanford, Iain Cockburn, Rebecca Henderson and Scott Stern, are researching the economic impact of artificial intelligence.⁷ They introduce a term based on their thesis that AI is able

to create better ideas than people. What GANs do is about a different way of inventing and creating things. GANs innovate innovation. In short: "invention in the method of invention".

"GANs innovate innovation. In short: invention in the method of invention."

Insilico now has partnerships with Astra Zeneca, Pfizer and numerous other major pharmaceutical companies. Insilico can be called an inventor of new ways in which inventions are made. The GENTRL model is available on GitHub⁸ for anyone who wants to start inventing in other ways.

The first creative machines with a patent

At the end of January 2020 it was announced that a medicine created by AI is being tested on humans for the first time. Not commissioned by Insilico this time but by the British start-up Exscientia and the Japanese pharmaceutical company Sumitomo Dainippon Pharma.⁹ In September 2019, a patent¹⁰ was granted in the US to a "creativity machine" running on GAN technology. For two of the inventions made by this machine, patents have also been applied for at the European Patent Office. The applications were rejected in January of this year. Professor Ryan Abbott of the University of Surrey, who is behind the patent application, comments that if the GAN had been a human being, nobody would have doubted attributing the inventions to that person.

1.2 GANs: behind the scenes

Insilico's creativity machine, with which they made the discovery, consists of NVIDIA graphics processors that are also used in computer games.


A whole bunch of cloud solutions, such as AWS, Alibaba Cloud and Google Cloud, now comes into play. Deep learning frameworks from TensorFlow and PyTorch are added, and petabytes of data from many different sources are added, such as from a gene dataset, information (texts) from scientific research, enumerations of chemical

compositions, a library of automated workflows for biology and chemistry, a toolset for identifying molecular generation and applications for creating proof-of-concepts. But the most important role is reserved for Generative Adversarial Networks.

Generative Adversarial Networks and synthetic data

The real breakthrough is due to a new form of generative AI, baptized Generative Adversarial Networks by Goodfellow, as mentioned previously. The principle behind this is to allow two self-learning neural networks to compete against each other. The first network, known as the

The Industry's Most Advanced End-to-End Drug Discovery AI



Applications: Insilico is developing and deploying a large number of proof of concept and commercial applications for target ID, chemistry and digital medicine.

Services: Insilico developed a set of tools for target identification and small molecule generation for internal use.

Main Focus on DL Pipeline and Libraries: Insilico developed a set of deep learning libraries and automated workflows for biology and chemistry.


Data: Insilico aggregated and processed petabytes of data including gene expression data, chemistry and text.

Deep Learning Frameworks: Insilico relies primarily on TensorFlow and Pytorch.

Computational Infrastructure: Insilico relies on the GPU clusters and some AWS. Unique storage solutions.

Physical Hardware Level: Insilico relies primarily on NVIDIA GPU for DL but is evaluating other technologies.

Insilico's creativity machine consists of an extensive collection of technologies.



Generative Adversarial Networks create synthetic data with white noise. The goal is that the synthetic data has the same properties and correlations as the training data and therefore forms a good synthetic representation of reality.


creator or artist, produces artificial output, for example images of a dog. The second network, known as the discriminator or critic, then assesses whether the produced results are real or fake by comparing them to a dataset of training examples. One tries to fool the other. And the other doesn't want to fall for that. There is constant feedback to the other as to why it has or has not worked. This happens without human intervention; in AI terms this is called "unsupervised learning". After many iterations, the quality of the output is so good that it is (almost) indistinguishable from the original training examples. The result is an image (or a text, a molecule, a sound, a video – many forms are possible) generated by the computer. The output is called synthetic data. The interesting thing is that because of this, the adage "data is the new oil" no longer applies. Less data, but more data modelling is the new paradigm. It also puts direct pressure on the profession of data scientist. Whereas just a few years ago this was proclaimed the sexiest profession of the twenty-first century,¹¹ with the introduction of synthetic data this can be critically questioned.¹²

Generative. The AI model generates new – synthetic – data based on a given set of training data. For example, if the training set consists of molecules, the AI model generates synthetic molecules. The goal is that the synthetic data has the same properties and correlations as the training data and therefore forms a good synthetic representation of reality.

Adversarial. This refers to the aspect of competition between the two algorithmic models: the creator and the critic.

Networks. Depending on the complexity of the GAN, it is a relatively simple neural network – a feed forward – or convolutional network that is mainly used to process images.

Synthetic data. Computer-created artificial data, the output of GANs. Synthetic data can generate ideas, accelerate learning models, bypass privacy issues, and overcome language barriers between different media forms. We will discuss this in more detail in Chapter 2.



1.3 Pharma's AlphaGo moment

Insilico spoke of "Pharma's AlphaGo moment" in its press release, referring to the historic 2016 victory of Google computer AlphaGo over humans in the board game Go. Of the five games that the computer played against world champion Lee Sedol, it won four. Experts thought that it would take another 10 years for the computer to succeed. And now the pharmaceutical industry, like the gaming world, has its AlphaGo moment. The question is whether we should or can identify a winner. If we go back to the game in question, there are two moments that stood out at the time. In the second game, AlphaGo made a strange move in the 37th turn, which the Go experts thought was a mistake. Later they were waxing lyrical about its beauty. This move had never been performed by a human before in a game. In the end AlphaGo won thanks to this ingenious inspiration. In the fourth game, at move 78, Sedol did something similar. The question now is whether he would have made this move if he hadn't had move 37 in front of him. In other words, man learns from the computer, just like the computer learns from man. It's not a competition, it's a play of circumstances. Three years after his loss, Sedol quit. The computer can't be beaten, he said. The fun is gone for him.

"I'm not at the top even if I become the number one."

Lee Sedol

Sedol's disappointment is understandable, but the conclusion that man's role has been played out is unsubtle and incorrect. Chess Grandmaster Garry Kasparov made a different choice after his defeat against Deep Blue. He teamed up with the computer and introduced a new game, Centaur Chess, in which people work together with computers to achieve better results. The ultimate impact of AI on the chess world is that the human level of play has increased across the board. Everyone now has access to a digital top coach. The technology has a democratizing and creativity-promoting effect.

"AlphaGo made an *original* and highly *effective* move. Within the definition of creativity, these are the two key concepts."

Thomas Edison used his own definition of creativity. According to him, human genius stems from 98 percent of

hard work and he added that creativity requires a lot of perspiration. In popular definitions of creativity, we still read the words of Thomas Edison: “Creativity is 98 percent perspiration and 2 percent inspiration.”

For most of the examples we describe in this report, the computer has to work hard to come to its original ideas and we can conclude that without man’s help and ingenuity nothing would come out of its hands.

1.4 The structure of the report and fifteen GAN conclusions

It is a fascinating thought that we can create something in no time. Medicine, but also all kinds of other products and services, can feel the GAN effect. We have, as it were, invented a kind of time machine, a machine that compresses time-to-create to make discoveries. Normally, in the world of computers, we talk about efficiency, effectiveness, brutal computing speed and cognitive ability. But what if the creative process itself can be improved with machines? That raises many questions. The three main questions are dealt with separately in the three chapters that follow.

How does the compression of time-to-create work?

In Chapter 2, we will present numerous GAN examples, divided into science, design, industry and architecture, art and media. Not all of them are equally groundbreaking, but they are often spectacular. It is interesting that people have been given this new digital possibility. After all, innovation is a matter of cleverly combining. The new possibilities of combining, in this case GAN creativity with human creativity, create new opportunities for innovation. Above all, we encourage you to think outside the box, because applications in the various sections interconnect gracefully. Time compression has many faces, as you will see.

Can computers dream and be creative?

In Chapter 3, we return to the existential question: can computers be creative at all? If we impose the restriction that computers have no will of their own, let alone the will to invent something, our answer is yes. It is a soulless variant of the human form of creativity. This sounds rather limited or a bit spooky, but eventually we will have to figure out over the next decade how these two forms of

creativity can reinforce each other, because human creativity also has its limitations. For the human-machine collaboration, or HUMAN-GAN interaction, we will see that in many places in the creativity process, the ball is in the court of mankind. Above all, the computer works, often at high speed, through numerous sources, something we as humans cannot match. If creativity is a combination of inspiration and perspiration, then much of the perspiration is on the side of the computer. The ultimate effect will be to increase the creative capacity in society – and possibly in your organization as well. So, what do you think?


Do CIOs still have dreams?

We conclude Chapter 4 with the question of whether CIOs still have dreams. It’s not the first time we hear that CIOs have a lot on their minds. This topic just adds to the list of things to do. Maybe even reading this bulky report is too much, because CIOs are under pressure. According to McKinsey,

this is the worst time imaginable for an average CIO to be at the helm. There are so many things that need to be dealt with and, at the same time, so much is happening in society that has to be taken into account. In this chapter we will present the above-average CIO and show five dreams (big and small) plus two nightmares in which we make the issue of creativity manageable.

Reading warning

Before we start, a few words of warning. Firstly, many of the cases of generative AI are not as grandiose as those of Insilico. Take, for example, the algorithm that converts music into dance, or the online retailer Zalando who uses a GAN to design dresses. We will also look at examples of computer paintings produced on an assembly line and how a GAN transforms old black-and-white video images into a high-resolution color film. Secondly, we have to warn you that GANs have only recently been included in Gartner’s well-known hype cycle. The term dangles somewhere at the bottom left, a lot lower than delivery drones, to name just one. Just above those drones we also find so-called synthetic data on the curve. Since this data is generated by GANs, we can consider it as part of the same development. The fact that it might not yet include an appealing case from your own sector does not necessarily mean that it isn’t relevant to your organization. We are simply in the early development phase of this new application. The question is whether you have the creativity to look through this matter, find your own application and




GANs and synthetic data are only at the beginning of Gartner's hype cycle.

shorten the speed to generate new ideas, products and processes.

Let's quickly dig deeper

In the list of conclusions that can be found throughout the report, you will immediately get a picture of the specific properties and operation of GANs. The examples will vary in maturity. It helps if you keep in mind that there are actually two sides to the coin: GAN input and GAN output. The output is about the ideas generated by the network, such as a molecule in the example of Insilico. It could also be a painting or a piece of music, or a design of a building. One specific application of GANs is that they can turn a piece of text into a drawing that expresses what is described in the text. In that case it is a kind of interpreter-translator from one medium (input) to another (output). The GAN input has a completely different application; the synthetic data generated by the GAN is added back into the neural network. This kind of synthetic data augmentation makes sense, because it makes the neural network function faster, better and cheaper. A second example of GAN input is very specific; synthetic human data files are created that are used as input for computer systems that need to be tested, but without any privacy-sensitive data being involved. GANs therefore bypass privacy issues.



Our fifteen GAN conclusions

page 14 Conclusion 1: Self-learning ability accelerates creativity

GANs are able to master laws such as quantum physics or traditional physics, bringing new scientific discoveries within reach.

page 14 Conclusion 2: GANs have a keen eye for what people have overlooked

GANs make hidden knowledge visible and point us to things that need to be investigated further. They help us to use people's creative abilities in the right way.

page 15 Conclusion 3: Data augmentation accelerates the learning process of GANs

The synthetic data produced by GANs can be used to accelerate neural network training. This saves money and time.

page 16 Conclusion 4: Brain interfaces have become more realistic

GANs work as an interpreter-translator from one medium to another. The decoder-brain interface described translates brain activity into images of relatively high quality.

page 17 Conclusion 5: GANs bypassing privacy issues

Generating a synthetic dataset that replaces real data allows systems to be tested without the risk of privacy issues and sensitive data ending up on the street.

page 18 Conclusion 6: Procedural modelling instead of fumbling around by yourself saves time

GANs show the future of (game) design. Time and money are saved by less doodling around and the designer can focus more on concepts and procedural modelling.

page 19 Conclusion 7: GANs significantly reduce design time

More and more GAN tools are coming onto the market that shorten design time and make suggestions for new designs. From initial sketch to unique product has turned into a quick process.

page 20 Conclusion 8: Provision of style advice by the computer to humans works well

GANs can advise people what to wear to look more fashionable, without having learned what is fashionable.

page 20 Conclusion 9: The first steps in architecture have been taken

GANs have taken the first steps for the architecture of physical buildings. This allows us to dream about the future possibilities of automated creativity, but we're not quite there yet.

page 21 Conclusion 10: GANs and 3D printing in industry offer opportunities

Generative design of industrial products has been around for a long time. Generative AI is an extension of this. The combination of GANs and 3D printing is particularly interesting.

page 22 Conclusion 11: GANs predict and detect errors

GANs play a role in the more efficient detection of defects, for example in solar panels, in predicting where problems will arise (predictive maintenance) and in preventing errors through simulations.

page 25 Conclusion 12: GANs make the invisible visible

GANs open the doors for synthetic media. Computer generated or modified images make the invisible visible, like a black hole, or make the world real by coloring black and white images or artificially increasing the quality of images.

page 26 Conclusion 13: GANs are interpreter-translators of media languages

GANs can function in the media as lingua franca so that text-to-image, photo-to-painting and music-to-choreography, for example, are automatically translated.

page 28 Conclusion 14: Fake is the new real

Deepfakes are perhaps the most famous and notorious application of GANs. Manipulative fake videos and fully automatic synthetic propaganda machines show that a boost in creativity is not always positive.

page 31 Conclusion 15: GAN art raises IP questions

GAN art undoubtedly raises new questions about originality, creativity and creative power that are also relevant to patent and IP issues.



2 Compression of time-to-create

From medicine to art, design and media, creativity is everywhere and has made our life on this planet what it is today. Imagine if we named all the results of human creativity. All in all, it doesn't even fit into a reference book. There are so many fields of application, there is so much diversity in discoveries and ideas, it provides a rich and colorful potpourri of ingenuity and inventiveness. From the work of composers such as Ludwig van Beethoven and Amadeus Mozart and painters such as Van Gogh and Gauguin, to brilliant computer experts such as Ada Lovelace and Alan Turing, pioneers of discovery such as Copernicus and Louis Pasteur and physicists such as Isaac Newton and Albert Einstein. In our modern age we also include Steve Jobs and Elon Musk. Musk, who as an advertising stunt shot one of his


Tesla's into space in a rocket around the sun, accompanied by David Bowie's "Space Oddity", is perhaps the most extravagant example of human creativity.

Although these are completely different people, what they have all done can be traced back to an idea and the imagination that such a thing is possible in real life. From there, it all comes down to the creativity of the execution. Something emerges, an artefact, something made by man. All these artefacts together make the world what it is today; we live in a human built world. From Van Gogh's sunflowers printed on a calendar or serving as an Instagram background, to Beethoven's Fifth as a rap performance by DJ Wilfred on your Spotify playlist. All these inventions and creations are now embedded in an extremely refined way in everyday life. What is ordered now will be delivered to your home before ten o'clock tomorrow, produced and assembled in different continents. What becomes clear at the expense of all those creative minds is that their creations have propelled the world and the economy in many new directions.

Art, culture, science, technology, education: everything can be traced back to the wealth of ideas and creativity of mankind that brought us where we are today. Suppose the GANs in this report ultimately make humanity 1 percent more creative, or 10 percent, or 20 percent? It is still too early to answer the question of exactly how much difference it will make. However, the potential to become more creative with the help of machines is definitely there. If only because GANs allow us to automate and accelerate a new part of the creative process.

2.1 Science: analyze, predict, synthesize

GANs are widely applicable, from art and design to science, industry and media. In this chapter we will specifically look at these domains. There is a strong connecting element: they shorten the time of the creative process, the time-to-create. This happens in different ways; by coming up with ideas and suggestions, as in the Insilico case, but also by other outcomes of the application of synthetic data, such as bypassing privacy issues or speeding up learning models through data augmentation. We will start with four exam-



ples of GAN applications from science: physics, life sciences, neuroscience and social science.


A) Physics: discovery of thermoelectric material

Researchers from the U.S. Department of Energy¹³ may have discovered new thermoelectric materials with the help of GANs. They believe there is a good chance that there are decent candidates among the suggestions, somewhat similar to the first phase of the Insilico case. This would be an important discovery, because if the efficiency with which heat can be converted into energy by a new material increases, this means that we are a little closer to solving the energy problem.

How did the researchers proceed and on what do they base their enthusiasm? On the basis of 3.3 million

scientific articles, GAN¹⁴ made suggestions for the development of new thermoelectric materials. The algorithm came up with solutions by making new connections between the research projects that were not previously found by humans. Without being told anything about materials science, the algorithm learned concepts such as the periodic table and the crystal structure of metals. GAN's self-learning and creative ability combined with the brutal computing power of computers provides a special kind of meta-analysis. With superhuman speed and volume.

There is a chance of success, because GAN has also identified materials from the many studies that have proven to work in the past. On this basis, it is believed that the new materials now suggested by GAN have the potential to work. Read how the researchers have achieved success with GAN in the article "With little training, machine-learning algorithms can uncover hidden scientific



Berkeley Lab researchers discovered that text mining of materials science summaries could yield new thermoelectric materials.

knowledge”.¹⁵ Anubhav Jai, the project’s leading researcher, tells us where the strength of the GAN approach lies. People are limited on time:

“The most interesting thing we’ve discovered is probably that you can use this algorithm to address gaps in materials research. Things that people should have researched but haven’t yet.”

This example is reminiscent of a study published in *Nature* with the title “AI-Copernicus discovers that the earth orbits the sun”.¹⁶ Not that we didn’t know this, but the scientists point out that this GAN has independently mastered the laws of quantum physics – just like the periodic table in the previous example. The fact that computers now discover for themselves what Copernicus already knew in 1543 is therefore very promising. The researchers say they are now ready for the bigger work, looking for discoveries that mankind hasn’t made yet.

Conclusion 1: Self-learning ability accelerates creativity


GANs are able to master laws such as quantum physics or traditional physics, bringing new scientific discoveries within reach.

Conclusion 2: GANs have a keen eye for what people have overlooked

GANs make hidden knowledge visible and point us to things that need to be investigated further. They help us to use people’s creative abilities in the right way.

B) Life sciences: learning acceleration through synthetic data

GAN techniques are also used in the life sciences sector. In this example they contribute to the solution of a general problem in the application of AI: the training of AI models. Identifying abnormalities in MRI scans of the brain is the work of X-ray scientists. Every picture has to be examined by such a specialist to detect abnormalities, such as blood clots or tumors. The more is diagnosed, the better the AI model will be able to make independent diagnoses. By producing synthetic MRI scans of the brain with GANs, the learning process is accelerated. Again, this works according to the GAN method: the discriminator (critic) and the generator (artist) work together to produce fictitious MRI scans (synthetic data)¹⁷ with correlations (what do healthy brains, blood clots, and tumors look like?) similar to real MRI scans. Suddenly a lot of scans can be produced that can be exchanged without problems and used to train other models.






A Deep Convolutional GAN (DCGAN) created synthetic MRI scans of the brain, based on 96 real MRIs. Indistinguishable from real scans by a lay person, even trained radiologists frequently misjudged which one was fake and which one was real. The researchers concluded that these synthetic datasets are a godsend for “data-hungry” techniques, such as supervised machine learning. According to a report by market analysis agency Cognilytica, the market for this third-party labelling of data will grow from \$150 million in 2018 to more than \$1 billion in 2023.¹⁸

C) Neurosciences: a brain decoder that sees what you see


Two researchers at the brain institute in Toulouse have found a way to use GAN technology to reconstruct images of people's faces directly from neural activity in the visual cortex. In this case, GANs are therefore able to act as translator-interpreter between brain and image; they quickly switch from one medium (brain waves) to another (image). It was extensively described in an article in *Nature* called "Reconstructing faces from fMRI patterns using

deep generative neural networks".¹⁹ By means of fMRI-technique, a 3D image of the brain is generated, after which the brain activity is converted into an image. In total, the dataset consisted of 8000 photos, which were rapidly displayed one by one. In the left column you see the faces that the participants got to see, the right two columns show reconstructed pictures from the measured brain signals. The middle column shows the results of the GAN method, the right column shows the results of another method (PCA: Principal Component Analysis). The GAN method clearly delivers superior quality.



Conclusion 3:
Data augmentation accelerates the learning process of GANs
 The synthetic data produced by GANs can be used to accelerate neural network training. This saves money and time.

By means of fMRI and VAE-GAN, a 3D image of the brain is generated, after which the brain activity is converted into an image.



Brain interfaces are not new, but the GAN approach delivers much higher quality. Look, for example, at the difference between a brain interface from about ten years ago in a study by Berkeley University.²⁰ The woman in the video (image below) is just a shadow, while we can see on the previous page that the new GAN images produce a recognizable result. What is striking is that with the faces of people with an Asian appearance, the resemblance is much less, see for example the fourth and sixth faces in the first column of the picture on the previous page. The reason for this may be that there were not enough Asian faces in the training set for GAN, which caused bias to get into the system.



For years, efforts have been made to translate brain activity into images. This example is from about ten years ago in a study by Berkeley University.

Conclusion 4: **Brain interfaces have become more realistic**


GANs work as an interpreter-translator from one medium to another. The decoder-brain interface described translates brain activity into images of relatively high quality.

D) Social science: GANs bypassing privacy issues

There are limitations when working with systems that contain personal data. Testing those systems is difficult because there are always conditions and risks associated with privacy legislation. Synthetic data is, by definition, anonymous because it does not relate to people who actually exist.²¹ For this reason, the Norwegian population register now works with fictitious residents to test their software systems. The British government is also showing interest. In a study carried out in Great Britain,²² four different GAN architectures were tested and one, in this case WCGAN (Wasserstein Conditional GAN), proved to be the best to mimic the distribution of population data. After 5000 iterations of the model, you will see the real data on the left and the GAN data on the right. Right away you can see that the different data sources do not differ much from each other, which is ideal to avoid problems with the General Data Protection Regulation (GDPR) when developing and testing systems that work with sensitive personal data.

Sogeti's data sciences team has launched a new service to meet this need. The Artificial Data Amplifier (ADA)²³ is used to amplify existing data. Although the focus of ADA is on creating synthetic data to solve privacy issues, ADA also offers more possibilities for the deployment of this GAN technology. Because GANs are so widely deployable, it is often best for marketing reasons to choose a specific focus, in this case an issue that many organizations have to deal with.


Conclusion 5:
GANs bypassing privacy issues
 Generating a synthetic dataset that replaces real data allows systems to be tested without the risk of privacy issues and sensitive data ending up on the street.



In a study carried out in Great Britain, the WCGAN (Wasserstein Conditional GAN) was found to be the best to mimic the distribution of population data. After 5000 iterations of the model, you can see the real data on the left and the synthetic GAN data on the right.

2.2 Design, fashion, and style

We now make the step from science to the world of (digital) design and fashion. We start with a good example of how time can be gained in the design process of games by using GANs and discover that the GAN market for design is becoming quite professional. More and more tools are coming to the market. For a quick overview of the possibilities, we refer you to the Algorithms.design website of a Russian designer who keeps track of all developments in the field of design and GANs. The tools are neatly classified into five categories: creating user interfaces, content creation, personalizing the user experience, graphic design and industrial design & fashion.



Conclusion 6:
Procedural modeling instead of fumbling around by yourself saves time
GANs show the future of (game) design.
Time and money are saved by less doodling around and the designer can focus more on concepts and procedural modelling.


A) Design of games: an end to fumbling

It takes a lot of time and effort to make a digital design for an object in a game. Andrew Price, a young, famous game developer, calculated how much it costs and sees the future of GANs in game design. He made a simple calculation. To create a building, from design to texture to lighting, an average designer takes about 66 hours. At an average hourly rate, a ready-made house costs about \$3,900. A fully equipped street will soon cost just under \$200,000, which contributes significantly to the enormous cost of developing a video game.



Andrew Price gave all items a price tag in this virtual street and came to the conclusion that it would cost about \$200,000 to build the entire street.

Price explains that when he builds a 3D home interior, he spends most of his time (50 to 70 percent) fine-tuning the design to achieve perfection. For example, to determine whether or not other lighting in the room is better, whether or not wine bottles should be on a table, whether or not there should be a refrigerator in the room, to ultimately get the best picture. This fumbling around takes a lot of time, because every change has to be examined and built in – and possibly thrown away again. With GANs, these things can simply be generated – and used or not – by a computer. The computer helps him in his creative process and that can save a lot of time. As another time saver, Price mentions procedural modeling as a replacement for a static method of working, in which everything has to be built up from scratch every time. A number of parameters of a building is specified: the number of windows, doors and floors and the texture of the building, and then the AI algorithm generates a series of examples from which you just have to choose. This modelling can also be done in other ways, as we see in the following example.



The twelve images of birds were generated by the AI. The same goes for the design of the bags and dresses.

B) Good taste: clothing design and advice

We've already told you that more and more GAN tools will be introduced to the market for design. Adobe showed the "Project Image Tango" at its MAX conference in November 2019. Based on GAN technology, the software generates new ideas for the design of bags, dresses and, for example, birds. It was continually emphasized that these designs did not yet exist and that they are nowhere to be found on the internet. A simple sketch of a bird, for example, is enough to generate a stunningly natural (synthetic) image of a non-existent bird. The desired style and colors are imported from examples, after which the software generates unique copies time after time.²⁴

A concept similar to the Adobe software is applied to a project by Zalando Research. In their paper *Disentangling Multiple Conditional Inputs in GAN's*,²⁵ the researchers change the texture and color of the garment, but they are also working on an interface to design a garment based on a simple drawing. This is already very useful for designers, but possibly in the future also for customers. The idea is that labor-intensive steps are skipped on the way from drawing to design, because they are automated with the AI.



Above, the created dresses by a GAN, where the shape and the color combination have been determined beforehand. Below, the black picture is a hand sketched garment. The GAN automatically turns it into a plausible clothing item.

Conclusion 7:
GANs significantly reduce design time
 More and more GAN tools are coming to the market that shorten design time and make suggestions for new designs. From initial sketch to unique product has turned into a quick process.

GANs that give style advice and make people streetwise

A step beyond designing a garment is giving advice on what to wear. Is it better to tuck in your shirt inside your pants or wear a sweater instead of a blouse? A number of academics have looked into this issue and eventually managed to capture "good taste" in a model and give clothing advice. The system is called Fashion++ and is based on a set of 10,000 photos from the collection of Chictopia.com, "a content archive for the history of the fashion community", which contains pictures of influencers in the latest outfits. The AI system is able to recognize details such as

the texture, shape and color and the pattern of the clothing. The system then removes the clothing and deliberately adds non-fashionable clothing, which is necessary to eventually create a so-called “fashionability classifier”. If someone adds a photo with their clothing, the computer comes up with an alternative that is more stylish than that depicted in the photo. The special thing about Fashion++ is that no one has said what is fashionable or not, the system has taught itself. Why something is fashionable, or more fashionable than something else, is often difficult to put into words.

Conclusion 8:
Provision of style advice from the computer to humans works well
GANs can advise people what to wear to look more fashionable, without having learned what is fashionable.

2.3 Design in architecture and industry

In this section we will present three cases that sketch a picture of possible applications in architecture and industry. We will start with architecture. As we saw in the design of computer games, there are gains to be made in the design process of industrial products and in the design of real buildings. The first steps have been taken with the development of a so-called archi.GAN. Next, we will look at the design of products. Generative design is a term with which the industry has been familiar for some time. Generative AI fits perfectly with this development. Finally, we will look at the detection of defects for maintenance and predictive maintenance.

A) ArchiGAN shows the future of architecture

The way living spaces are designed has long since shifted from paper to computer screens and from pencil to CAD software. NVIDIA now has a step-by-step plan in place²⁶ to let GANs do the work. It's called *ArchiGAN: a Generative Stack for Apartment Building Design*. It's all still at an early stage. For example, NVIDIA points out that they cannot

guarantee the calculations of the weight walls can carry when building multiple floors. Stanislas Chaillou, working at NVIDIA, shows in his graduation project²⁷ at Harvard how he uses GANs to generate and analyze floor plans for a building. The texture and thickness of walls, orientation, circulation and connectivity are, among other things, included in the design. Chaillou prefers a so-called “Grey Boxing” approach, in which people can control the design at several points. This is the opposite of a black box, where people provide input in advance and then have to deal with the output afterwards. His GAN models can be flexible, for example, when people have certain thoughts about furniture and other dimensions.



The user interacts with the design via a handy UI. With a simple click of the mouse, the room, doors, windows, etc. are changed and the GAN then creates a possible layout.²⁸

Conclusion 9:
The first steps in architecture have been taken
GANs have taken the first steps for the architecture of physical buildings. This allows us to dream about the future possibilities of automated creativity, but we're not quite there yet.



Autodesk's generative design for NASA's exploration of Jupiter and Saturn that is lighter and stronger and can only be fabricated by 3D printing.



Above, "generative designs" for the famous Volkswagen T2 van; below, designs for a General Motors seat brace to which the safety belt is attached.

B) Industrial design coming up with organic solutions

Because NASA has set its sights on the exploration of the moons of Jupiter and Saturn, the space organization needs materials that are even lighter but at least as strong. To this end, NASA is collaborating with Autodesk and their creative generative AI, among others.²⁹ The algorithm presents a surprising but more optimal solution. It is interesting to note that the structures generated cannot be realized through conventional manufacturing but require 3D printing technology.


Car manufacturers are also increasingly experimenting with the creative design of generative design processes. Often the result looks almost "organic", because the software makes much more accurate calculations about the forces that the components have to withstand. Left, you can see the designs for the well-known Volkswagen T2 van,³⁰ and below you can see the design for a General Motors seat brace to which the safety belt is attached. The design is 40 percent lighter and 20 percent stronger than the original.³¹

Conclusion 10: GANs and 3D printing in the industry offer opportunities

Generative design of industrial products has been around for a long time. Generative AI is an extension of this. The combination of GANs and 3D printing is particularly interesting.

C) Detection of mistakes, predictive maintenance in simulation

There are two methods to investigate the energy efficiency of solar panels. One is by looking at energy production, but unfortunately it does not detect all the defects. The other method is called electroluminescence,³² where “an image” is made of the energy production. It is a commonly used method; the disadvantage is that it has to be pitch-dark to make those images and labelling the mistakes requires a lot of specialized human knowledge. At an international conference on the future of energy supply, a GAN model was shown that improved the quality of electroluminescence by 14 percent. The same synthetic data augmentation technique was used here as with the MRI scans in the example on page 14.



The top row shows original images of deviations in solar cells. The lower row shows the images generated by GAN. In the end, with the addition of this synthetic data, the model was 14 percent better at recognizing abnormalities in solar cells.

Predictive maintenance

Predicting when defects will occur has long been a trend, especially in the automotive industry. Fujitsu developed an AI system that detects abnormalities at an early stage based on in-flight sensor data from Airbus helicopters. The applied algorithm (MAD-GAN: Multivariate Anomaly Detection with Generative Adversarial Networks) helps flight engineers analyze a huge amount of test flight data and quickly focus on possible problems, which saves an enormous amount of time.³³ In addition, the GANs that develop the best and most efficient synthetic data to optimize predictive maintenance are being researched at various locations.^{34,35}

Preventing mistakes by simulation

It appeals to the imagination: a fleet of self-driving cars no longer has to travel millions of miles for testing purposes; it can simply stay in the garage. This is because data simulations using GAN technology make it possible to simulate real-life situations.³⁶ They can easily turn winter into summer, or dry weather into rainy weather. This is a useful application for testing self-driving cars in simulated worlds.



GANs can be used to change weather conditions in a photo or video, resulting in more training data for self-driving cars.

Conclusion 11:
GANs predict and detect mistakes
GANs play a role in the more efficient detection of defects, for example in solar panels, in predicting where problems will arise (predictive maintenance) and in preventing mistakes through simulations.


2.4 Synthetic media and art: fake is the new real

GANs are now gratefully used for conveying messages, be it through media or art – whether or not in an experimental setting. Here we outline GAN's qualities as an interpreter-translator from one medium to another. We mentioned it before: GANs can serve as a kind of lingua franca and we show even more special media qualities, such as the upscaling of old films up to and including making visible what is invisible. A more important observation is that reducing time-to-create is not always a meaningful goal. In the case of art, one can say that the time an artist has put into the work is part of the value attributed to it. A second observation is that deepfakes – which can be used to manipulate people – have a disruptive character and consequences that are not yet fully understood.

A) Making the invisible visible; erasing it and upscaling it

The media moment of 2019; the photo of a black hole. For the first time, people managed to take a picture of this phenomenon. When seeing the photo, most people don't realize how the image was created. The photo is a composition synthesized from various images taken by the Event Horizon Telescope (EHT), a large composite telescope consisting of a worldwide network of radio telescopes. It is a virtual radio telescope as large as the earth itself. By means of an atomic clock, all telescopic sensors were able to shoot an "image" of the black hole in galaxy M87 at exactly the same moment. The sensors used only detect radio waves, which are not detectable by the human eye. The electromagnetic information then had to be converted into digital information before it could be used to generate a visual image using artificial intelligence (including GANs). The entire process took 2 years of computer time and 347 scientists were involved worldwide.

The photo is a milestone, and not just for science. It also marks the beginning of a new type of media, called "synthetic media". These are media manipulated or created by artificial intelligence, such as images, sound, text, video and virtual experiences. These media are often so realistic that it is impossible for a lay person to tell whether something is real or fake. The image of the black hole helps to better understand the nature of synthetic media. This illustrates that synthesized information is not necessarily unreliable, deceptive or false. The image is reliable to such



For the first time, people managed to take a picture of a black hole. The photo is an example of a new type of media, so-called "synthetic media": media manipulated or created by artificial intelligence, such as images, sound, text, video and virtual experiences.

an extent that it leads to an increase in our human knowledge and science can now build on it. With the help of this image, existing theories can be tested, and new ones can be created.

It can also be the other way around: making what is visible invisible. In the two images depicted right, the toolkit of technology company NVIDIA shows us how images are erased (left) and the original image is automatically filled with lifelike details. You can no longer see this is Photoshopped. We are losing our grip on reality in this case.

An intermediate form is called the upscaling of material. Anyone who has seen the latest *Lion King* movie knows how close these synthetic beasts are to reality. Old films and photos usually don't look very real because of the low resolution and the small number of images per second. A recent but already iconic example of upscale art is the *Arrival of a train at La Ciotat station*. This has often been mentioned as the birth of film. In the 50-second production of the Lumière brothers we see trains arriving and people boarding in grainy and jerky images. It's hard to imagine that in 1896 people fled the auditorium for fear of the approaching train. At least, that's how the myth

goes. On YouTube (see the link in the notes³⁷) you can see what developer and YouTuber Denis Shiryayev made of it with the help of Gigapixel AI. The neural network DAIN then added new images to the original frames of the video by imagining what the missing images should look like. The end result is a 4K (a resolution of 3,840 by 2,160 pixels) video clip with a frame rate of 60 frames per




Elements are erased (left) and the original image is automatically filled with lifelike details. In this case, the machine takes over the necessary skills for photo and video editing.



The 125-year-old film *Arrival of a train at La Ciotat station* has been updated via GANs to 4K resolution and 60 fps and colorized with DeOldify software.

second. But there was still room for improvement. Shortly after uploading the original video, Shiryayev shared a colorized version made with DeOldify software, which basically looks like a short film taken in the 1970s.³⁸ An excellent re-interpretation and proof of the power of synthetic media to manipulate video images.

Besides old films, Albert Einstein, Marilyn Monroe and the Mona Lisa are also brought to life with the help of GANs. Researchers from Russia at Samsung's AI Lab used GAN technology to animate still images so that the images could smile or move their lips as if they were talking.³⁹ This is a new form of photo and video editing that was unthinkable until recently but has now become relatively simple.



Conclusion 12:

GANs make the invisible visible

GANs open the doors for synthetic media. Computer generated or modified images make the invisible visible, like a black hole, or make the world real by coloring black and white images or artificially increasing the quality of images.

B) Interpreter-translator: from one medium to another

Several researchers have looked into the question of whether you can turn text into an image or a video production in a straight line. This way you can skip an entire process of making animations in one fell swoop. The examples given symbolize the lingua franca function of AI, from text to image, from image to text or from text to sound. If a musician can conjure up a more up-tempo background of a saxophone just by asking, composing music becomes a lot easier. You can then have a GAN draw choreography that goes with the video.



The translation from music to dance is made possible thanks to the Music-to-Movement-GAN. The MM-GAN translates music directly into a matching choreography. The generated choreography is attuned to the beat, the musical style and the diversity of the steps.⁴⁰

Researchers from Microsoft and Tencent, among others, developed a way to let GAN tell a story in pictures by just entering text. They appropriately called it "StoryGAN".⁴¹ Scriptwriters for animated films will soon be able to generate a (concept) film directly from their written text.

The same trick was shown with videos of the Flintstones, which can be read in the scientific paper with the appealing title *Imagine This! Scripts to Compositions to Videos*.⁴² We see Fred and Wilma Flintstone in action as soon as the sentences in the script are typed in. "Fred opens the fridge" or "Wilma walks out of the house" is immediately visualized and converted into a moving animation. Or, where Fred is talking to Wilma in the dining room while she is reading a newspaper and ignores Fred, as seen on the right.

Changing from one medium to another also happens from image to image. Using the CycleGAN,⁴³ cats are made into dogs, apples into pears or, as in the image below, horses into zebras (top left; in still image and in moving video). The CycleGAN also transforms a painting into a photo (top right) or a photo into a painting in the style of your favorite artist (bottom).



Film writers may soon be able to generate a (concept) film directly from their written text. The technology is not yet perfect, but seems promising.

C) Creating new realities

For those who still wonder how truthful the synthetic media generated by GANs are, it is interesting to look at the astonishing improvement in recent years of GANs generating images of human faces. A well-known example is the website Thispersondoesnotexist.com, where visitors are confronted with the authenticity of fake faces. There are also stock photos generated by GAN, for example the 100,000 on <https://generated.photos/>, where you can find images of synthetic people for your (marketing) purposes. Some are clearly fake, and others are indistinguishable from real. (The creator of the website exempts all creations from copyright and distribution rights.)

Conclusion 13:
GANs are interpreter-translators of media languages
GANs can function in the media as lingua franca so that, for example, text-to-image, photo-to-painting and music-to-choreography are automatically translated.



Input

Monet

Van Gogh

Cezanne

Ukiyo-e



With the so-called CycleGAN, horses are made into zebras (top left; in still image and in video), a painting is transformed into a photo (top right) or a photo into a painting in the style of your favorite artist (bottom).




The enormous leap in quality that synthetic data has made since Ian Goodfellow came up with his GAN concept in 2014. From left to right computer-generated faces of people who do not exist. We see the quality improve drastically each year.

A powerful example of the creation of new realities is an audio fragment of the speech John F. Kennedy had ready, but which he never gave because of his assassination. Listen to *JFK Unsilenced* and hear how lifelike he delivers his speech post-mortem.⁴⁴ To warn of the danger of these possibilities, the MIT Center for Advanced Virtuality has made a video of Richard Nixon. We listen to the voice of a voice actor and watch – in all likelihood – synthetic images created by GANs. Nixon announces in the video that the people who have landed on the moon will stay on the moon (the video is called *In Event of Moon Disaster*). It's a strange message that perpetuates the myth that Nixon had actually prepared such a text in case something went wrong with the mission to the moon.

You could say that it's all just a game. Let's enjoy the possibilities and have fun with the remakes of old movies, explore galaxies or take a quick snapshot of some text. However, if you can no longer believe and trust your eyes and ears, alarm bells will go off. This should certainly also happen with Chief Information Officers who trust dates and facts blindly. If only to prevent super-sophisticated phishing methods based on these media tricks from leaving the gate wide open to hackers. Of concern are phenomena ranging from advanced phishing methods to misinformation and the so-called "liar's dividend": the defendant can deny everything by claiming that the authentic video is a deepfake.⁴⁵ In the meantime, Facebook has announced that deepfake videos are no longer allowed on its platform.⁴⁶ In December 2019, it



GANs are often also the technology behind deepfakes. For example, we hear John F. Kennedy reading a speech he wrote but never delivered, and we see Richard Nixon announce that the people who landed on the moon will stay on the moon.



removed another network of more than 900 fake accounts from its platforms that allegedly used computer-generated faces as profile photos and were brought to life to promote right-wing ideologies.

Getting started with fake news


Using the recent open-sourced synthetic text generator GPT-2, you can see for yourself which way this type of generative AI is going. The text generator is called: <https://talktotransformer.com>. The system is trained to predict the next word when we're typing. Based on that knowledge it produces coherent (nonsense) texts of a high level. It is open sourced to allow scientists to work on an anti-serum, something that can easily detect fake texts of this kind. The creator of GPT-2 is OpenAI, a research institute in San Francisco⁴⁷, which was founded in 2015 by, among others, Elon Musk, with the aim of promoting and developing people-friendly AI that benefits mankind. At the beginning of this year they stated that the model could not be made public because it was too risky.

I have created a monster

One of the interesting recent – and potentially worrying – applications building on the GPT-2 model comes from Dr. James Howard, a cardiologist in training at Imperial College in London. He has trained GPT-2 using more than 30 million quotes from the biomedical literature.⁴⁸ His creation is now able to independently generate fake summaries and conclusions of fake research – including fictitious clinical test data and false references to government registration – by simply entering a random fictitious title for a medical research paper. The experiment takes the fake news debate to a new level. “I have created a monster”, Dr. Howard tweeted when he saw the results.

Conclusion 14: Fake is the new real

Deepfakes are perhaps the most famous and notorious application of GANs. Manipulative fake videos and fully automatic synthetic propaganda machines show that a boost in creativity is not always positive.



Super-fast GAN art that disappears after an hour. The neural network continually generates art.

D) Artists embrace GAN possibilities

Artists are creative minds open to their surroundings and developments. It's not that surprising therefore that extensive experimenting is going on in art applications using GANs. It's interesting to see how some artists look at this new mix of technology and their own creative ability. The conclusion is quickly drawn: GANs are a valuable addition to an artist's arsenal for expressing themselves. This is a valuable lesson that we can take to heart for the other domains in which GANs will play a role. As with media, artists sometimes want to approach reality in detail. The so-called hyper-reality in painting is totally hip and happening again these days. But the theme, compression of time-to-create clashes with the concept of art. When assessing a work of art, for example a painting, we often consider the skills needed in the process and the difficulties the artist had to overcome in order to come to their expression. In short, we appreciate the time and dedication of the artist and take that into account in our considerations. It often becomes even more special when the artist has mixed their own colors and oils to make the paint instead of using ready-made paint tubes from the shop. Or we appreciate the suffering, the intention or the


progressive vision of the artist. According to Margaret Boden,⁴⁹ research professor of cognitive science at the computer science faculty of the University of Sussex, this is one of the problems people have with digital art. New skills such as innovative programming, creative pre-processing and tweaking a neural network are often invisible, difficult to assess and less palpable.

In GAN art, it comes down to the fact that the neural network trained by man can continually produce new art. The digital art gallery 9gans.com, for example, does that. It delivers new AI paintings every hour that can be printed on canvas. However, you need to react quickly if you see something nice on the website, because the generated artworks are not saved and are automatically deleted after an hour. You can call these disposable paintings conceptual art because of the disposable effect, similar to the work of the famous artist Banksy, who automatically lowered a painting into a shredder immediately after it was hammered at an art auction.

You can leave the choice of exhibited art for a virtual gallery to the computer, but most of the time the artist is in charge of the curation. This takes time, just like training

the models and the style choices you make as an artist. The paintings shown on the previous page were made by someone who calls themselves a “hacker and scraper” and is not an artist. But the works look very arty and would not be out of place with the average art loan. There are also plenty of serious artists working on GAN creations. For example, anyone who visited Hito Steyerl’s video installation at the Venice Biennale will immediately see that GAN art has more to it than a simple push of a button. Beautiful synthetic video images (GANs) of the city and nature come to you like a crazy dream. There are screens everywhere, there are vistas and you can walk across platforms straight through the artwork. Mario Klingemann, another artist, also puts a great deal of work into it. In one of his works he explains how he used different neural networks to make a video clip. The faces are made with a pix2pix GAN,

a GAN specifically designed to translate images into images. Klingemann has biometric facial markers directly connected to music and to special effects and has a self-designed style transition technique he calls GAN2GAN, in which he has connected several GANs, trained for enchanting facial transitions. Klingemann has a radical vision of creativity. “People are not original,” he says, “they elaborate on what they have learned before and on what others have built before. Machines, on the other hand, can create something from nothing”. The question of who the spiritual father of AI art is ultimately, is becoming increasingly difficult to answer. It is obviously an interplay, but the signature of the algorithm under the portrait of Edmond de Bellamy makes it clear that we are dealing with new ownership issues. Who has intellectual property (IP)? In our introduction we pointed out that the European Patent Office is not yet ready for IP in the hands of a machine.



The portrait of Edmond de Bellamy was sold for \$432,500 in 2018 and raised new questions. A GAN created the painting, but the idea to feed the GAN with a dataset of 15,000 portraits painted from the 14th to the 20th century came from the French art collective Obvious.⁵⁰ The GAN algorithm was written by Robbie Barrat and taken from GitHub. Who exactly is the creator of the painting here? Do we speak of machine creativity in this case?

Professor Ahmed Elgammal explores the boundaries of what exactly those machines can create for art. Elgammal is Professor of Computer Science at Rutgers University and director of The Art & Artificial Intelligence Lab.⁵¹ Whereas GANs generate countless variations on the original training data, he wants to deviate from this training data. With his AICAN (Artificial Intelligence Creative Adversarial Network) he tries to deviate just enough from the training data to create out of the box works that are still attractive. His AI artworks were exhibited in a gallery in Chelsea, New York in February 2019 and he is regularly in the media. He does not see computer art as a substitute for artists, but as a creative instrument or partner. Call it a creative cobot. "Artists have always been using new technology in making their art, and every time some technology came, it changes the way art is going. From the invention of oil paints, to the invention of print making, to the invention of photography, and I totally believe using AI in making art will be integrated into the mainstream of making art in the future."

Conclusion 15:
GAN art raises IP questions
GAN art undoubtedly raises new questions about originality, creativity and creative power that are also relevant to patent and IP issues.



Art generated by AICAN. When people were asked about the creator of these works of art, 75 to 85 percent of those questioned thought that they were created by a human artist. They added words like "intentional", "inspiring" and "communicative" to the art.


2.5 In conclusion: the short nose of innovation

GANs are able to do something that a creative person normally needs more time to do – in art and design, science and the media – in a much shorter period of time. Whether it's coloring a black and white film, looking for a new medicine, designing a house or designing a new product; when using GANs we compress time and come up with a result sooner. This drastically shortens the time-to-market from idea to launch and reduces the required investment. All GAN applications therefore have one thing in common: they compress the time-to-create. Bill Buxton, pioneer in human-computer interaction and principal researcher at Microsoft, has focused on this time problem and elaborated it in a theory he calls "The Long Nose of Innovation". In short, new innovations take about twenty years to mature. If we focus on this and compare it to the examples of the many GAN applications, the time span of twenty years will soon be shortened, depending on the inventions

we're talking about and the GAN applications, of course. Buxton's long nose will therefore be a lot shorter. The same conclusion is drawn by the science trio we mentioned in the introduction (Iain Cockburn, Rebecca Henderson and Scott Stern). They came up with "the invention of the methods of invention" and even promised exponential economic growth based on the new AI possibilities.

Each domain has a different meaning. We have seen examples where GANs help to invent (from medicines to clothes) and improve existing products (design lighter materials, better product properties) and help to give other technologies traction, such as increasing the efficiency of solar cells. The synthetic data as output of GANs can serve as a learning or idea accelerator for other models and as a catalyst for creative processes. Sometimes they are even the creative end product itself. If the creative end product is a media appearance, such as voice, image, text or video, it has two faces. It can be an enhancement of the original, such as old films being

The Long Nose of Innovation




GANs compress creativity by directly interfering with the method of invention. In this sense they shorten Bill Buxton's "The Long Nose of Innovation".

scaled up, or even a first imagination of the original, such as the image of a black hole that is not perceptible with our eyes. It can also be a creation of an original, of something that did not exist before, such as a face of a non-existent human being or a voice of someone who actually does exist. This last category of synthetic media is not only entertaining but also disturbing. If the time-to-market for confusing people and making them believe something other than the truth is shortened by GANs, we have a problem on our hands of which the consequences we cannot fully comprehend. The short nose of innovation emerges as the nose of Pinocchio.

Our final note is also an introductory look at the chapter that follows. The arrival of GANs offers a new opportunity to come up with ideas; it is a new tool from the rich toolset available to man to innovate. This touches the essence of innovation, because innovation is combining. This special tool simply increases the number of combinations, because it can create an explosion of new thoughts, expressions and inventions. We're talking about what Hal Varian, Google's Chief Economist, calls "combinatorial innovation": the ability to innovate by making combinations of already existing components. The graph below shows at a glance how linear, exponential and combinatorial functions relate to each other. In just a few steps we're exposed to staggering numbers and outcomes that far exceed human imagination.

GANs and the artificial data they generate bring us into a "mirror world" of the real world. Numerous scenarios, simulations and virtual and alternative realities and even hyperreality can be generated at the touch of a button. People only have to make a selection, which can then be tried out and rolled out in the real world. The present is no longer tied to what is there; it is determined by the extent to which the imagination of these neural networks succeeds in rearranging the material characteristics of the present. We see the emergence of a new form of cooperation between man and machine, which is the subject of the next chapter.



The slogan that technological growth is exponential rather than linear is now surpassed by the statement that a combinatorial explosion awaits us. Not only will developments be faster, there will also be more ideas that can be combined again. The functions shown in this graph and table symbolize the difference in acceleration.



3 HUMAN-GAN-creativity


The look behind the scenes we provided at the beginning of this report brings perhaps a sense of disillusionment. You will find no aspirations, no intentions, no passion, no emotion, no goals and no awareness in technology. GANs are advanced analytical calculators that search for patterns and create new synthetic output based on these correlations. But if we look at some of the examples of output mentioned in Chapter 2, they evoke wonder and amazement. In several experiments in which people unacquainted with music or art experienced computer-generated instances, they became emotional and added all kinds of emotions and intentions to the synthetic creation. How does this view of technology and emotions evoked by the output relate to the question that is constantly lurking in the background in this report?

3.1 Are computers actually creative?

It will not surprise you that we answer yes to this question. After all, despite divergent opinions, the frequently quoted meta-study *The Standard Definition of Creativity* shows that since 1953 a simple two-part definition has prevailed in academic publications. This is also called the Standard Definition of Creativity and it simply reads: “Creativity requires originality and effectiveness.”⁵² And yes, computers can certainly come up with an original idea and be effective, like the incredible 37th move that the AlphaGo algorithm put on the Go board.

There is, however, an important “but” that is often overlooked. Computer creativity is in fact uninhibited creativity, unimaginative, submissive and without context. Computers can be creative, but that does not give free rein to all sorts of anthropomorphist projections about the world of experience of these machines. Again and again, human beings set the goal and the preconditions and determine the question that needs to be answered. The GAN then analyzes, iterates, varies, generates and produces it all. It would therefore be nice to see computer creativity as just another, complementary form of creativity, instead of an enemy, or a creative god. Just as we distinguish cognitive and social intelligence or human and artificial intelligence, we also distinguish human creativity and computer creativity, or as Oxford University calls it: Creative Artificial Intelligence.⁵³ We define computer creativity as follows:

Computer creativity is the soulless creative power (of synthetic data) that requires originality and effectiveness.




The Lovelace 2.0 test, the lesser-known sister of the Turing test, offers a different perspective on computer creativity.⁵⁴ In this test, a person sits behind the computer and knows they are talking to AI. They first ask the program for a specific creative form of expression (e.g. a painting, story or poem) and then set the criteria of what it should be about (e.g. 'a man who saves a cat from the tree'). If the human being is satisfied with the result, they will ask a next, more difficult question. This continues until the AI fails or the human being is completely satisfied. The goal is to come to a score of (creative) intelligence instead of a right or wrong judgement.

3.2 HUMAN-GAN: a closer look

If we visualize the workings of the creative machine once more, as in the image below, we immediately notice the central position of the human being. Human beings determine the goal and give direction. This starts with pre-curation, where human beings determine the input. Is it a million scientific papers, all compositions by Bach or Beethoven or all flora and fauna on this planet? GAN generates and comes up with a synthetic representation, a sound, a model of a mineral or a photograph of a person who does not exist. This result of the algorithmic GAN game between the creator and the critic (discriminator) is presented to the human. Essentially, we have a fast analyst and imaginer – GAN – who helps the human in their creation process. There is always need for a representation of an idea, if only to keep our own flow of thoughts going.

The human being is the second discriminator who has to make decisive choices in post-curation. In the example of AI art this is still the most ostentatious. Which painting of the thousands produced by GAN is worth showing to the world? Or what else are we going to do with that outcome, which twist do human beings ultimately give to the computer idea? Is there something that still needs to be changed, added or finished? Also, in the case of the digital art gallery we showed, where GAN automatically uploads its work, it's the human being who determines the rules of the game.



GAN is an advanced inanimate calculator, while human beings have intentions, aspirations and goals. GAN learns the correlations present in the human-selected input and creates new synthetic variants. This creates a range of simulations and creations that fit within the presented parameters (the input).

3.3 A new collaboration, to what level?

To our own surprise, we conclude that behind our original question of whether computers are creative is a much more relevant question: what is the impact of these generative systems on our human creativity? Because, apart from the fact that creativity is important for our well-being⁵⁵ and beyond the question of definition, we see the emergence of new creations in numerous examples, creations that would never have seen the light of day if human beings had stood alone – without analytical computer creativity – or if the computer had stood alone – without the goals and intentions of human beings. In all these new HUMAN-GAN collaborations, GAN learns the correlations present in the input and opens up many possibilities with automated speed with the same parameters as the input. If a human being supplies 10,000 faces for example, GAN learns which correlations turn a face into a face and can then make infinite variations. Where a person loses many man-hours to make one unique face, GAN is able to create many synthetic faces in a short period of time after training. The result is a form of automated “narrow creativity” on demand. Wider creativity on demand arises when there is more data diversity in the game, as we saw in the example of Insilico.

In an interview between Lex Fridman and GAN inventor Ian Goodfellow (which can be listened to in a podcast⁵⁶), Goodfellow fantasizes about the possibilities. He is challenged whether GANs are the harbinger of Artificial General Intelligence; the point that computers are no longer inferior to people. Goodfellow does not address the question of whether it is ever possible or whether we


are going to reach this point. He does lay out what it takes: much more data diversity to *truly* understand the world. A computer doesn't know what a car is if you show it 10,000 pictures of cars. For human beings, the car is a means of transport, a danger on the road sometimes, a romantic look back in the past when we see an old VW beetle driving, a simmering engine, the smell of exhaust fumes, etc. That whole car experience makes a car a car for people. If computers are ever going to reach that human level, whatever it is, all these things will also have to be parameterized.

The question is whether it is necessary to reach that human level. The important lesson, that computers are creative in their own way and humans in their own way, should be sufficient to strive for optimal cooperation: a new creative era in which HUMAN-GAN creation is central. The computer is better at getting data together and makes connections in its own way. If we go back to Thomas Edison's words, “Creativity is 98 percent perspiration and 2 percent inspiration”, a lot of perspiration will have to come from the machines. In turn, the output of the computer also inspires people. In English we would speak of new forms of ‘augmented creativity’. It is the beginning of the era of HUMAN-GAN creativity. Analyst firm Gartner argues that the advent of creative computers requires a completely different interplay between the AI-enabled professionals and the natural talents in the field of creativity in the organization. For the second group it is probably difficult to digest, but it is inevitable that autonomous creative machines will play an important role. All roles will continue to exist side by side and work together. The CIO must pave the way to that HUMAN-GAN future.

The Super Creative Ensemble Team of the Future



A natural creative



AI supported creative



AI creative (autonomous)

3.4 Finally: human creativity revealed


People who consider themselves creative may find it hard to accept that machines can also be creative. We learn from an experiment that human creativity can be a “mind fuck” and is by definition a product of the environment in which we operate. In a television program by mentalist Derren Brown, two marketers are asked to come up with an original idea. At the end, an envelope is opened with an almost identical image to that just made by these people. Mentalist Derren Brown performed this trick with two creatives from an advertising agency. He commissioned them to design a poster advertising a shop for stuffed animals. The poster had to contain a self-designed company name, slogan and logo. Below is the image of the shop concept that they came up with; they called it “Animal heaven” with the slogan “The best place for dead animals”. The logo is a bear playing on a harp.

Brown then shows his own pre-made poster. It is almost identical, with the name “Creatures Heaven”, the slogan “Where the best dead animals go” and the logo of a bear playing the harp. The creatives are stunned. How is this possible? But the taxi ride to the experiment was planned in such a way that they subtly saw images that inspired them to come up with their concept – a zoo, children wearing a T-shirt with a bear playing the harp, etc. These two creatives turned out to be less creative than they had thought and were blown away. The limited creativity from this experiment comes from the created bubble – the taxi ride and everything the creatives went through. Every day hordes of people go to work in the same bubble of a limited environment, trying to be creative and thinking they are creative, until a machine comes along that also seems to be able to cough up ideas.



Mentalist Derren Brown manipulated the creative process of two marketers so that he could predict the output.





4 In conclusion: a future CIOs dream of


The entire development of GANs and computer creativity, from 2014 to the present day, can be summed up in a chain of dreams. Dreams about the possibility of making a “dream computer” have become dreams about the possibilities of computer creativity, which in turn have become dreams about what this will bring. If Ian Goodfellow’s students hadn’t asked that one question, if computers can dream, we probably wouldn’t be having this discussion with you right now. After a night of intense work, Goodfellow provided initial proof that computers are potentially creative. By now, there are over five hundred different types of GANs, with illustrious names like 3D-RecGan, CowboyGan and SiftingGan, that can be found on GitHub for use directly by anyone. Which brings us to the question of whether the CIO, like the scientists, artists and designers who are at the forefront of this technology, is dreaming about GANs, about computer creativity, perhaps even about an infinitely more creative organization. In other words:

Can CIOs still dream?

An often-heard cry is: “There still is so much to do.” If you’re dreaming, it’s more like nightmares, because how do you keep all the balls in the air at once? According to consulting firm McKinsey it is the worst time for the average CIO to practice their profession, “There’s no worse time than now to be an average CIO” is the opening sentence of a report in which McKinsey mentions the CIO challenge. There is so much to do on so many fronts. The profession has become much more complex. The technological possibilities are much greater, social and cultural changes have accelerated, there is a talented young generation with a different value pattern and different wishes. Adding it all up, we can conclude that there is no longer room for an “average” CIO, we need a new technology leader instead.⁵⁷ It wouldn’t be out of place to add a name to the position of CIO, the Chief Creativity Officer. After all, everything that fits the profile of the new CIO requires an enormous dose of organizational creativity.

The tech leader according to McKinsey

This new tech leader is *transformative*, as we can read in McKinsey’s “The CIO challenge: Modern business needs a new kind of tech leader.” It’s a *business man or woman* who knows everything about the customer, who’s a *change*



agent addressing the “why” question, a *talent scout* whose main task is to *reimagine* what it takes to attract and retain talent, a *cultural revolutionary* who knows everything about community building and breaking down silos, and finally a *tech translator*, someone who helps the organization better understand technology and business implications.

If the CIO isn’t dreaming of a creative future anymore, who is?⁵⁸ You’d say there’s no better topic to tackle than the advent of computers that stimulate creativity. Because the challenges organizations face today demand creativity on all fronts. We previously discussed this subject in *Utopia for Executives*, our earlier report.

The conclusion we drew in that report is that the CIO must be creative on three different fronts: lead the transformation of the organization; intensify digital technology and extract business value; and give direction to a new course – purpose directionality – that fits in with the new value patterns in society.

In this final chapter, we present seven dreams and nightmares, from large to small, putting them into perspective in a world where the arrival of creative machines can be accommodated.

1 A big dream: preparing the organization for its AlphaGo moment

World Go champion Lee Sedol losing in his game against the computer has become a symbol of the creative breakthrough of machines. This dream of computer creativity (which ends with your own AlphaGo moment) has a different course than that of Sedol. He stopped playing Go three years after his loss because he couldn’t win anymore. Kasparov, who experienced his “Go moment” earlier in his confrontation with chess computer Deep Blue, chose another option: he wondered what we could learn from the computer and how we could change the game (chess) with it. Organizations are no chess pieces. The game is not about winning, it’s about staying relevant. Simon Sinek, organizational consultant and known for the why-what-how theory, introduces two types of organizations in his latest book: finite and infinite organizations. Organizations that play the game to win are the losers in the long run because they have no relevance. In his book *The Infinite Game*, he describes how higher goals bring long-term success to the organization. In that sense, the AlphaGo moment is not an end point, it’s only the starting point. For the transformative CIO we introduce here, it must be grist to the mill. The AlphaGo moment is an

excellent dot on the horizon. A new goal has been set for the digital intensification of the organization. The organizational transformation – cross-silos and in agile teams – thrives well with GAN technology. And what the ‘purpose’ is – the goal that serves all that creativity – is pre-eminently a conversation you want to enter into with the young talent on the job market.

2 A small dream: getting started with GANs

The fact that there are now creative machines is no guarantee that organizations will become more creative. Technology is a social construction. Success or failure depends on the soft factors. This dream could be about the hackathons you’ll be organizing in this area, how to break through the silos in your organization, how to attract and inspire new talent from outside, how to set up AI living labs. Of course, you could make a case for combining specific characteristics of synthetic data and GANs and other AI possibilities. Many organizations are now making steps with AI and this development can be taken into account. This is a dream that starts small and, hopefully, ends big. If we look at the examples in this report, we can conclude that humanity has been given an extra tool, like the wheel, hammers, chisels, books, stargazers, dykes, electricity, roads, cars and medicine. We live in a human-built world. In this case, however, we are dealing with an instrument from the digital subcategory. Digital tools have special capabilities that we still look at with amazement when, for example, we FaceTime with a family member in Australia. The ability to bring together things that do not originally belong together – such as bridging differences in time and place via Facebook – is just one example of what digital tools can and physical tools cannot do. Now we are discovering a new capacity: the extraordinary ability to generate unique and usable output, the mechanical ability to boost our human creativity. In concrete terms, getting it to work means keeping both feet on the ground and, at the same time, daring to be open to the remarkable characteristics of digital and learning to understand them better.

3 A nightmare: information inflation and manipulation

The boss of an English energy company thought he was on the phone with the director of the German parent company. At his request, he wrote a cheque of €220,000 to a Hungarian supplier. The voice of the parent company turned out to be a GAN creation, the Hungarian supplier was part of a criminal network. It is a remarkable fraud

case and Forrester Research expects that the damage caused by this type of case will run into the hundreds of millions in the not too distant future. According to Forrester, the cure for these problems is to get the most out of deepfake education for the organization. There is also an emerging market for detecting deepfakes. But if the phishing emails become phishing phone calls or videos where we hear people we know saying and doing things, what do we do? The energy company got off easy with a simple bank transfer. What if the CEO had given access to the information systems?

Incidentally, the nightmare is shared by 77 percent of CISOs (Chief Information Security Officers), according to a survey of 100 of these people.⁵⁹ The main fear they have is that facial recognition software becomes easy to hack. On networks outside the organization, such as YouTube, LinkedIn and Facebook, employees make connections with fake people made out of synthetic material, with fake histories and fake intentions. We are heading for a world where your senses can no longer be trusted, where a photo of damage claimed can be synthetic, where a phone call and the voice that is recognized can’t be trusted. For anyone who has the word “information” on their business card, the alarm bells must have begun ringing long ago. Identification becomes “check, check and double check” and information education becomes lifelong learning. The question is: are people able to cope with the erosion of truth and what happens when democracies come under constant attack from a cordon of true-to-life lies?


4 A dual dream: creativity has a democratizing and monopolizing effect

Democratization and the internet were once almost synonymous. But now we know better. A lot of power is in the hands of a few technology giants. Policymakers and governments have woken up and started to fight the supremacy of this power base. In the latest technology development, we are once again talking about democratization. It is undeniably true that the technology has a democratization effect. For example, you can now use the same top-level artificial chess coach anywhere in the world. There are numerous GAN tools that make the profession of designer easier. The GANs for numerous applications are available on GitHub. You can use them in no time at all, and that sounds democratic too. However, there is a duality in that we also know from the history of the internet that providing the tools and using them smartly are two different things. In this dream, you’re looking at both the possibilities for broad use of

computer creativity and the side effects that can be expected. Concentrations of power of future GAN players that are going to dominate are to be expected. Dream about the possibility of taking advantage of the democratization of creativity and be wary of these new strongholds of power.

5 An ethical dream: giving direction to creativity

Creativity sounds positive but, in reality, is neutral. You can use it for something good, but just as easily for bad things. Now that organizations are under a magnifying glass to answer the question of what is good or bad (from their ecological footprint and other environmental effects, to gender-neutral policy, inclusiveness, etc.), the question arises about the direction in which computer creativity should be deployed. This CIO dream is about what the good things are and what ethical questions need to be asked. Principal economist of the World Bank Paul Romer has an outspoken opinion on this issue. He won the Nobel Prize in Economics in 2018 for his research into innovation and growth. His lesson is that the creative forces must be steered in the right direction. That leads to *creative efficiency*. You can also steer in the wrong direction. As an example, Romer mentions fracking, the hydraulic fracturing of the soil with chemicals. Someone has sent creative minds into this dead-end street, because fracking is extremely bad for the environment and has no future. The same people could, for example, have developed new photovoltaic cells. This issue will only increase in the future. Generative AI can, for example, consciously (deep-fakes) and unconsciously (e.g. biases in the data) lead to harmful consequences. The bitcoin, to name another technology, consumes as much energy as a country like Austria. And the following example shows that it pays to think about the purpose for which AI is used.



Run three nuclear power plants for one hour

When OpenAI researchers demonstrated an algorithm at the end of 2019 that was capable of solving a Rubik's Cube using a robot hand, it didn't take long to calculate how much energy it took to train this algorithm. It required more than 1000 desktop computers, plus a dozen or so machines running specialized graphics chips that made intensive calculations over several months. Evan Sparks, CEO of Determined AI, a start-up that creates software to manage AI projects, tells Wired Magazine that the algorithm probably consumed about 2.8 gigawatt hours of electricity,⁶⁰ which is equal to the output of three nuclear power plants for one hour.⁶¹ It is predicted that if we don't work on more efficient technology, 51 percent of global electricity needs will come from communications technology by 2030.⁶²

6 A second nightmare: creative machines take over humans

Get people out of this nightmare as quickly as you can. Computers aren't going to take over from us, if only because they don't feel like it and because they don't have any sense of awareness. The dream or nightmare in which this subject can end up is unproductive. The pitfall towards which we are heading was described in crystal clear fashion by Sigmund Freud in *Das Unheimliche*⁶³ over a century ago in 1919: people attribute human traits to machines that aren't there. In Freud's time, it was moving dolls that brought us to these fallacies. Humans are psychologically defenseless against these so-called anthropomorphic thoughts. Creative machines are getting better and better at feigning humanity, as we saw with the text generator Talktothetransformer, or as shown in the synthetic human images of Thisisnotaperson. The mental step from seeing something human to suspecting that these things also have a consciousness is quickly made. Getting out of the dream is therefore something that is integral to the CIO's task, because the fantasy images come up automatically. The trap is called: "What if ...? Yeah, but what if computers gain consciousness now?" This "What if" is a Hollywood construction, not a construction of reality.⁶⁴ The creative powers of the machines will undoubtedly continue to increase and will amaze us time and time again. Let's hope that humans are smart enough to push this new creative potential in the right direction and solve the concrete problems – big and small – that we face today. Anyone who is afraid that the challenges will

run out when the computer keeps coming up with creative solutions is simply not looking hard enough. The human imagination is virtually infinite, creativity can come infinitely far within the limits of physical laws, although challenges will continue to present themselves. As long as humans still have questions, the computer must come up with the answer. No human or machine ever has to be bored again.

7 A future dream: the (in)finite creative future

The above-average CIO is a *change agent*, *transformer* and *cultural revolutionary* according to the McKinsey profile description. You stick out your neck if you also position yourself as a "creativity booster". That's why we conclude by speculating a little with you about the possible limitations that lie ahead for us in the creative field. In their research paper *Are ideas getting harder to find?* (2019), economists Nicholas Bloom, Charles I. Jones, John Van Reenen and Michael Web of Stanford University and MIT argue that ideas are getting harder to find. Based on their analysis of different industries, they draw the conclusion that the number of people doing research around the world is increasing noticeably, but that it is producing fewer ideas than before, which is slowing down economic growth. The Stanford/MIT researchers explain that the easy ideas have already been conceived and that we are now only saddled with the difficult issues. The "low-hanging fruit" in terms of inventions has now been picked. *New Scientist* reported back in 2005 that we were heading towards a "dark age of innovation". The low point would be around 2024, when innovations would dry up. Science journalist John Horgan is a source of inspiration for these kinds of thoughts. In his book *The end of science*⁶⁵ (1996) he argues that science is drawn-out; it only has to deal with really big issues like quantum physics and the big-bang theory, we (humans) solved the rest. But the really big things are one size too big for humans. A philosophical school, called the *new mysterians*, endorses this idea. Certain things will always remain a mystery to mankind. The anarchist philosopher David Graeber wrote an essay on the tantalizing proposition that we are not getting along anymore: *Of Flying Cars and the Declining Rate of Profit*⁶⁶ (2012). He, too, wonders why progress has stopped: "Where are the flying cars? Where are the force fields, traction beams, teleportation cells, anti-gravity sledges, tricorders, immortality pills, colonies on Mars, and all the other technological wonders that every child – who grew up in the middle to the end of the twentieth century – expected to have by now? Even the inventions that were about to break through – such as cloning and

cronyism – didn't keep their promises. What happened to them?"

In our dreams, we are teleported from one planet to another without having to worry about Albert Einstein's Theory of Relativity. Our imagination is limitless, but whether we can turn what we imagine into action depends on our creative ability and its possible limits. With the advent of a new creative potential – the GAN computer – we can, at least, see that these limits have been stretched further than ever.



Back in 1965, futurologists promised that we would move around in flying saucers within ten years. They're still not here.

Notes

1. Wired (2015). Google Deep Dream. <https://www.wired.co.uk/topic/google-deep-dream>
2. Osborne, M., Eynon, R., Ploin, A., & Duckworth, P. (n.d). Creative algorithmic intelligence: Capabilities and complementarity. Oxford Internet Institute. <https://www.oii.ox.ac.uk/research/projects/creative-algorithmic-intelligence?overview>
3. IBM (n.d.). The quest for AI creativity. <https://www.ibm.com/watson/advantage-reports/future-of-artificial-intelligence/ai-creativity.html>
4. Miller, A.I. (2019). *The artist in the machine: The world of AI-powered creativity*. MIT Press. <https://mitpress.mit.edu/books/artist-machine>
5. Hindupur, A. (n.d.). *The Gan Zoo*. GitHub. <https://github.com/hindupuravinash/the-gan-zoo>
6. Freedman, D. (2019, December 18). Hunting for new drugs with AI: The pharmaceutical industry is in a drug-discovery slump. How much can AI help? *Nature*. <https://www.nature.com/articles/d41586-019-03846-0>
7. Cockburn, I., Henderson, R., & Stern, S. (2019). The impact of artificial intelligence on innovation: An exploratory analysis. In Agrawal, A., Gans, J., & Goldfarb A. (Eds.), *The economics of artificial intelligence: An agenda* (pp. 115-146). University of Chicago Press. <https://www.nber.org/chapters/c14006>
8. Insilico Medicine (2019, September 9). Generative Tensorial Reinforcement Learning (GENTRL) model. GitHub. <https://github.com/insilicomedicine/GENTRL>
9. Wakefield, J. (2020, January 30). Artificial intelligence-created medicine to be used on humans for first time. BBC. <https://www.bbc.com/news/technology-51315462>
10. Thaler, S.L. (2019, September 24). Electro-optical device and method for identifying and inducing topological states formed among interconnecting neural modules. United States patent and trademark office. <http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO2&Sec2=HITOFF&p=1&u=/netahtml/PTO/search-bool.html&r=1&f=G&l=50&co1=AND&d=PTXT&s1=Stephen.INNM.&s2=Thaler.INNM.&OS=IN/Stephen+AND+IN/Thaler&RS=IN/Stephen+AND+IN/Thaler>
11. Davenport, T.H. & Patil, D.J. (2012, October). Data Scientist: The sexiest job of the 21st century. *Harvard Business Review*. <https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century>
12. Reddy, B., Sundararajan, A., & Naidu, S. (2020). Reality Engines.AI. <https://realityengines.ai/>
13. Tshitoyan, V., Dagdelen, J., Weston, L., Dunn, A., Rong, Z., Kononova, O., ... Jain, A. (2019). Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571, 95-98. doi:10.1038/s41586-019-1335-8
14. The specific GAN used here is called word2vec.
15. Chao, J. (2019, July 3). With little training, machine-learning algorithms can uncover hidden scientific knowledge. Berkeley Lab. <https://newscenter.lbl.gov/2019/07/03/machine-learning-algorithms-can-uncover-hidden-scientific-knowledge/>
16. Castelvechi, D. (2019, November 7). AI Copernicus 'discovers' that Earth orbits the Sun. *Nature*. <https://www.nature.com/articles/d41586-019-03332-7>
17. Kazuhiro, K., Werner, R.A., Toriumi, F., Javadi, M.S., Pomper, M.G., Solnes, L.B., ... Rowe, S.P. (2018, December). Generative Adversarial Networks for the Creation of Realistic Artificial Brain Magnetic Resonance Images. *Tomography (Ann Arbor, Mich.)*, 4, 159-163. doi:10.18383/jtom.2018.00042
18. Cognilytica Research (2019, January 31). Data engineering, preparation, and labeling for AI 2019. <https://www.cognilytica.com/2019/03/06/report-data-engineering-preparation-and-labeling-for-ai-2019/>
19. VanRullen, R. & Reddy, L. (2019, May 21). Reconstructing faces from fMRI patterns using deep generative neural networks. *Communications Biology*, 2, 193. doi:10.1038/s42003-019-0438-y
20. Nishimoto, S., Vu, A.T., Naselaris, T., Benjamini, B.Y., & Gallant, J.L. (2011, September 22). Reconstructing visual experiences from brain activity evoked by natural movies. *Elsevier*. doi.org/10.1016/j.cub.2011.08.031
21. Rocher, L., Hendrickx, J.M., & de Montjoye, Y. (2019). Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications*, 10, 3069. doi:10.1038/s41467-019-10933-3
22. Joshi, C. (2019, February 21). Generative adversarial networks (GANs) for synthetic dataset generation with binary classes. Data Science Campus. <https://datasciencecampus.ons.gov.uk/projects/generative-adversarial-networks-gans-for-synthetic-dataset-generation-with-binary-classes/>
23. Sogeti (n.d.). The power of synthetic data. <https://www.sogeti.nl/nieuws/artificial-intelligence/blogs/de-kracht-van-synthetische-data>
24. Thacker, J. (2019, November 8). Sneak peek: Adobe's project sweet talk and go figure. CG Channel. <http://www.cgchannel.com/2019/11/sneak-peek-adobes-project-sweet-talk-and-go-figure/>
25. Yildirim, G., Seward, C., Bergmann, U. (2018, June 20). Disentangling Multiple Conditional Inputs in GANs. Arxiv. <https://arxiv.org/abs/1806.07819>
26. Chaillou, S. (2019, July 17). ArchiGAN: a generative stack for apartment building design. NVIDIA developer blog. <https://devblogs.nvidia.com/archigan-generative-stack-apartment-building-design/>

27. Chaillou, S. (2019, February 24) AI & architecture: An experimental perspective. Towards Data Science. <https://towardsdatascience.com/ai-architecture-f9d78c6958e0>
28. Chaillou, S. (2019, July 17). ArchiGAN: a generative stack for apartment building design. NVIDIA developer blog. <https://devblogs.nvidia.com/archigan-generative-stack-apartment-building-design/>
29. Davis, M. (n.d.). NASA JPL explores generative design for space exploration. Autodesk. <https://www.autodesk.com/autodesk-university/article/NASA-JPL-Explores-Generative-Design-Space-Exploration-Mark-Davis-2019>
30. Volkswagen (2019, July 11). A new era for Volkswagen innovation in America begins with the Type 20 Concept. <http://newsroom.vw.com/vehicles/a-new-era-for-volkswagen-innovation-in-america-begins-with-the-type-20-concept/>
31. Autodesk & General Motors (n.d.). Driving a lighter, more efficient future of automotive part design. Autodesk. <https://www.autodesk.com/customer-stories/general-motors-generative-design>
32. Luo, Z., Cheng, S.Y., & Zheng, Q.Y. (2019, November 29). GAN-based augmentation for improving CNN performance of classification of defective photovoltaic module cells in electroluminescence image. IOP Publishing. <https://iopscience.iop.org/article/10.1088/1755-1315/354/1/012106/pdf>
33. Fujitsu (2019, December 11). Fujitsu wins first prize for predictive maintenance in Airbus AI challenge. <https://www.fujitsu.com/nl/about/resources/news/press-releases/2019/emeai-20191211-fujitsu-wins-first-prize-for-predictive.html>
34. Zheng, S., Farahat, A., & Gupt, C. (2019, September). *Generative Adversarial Networks for failure prediction*. Paper presented at *The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. <https://ecmlpkdd2019.org/downloads/paper/23.pdf>
35. Li, D., Chen, D., Goh, J., & Ng, S-K. (2019, January 15). Anomaly detection with Generative Adversarial Networks for multivariate time series. Arxiv. <https://arxiv.org/pdf/1809.04758.pdf>
36. Powell, K. (2017, December 3). NVIDIA Researchers showcase major advances in deep learning at NIPS. NVIDIA. <https://blogs.nvidia.com/blog/2017/12/03/nvidia-research-nips/>
37. Shiryayev, D. (2020, February 3). [4k, 60 fps] Arrival of a Train at La Ciotat (The Lumière Brothers, 1896). YouTube. <https://www.youtube.com/watch?v=3RYNThid23g>
38. Shiryayev, D. (2020, February 4). [DeOldified] Arrival of a Train at La Ciotat (The Lumière Brothers, 1896). YouTube. <https://www.youtube.com/watch?v=EqbOhqXHL7E>
39. Zakharov, E., Shysheya, A., Burkov, E., & Lempitsky, V. (2019, September 25). Few-Shot adversarial learning of realistic neural talking head models. Arxiv. <https://arxiv.org/abs/1905.08233>
40. Lee, H.Y., Yang, X., Liu, M.Y., Wang, T.C., Lu, Y.D., Yang, M.H., & Kautz, J. (2019). Dancing to Music. Arxiv. <https://arxiv.org/pdf/1911.02001.pdf>
41. Li, Y., Gan, Z., Shen, Y., Liu, J., Cheng, Y., Wu, Y., Carin, L., Carlson, D., & Gao, J., (2019, April 18). 3StoryGAN: A sequential conditional GAN for story visualization. Arxiv. <https://arxiv.org/abs/1812.02784>
42. Gupta, T., Schwenk, D., Farhadi, A., Hoiem, D., & Kembhavi, A. (2018, April 28). Imagine this! Scripts to compositions to videos. Arxiv. <https://arxiv.org/abs/1804.03608>
43. Jun-Yan, Z., Park, T., Isola, P., & Alexei, A.F. (2017). Unpaired image-to-image translation using cycle-consistent Adversarial Networks. Arxiv. <https://arxiv.org/pdf/1703.10593.pdf>
44. Rothco (2018, April 25). JFK Unsilenced Speech Highlights. YouTube. https://www.youtube.com/watch?time_continue=5&v=fOiDO8BWzfk&feature=emb_logo
45. Chesney, R., & Citron, D.K. (2018, July 21). Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3213954
46. Shead, S. (2020, January 7). Facebook to ban 'deepfakes'. BBC. <https://www.bbc.com/news/technology-51018758>
47. OpenAI (n.d.). About OpenAI. <https://openai.com/about/>
48. Howard, J. [DrJHoward]. (2019, October 26). I've created a monster. I've re-trained @OpenAI's GPT2 transformer neural network on the Pubmed/MEDLINE database, so that if I give it an article's title, it spits out an abstract for me. I didn't teach it how to structure an abstract, how long it is, or any of the lingo. [Tweet]. <https://twitter.com/drjhoward/status/1188130869183156231>
49. Boden, M., & Edmonds, E.A. (2019). *From Fingers to Digits*. MIT Press.
50. See <http://obvious-art.com/>
51. Elgammal, A., Bingchen, L., Elhoseiny, M. & Mazzone, M. (2017, June 21). CAN: Creative Adversarial Networks generating "art" by learning about styles and deviating from style norms. Arxiv. <https://arxiv.org/pdf/1706.07068.pdf>
52. In the frequently quoted paper *The Standard Definition of Creativity*, cognitive psychologists Mark Runco and Garrett J. Jaeger investigate the origins of different definitions of creativity. They conclude that since 1953 a simple two-part definition actually prevails, the standard definition of creativity: originality and effectiveness. Runco, M., & Garrett, J. (2012, February 10). The standard definition of creativity. In *Creativity Research Journal*. https://www.researchgate.net/publication/254301596_The_Standard_Definition_of_Creativity

53. Oxford Internet Institute. (n.d.). Creative Algorithmic Intelligence: Capabilities And Complementarity. University of Oxford. <https://www.oii.ox.ac.uk/research/projects/creative-algorithmic-intelligence/?overview>
54. Bringsjord, S., Bello, P., & Ferrucci, B. (2001, February). Creativity, the Turing test, and the (better) Lovelace test. *Minds and Machines* 11, 3-27. <https://doi.org/10.1023/A:1011206622741>
55. Tan, O.S. (2015, April 24). Flourishing creativity: education in an age of wonder. *Asia Pacific Education Review*, 16, 161-166. <https://link.springer.com/article/10.1007/s12564-015-9377-6>
56. Goodfellow, I. & Fridman, L. (2019, April 8). Generative Adversarial Networks (GANs) | Artificial intelligence (ai) podcast. YouTube. <https://www.youtube.com/watch?v=Z6rxFNMGdn0>
57. Dhasarathy, A., Gill, I., & Khan, N. (2020, January). The CIO challenge: Modern business needs a new kind of tech leader. McKinsey Digital. <https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/the-cio-challenge-modern-business-needs-a-new-kind-of-tech-leader>
58. This is also the conclusion of our recently published report: *Utopia for Executives*. <https://labs.sogeti.com/research-topics/digital-happiness/#utopia-for-executives-4-of-4>
59. Scroton, A. (2020). Fintechs fear deepfake fraud. *ComputerWeekly*. <https://www.computerweekly.com/news/252477448/Fintechs-fear-deepfake-fraud>
60. Knight, W. (2020, January 21). AI can do great things—if it doesn't burn the planet. *Wired*. <https://www.wired.com/story/ai-great-things-burn-planet/>
61. An OpenAI spokesperson doubted the calculation, but OpenAI refused to disclose further details of the project or provide an estimate of the electricity consumed. In one fell swoop, a 'jolly project' becomes an 'energy-guzzling misplaced joke'..
62. Andrae, A., Edler, T. (2015). On global electricity usage of communication technology: Trends to 2030. *Challenges*, 6, 117-157. <https://doi.org/10.3390/challe6010117>
63. Freud, S. (1919). *Das Unheimliche*. Gutenberg. <http://www.gutenberg.org/ebooks/34222>
64. Also read our report: *The Frankenstein Factor: the anatomy of fear of AI*. <https://labs.sogeti.com/research-topics/machine-intelligence/#frankenstein-factor-anatomy-fear-ai>
65. Wikipedia (n.d.). John Horgan (journalist). Consulted in 2019 at [https://en.wikipedia.org/wiki/John_Horgan_\(journalist\)](https://en.wikipedia.org/wiki/John_Horgan_(journalist))
66. Graeber, D., (2012, March). Of flying cars and the declining rate of profit. *The Baffler*. <https://thebaffler.com/salvos/of-flying-cars-and-the-declining-rate-of-profit>

Image credits

- p. 3 (left): Miller, A.I. (2019). *The artist in the machine: The world of AI-powered creativity* [Book cover]. MIT Press. <https://mitpress.mit.edu/books/artist-machine>
- p. 3 (middle): Runway ML. (n.d.). [Screenshot of website]. <https://runwayml.com>
- p. 3 (right): Sautoy, M. (2019, February 18). *The Creativity Code: How AI is learning to write, paint and think*. 4th Estate.
- p. 4 (top right): The Ancaro Project. (2019, January 16). *They know you the most* [Photograph]. Unsplash. <https://unsplash.com/photos/6VQIKJp2vpo>
- p. 4 (bottom left): Insilico Medicine (n.d.). <https://insilico.com>
- p. 4 (bottom right): Zhavoronkov, A., Ivanenkov, Y.A., Aliper, A. et al. (2019). Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, 37, 1038-1040. <https://doi.org/10.1038/s41587-019-0224-x>
- p. 5: Zhavoronkov, A., Ivanenkov, Y.A., Aliper, A. et al. (2019). Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, 37, 1038-1040. <https://doi.org/10.1038/s41587-019-0224-x>
- p. 6: Insilico Medicine (n.d.). <https://insilico.com>
- p. 7: Sogeti
- p. 10 (top): Smith, D. & Burke, B. (2019, August). Hype Cycle for Emerging Technologies, 2019. Gartner. <https://www.gartner.com/en/documents/3956015/hype-cycle-for-emerging-technologies-2019#:~:text=The%202019%20Hype%20Cycle%20highlights,world%20and%20mimic%20human%20creativity>
- p. 13: Tshitoyan, V., Dagdelen, J., Weston, L., Dunn, A., Rong, Z., Kononova, O., ... Jain, A. (2019). Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571, 95-98. doi:10.1038/s41586-019-1335-8
- p. 14: Cognilytica Research (2019, January 31). Data engineering, preparation, and labeling for AI 2019. <https://www.cognilytica.com/2019/03/06/report-data-engineering-preparation-and-labeling-for-ai-2019/>

- p. 15 & 16: VanRullen, R. & Reddy, L. (2019, May 21). Reconstructing faces from fMRI patterns using deep generative neural networks. *Communications Biology*, 2, 193. doi:10.1038/s42003-019-0438-y
- p. 17: Joshi, C. (2019, February 21). Generative adversarial networks (GANs) for synthetic dataset generation with binary classes. Data Science Campus. <https://datasciencecampus.ons.gov.uk/projects/generative-adversarial-networks-gans-for-synthetic-dataset-generation-with-binary-classes/>
- p. 18: Price, A. (2018, November 5). The Next Leap: How A.I. will change the 3D industry - Andrew Price [Screenshot of video]. YouTube. <https://www.youtube.com/watch?v=FlgLxSLsYWQ>
- p. 19 (top): Thacker, J. (2019, November 8). Sneak peek: Adobe's project sweet talk and go figure. CG Channel. <http://www.cgchannel.com/2019/11/sneak-peek-adobes-project-sweet-talk-and-go-figure/>
- p. 19 (bottom): Yildirim, G., Seward, C., Bergmann, U. (2018, June 20). Disentangling Multiple Conditional Inputs in GANs. Arxiv. <https://arxiv.org/abs/1806.07819>
- p. 20: Chaillou, S. (2019, February 24). AI & architecture: An experimental perspective. Towards Data Science. <https://towardsdatascience.com/ai-architecture-f9d78c6958e0>
- p. 21 (top): Davis, M. (n.d.). NASA JPL explores generative design for space exploration. Autodesk. <http://www.autodesk.com/autodesk-university/article/NASA-JPL-Explores-Generative-Design-Space-Exploration-Mark-Davis-2019>
- p. 21 (middle): Volkswagen (2019, July 11). A new era for Volkswagen innovation in America begins with the Type 20 Concept. <http://newsroom.vw.com/vehicles/a-new-era-for-volkswagen-innovation-in-america-begins-with-the-type-20-concept/>
- p. 21 (bottom): Autodesk & General Motors (n.d.). Driving a lighter, more efficient future of automotive part design. Autodesk. <https://www.autodesk.com/customer-stories/general-motors-generative-design>
- p. 22 (left): Luo, Z., Cheng, S.Y., & Zheng, Q.Y. (2019, November 29). GAN-based augmentation for improving CNN performance of classification of defective photovoltaic module cells in electroluminescence image. IOP Publishing. <https://iopscience.iop.org/article/10.1088/1755-1315/354/1/012106/pdf>
- p. 22 (right): Powell, K. (2017, December 3). NVIDIA Researchers showcase major advances in deep learning at NIPS. NVIDIA. <https://blogs.nvidia.com/blog/2017/12/03/nvidia-research-nips/>
- p. 23: Event Horizon Telescope (2019, April 10). <https://eventhorizontelescope.org>
- p. 24 (top): NVIDIA (2018, April 23). Research at NVIDIA: AI Reconstructs Photos with Realistic Results. YouTube. <https://www.youtube.com/watch?v=gg0F5JjKmHA>
- p. 24 (bottom): Shiryayev, D. (2020, February 3). [4k, 60 fps] Arrival of a Train at La Ciotat (The Lumière Brothers, 1896). YouTube. <https://www.youtube.com/watch?v=3RYNThid23g>
- p. 25: Li, Y., Gan, Z., Shen, Y., Liu, J., Cheng, Y., Wu, Y., Carin, L., Carlson, D., & Gao, J., (2019, April 18). 3StoryGAN: A sequential conditional GAN for story visualization. Arxiv. <https://arxiv.org/abs/1812.02784>
- p. 26 (top): Gupta, T., Schwenk, D., Farhadi, A., Hoiem, D., Kembhavi, A. (2018, April 10). Imagine this! Scripts to compositions to videos. Arxiv. <https://arxiv.org/abs/1804.03608>
- p. 26 (bottom): Jun-Yan, Z., Park, T., Isola, P., & Alexei, A.F. (2017). Unpaired image-to-image translation using cycle-consistent Adversarial Networks. Arxiv. <https://arxiv.org/pdf/1703.10593.pdf>
- p. 27 (top): Future of Life Insitute (2018, February). The malicious use of artificial intelligence: forecasting, prevention, and mitigation. https://www.eff.org/files/2018/02/20/malicious_ai_report_final.pdf
- p. 27 (bottom): Burgung, H. (2019, November 21). In event of moon disaster - Nixon deepfake clips. YouTube. <https://www.youtube.com/watch?v=yaq4sWFvnAY>
- p. 29 (2020, January). 9gans: An art gallery created by artificial intelligence. <https://9gans.com/>
- p. 30: Obvious (collective) (2018). Edmond de Belamy. <http://obvious-art.com>
- p. 31: AICAN (n.d.). <https://www.aican.io>
- p. 32: iStock/Devrimb
- p. 37: Brown, D. (2012, September 12). Derren Brown advertising agency task. YouTube. <https://www.youtube.com/watch?v=YQXe1CokWqQ>
- p. 42: Vaughan, J. (2009, September 24). 1954 Flying saucers for everybody. Flickr. https://www.flickr.com/photos/x-ray_delta_one/3949834600/

About VINT labs.sogeti.com

VINT, the Sogeti research institute and part of SogetiLabs, provides a meaningful interpretation of the connection between business processes and new developments. In every VINT publication, a balance is struck between factual description and the intended utilization. VINT uses this approach to inspire organizations to consider and use new technology. VINT research is done under the auspices of the Commission of Recommendation, consisting of • F.M.R. van der Horst (chairman), Chief Executive Officer Retail Banking ABN AMRO Bank N.V. • A. Arooni, Chief Information Officer NN Group • K. van Baardwijk, Chief Operating Officer and member of the Executive Committee of Robeco • Jørgen Behrens, Head of Applications and Services at HERE Technologies • M. Boreel, Chief Technology Officer Sogeti Group • Paul Dirix, Chief Executive Officer AEB Amsterdam • D. Kamst, Chief Executive Officer Klooker • M. Krom, Chief Information Officer PostNL • T. van der Linden, Group Information Officer Achmea • Prof. dr. ir. R. Maes, Professor of Information & Communication Management Academy for I & M • P. Morley, Lecturer Computer Science University of Applied Science Leiden • J.W.H. Ramaekers, Head of Sogeti Netherlands • E. Schuchmann, Chief Information Officer Amsterdam UMC • W.H.G. Sijstermans, Member of the Board, Chief Information Officer Nederlandse Zorgautoriteit (NZa) • K. Smaling, Chief Technology Officer Continental Europe Aegon • J.P.E. van Waayenburg, Head of Capgemini Netherlands • M. der Weduwen, Chief Information Officer & Vice President Information Services, Holland Casino.

About SogetiLabs

SogetiLabs is a network of technology leaders from Sogeti worldwide. SogetiLabs covers a wide range of digital technology expertise: from embedded software, cybersecurity, AI, simulation, and cloud to business information management, mobile apps, analytics, testing, and the Internet of Things. The focus is always on leveraging technologies, systems and applications in actual business situations to maximize results. SogetiLabs also provides insight, research, and inspiration through reports, articles, presentations, and videos that can be accessed on the SogetiLabs website.

www.labs.sogeti.com

About Sogeti

Part of the Capgemini Group, Sogeti operates in more than 100 locations globally. Working closely with clients and partners to take full advantage of the opportunities of technology, Sogeti combines agility and speed of implementation to tailor innovative future-focused solutions in Digital Assurance and Testing, Cloud and Cybersecurity, all fueled by AI and automation. With its hands-on 'value in the making' approach and passion for technology, Sogeti helps organizations implement their digital journeys at speed.

www.sogeti.com