Vol. 9, Issue 2, September 2019

Journal of Unsolved Questions

JUnQ

The prospects and limitations of digitalization – Interview with Alex T. Steffen

How to respond to the potential malicious uses of artificial intelligence?

Some thoughts on artificial intelligence – Interview with Dr. Anton Bogomolov

Surprising creativity of digital minds

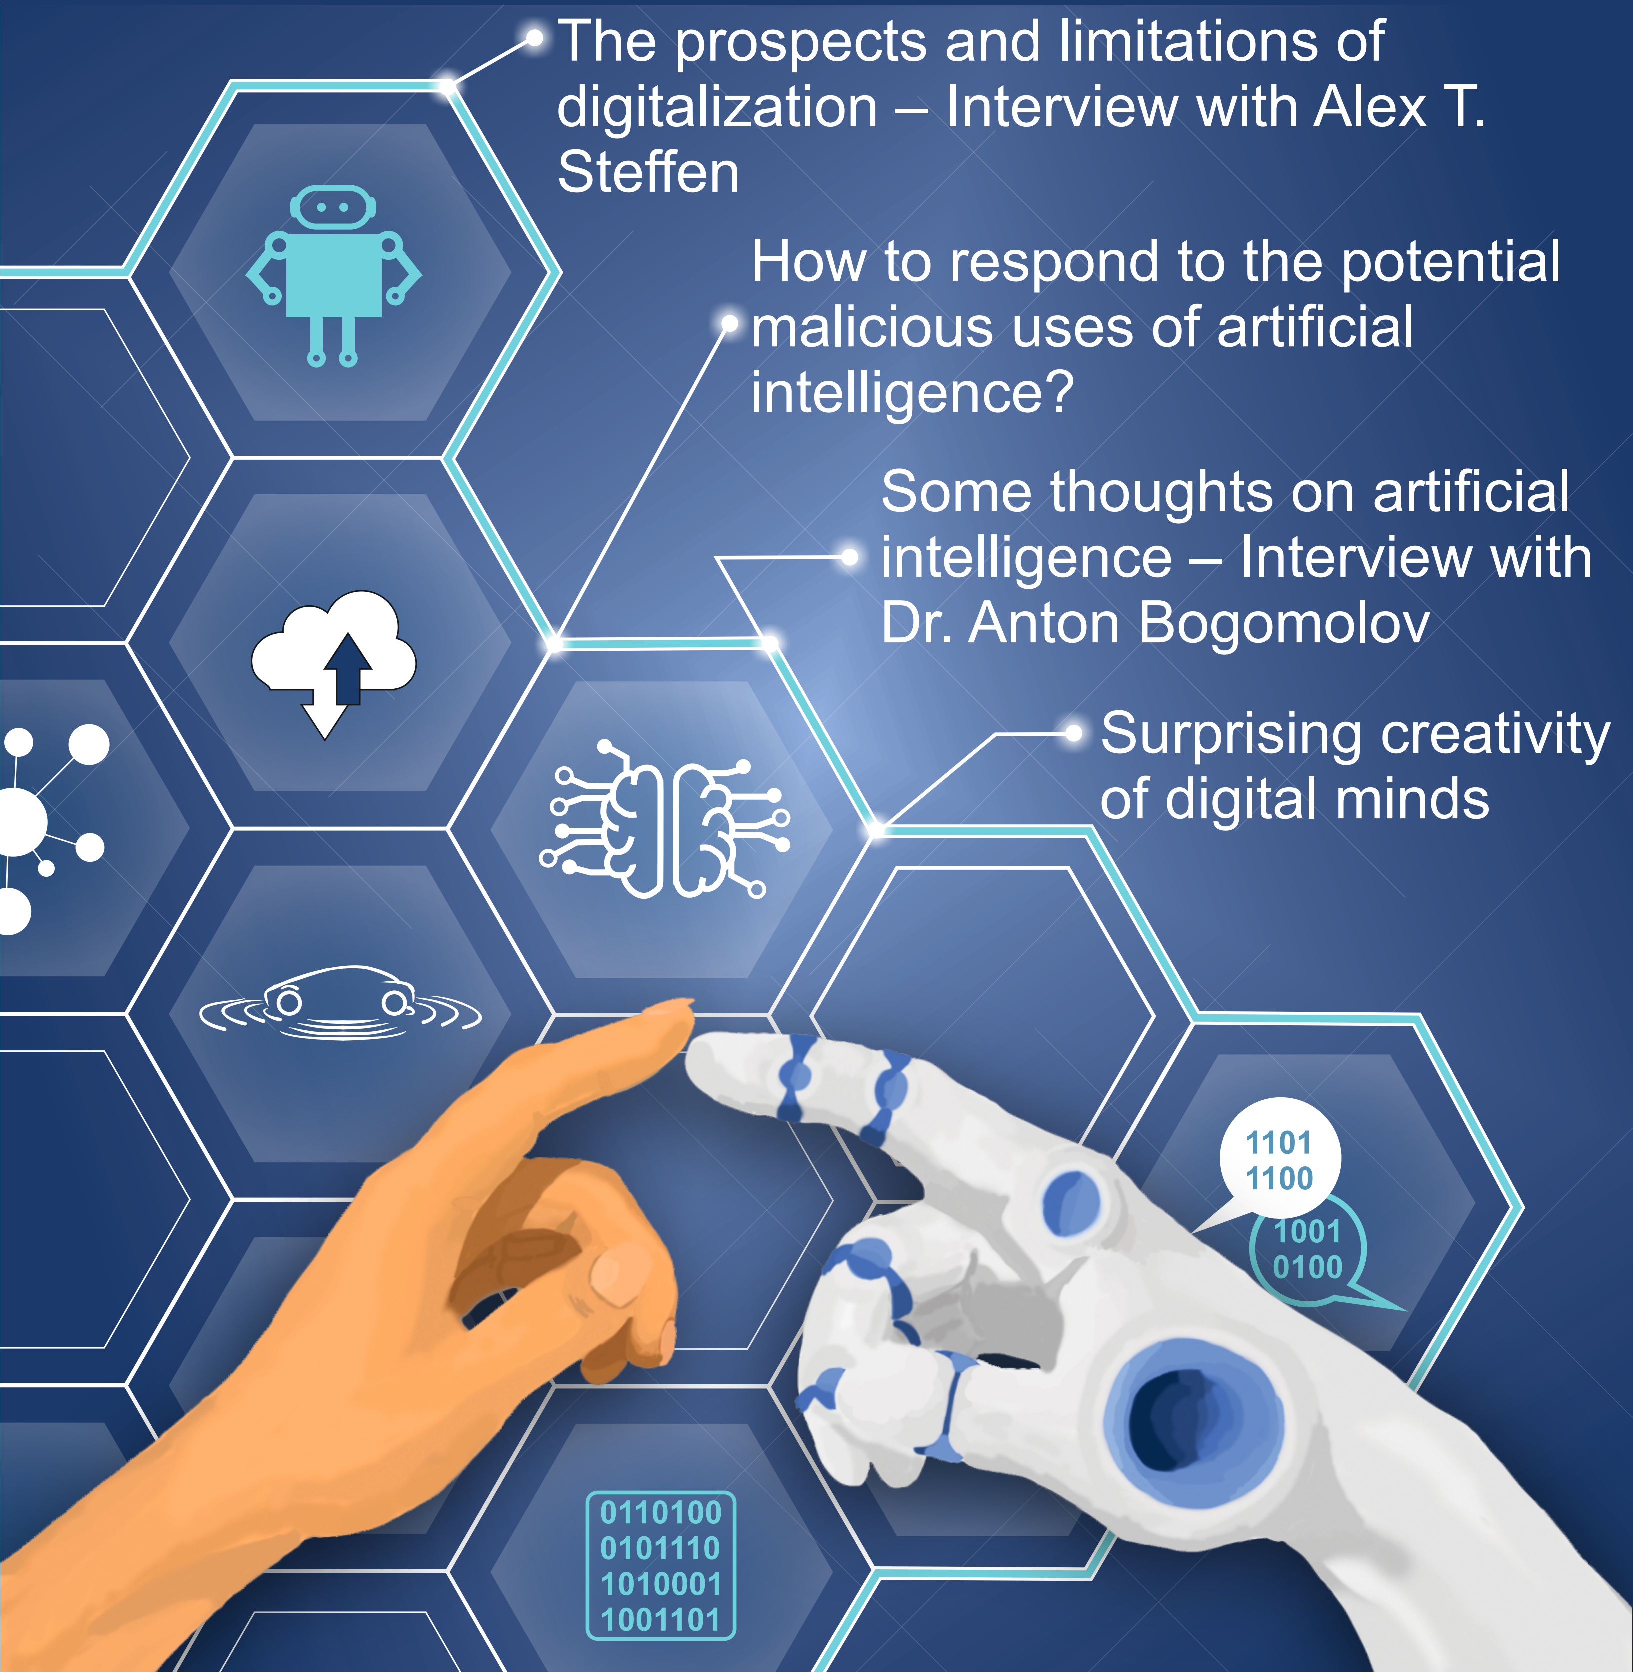

DIGITALIZATION IN SCIENCE

Table of Contents

Preface

Editorial Note

Dear Reader,

Not a single day passes that we are not confronted with problems arising from digitalization in every thinkable field: When will we travel with fully autonomous transportation? Is a digital war conceivable? Will artificial intelligence one day place itself above humankind and defeat us? Or can it treat patients way better than human physicians? Should schools use modern media to teach? How can a software beat record holders in the most advanced strategy games? Why is the internet connection so labile? Who is collecting my data and for what purpose? Did the gamers among us ever wonder how the background music of the latest games seamlessly adjusts to fit the gameplay? Have you ever heard of digital citizenship?

It is evident: everybody is worried about possible hazards of digitalization, yet everybody seems to become absolutely dependent on digital tools, be it in their social or professional lives. Self-learning systems are evoking both distrust and optimism for the future. Dr. Anton Bogomolov comments on concerns about AI and self-learning software in his interview on page 1. Malicious use of AI and counter actions are summarized in Haydn Belfield's essay on page 5. Take a look at future possibilities and responsibilities and read Alex T. Steffen's interview on page 7. Not everything coming out of a digital mind is a source of danger. Mariia Filianina's short overview about creative digital minds illustrates that AI can also be entertaining or perhaps revolutionize our way of thinking outside the box on page 11.

With this issue we hope to still some fears one might have in our digital age. After all, it is what we make of it and it might just be the next step of evolution. And as always: stay curious and dig through the JUnQ to find the hidden treasures!

—Tatjana Daenzer

Opinions

Some thoughts on artificial intelligence

Anton Bogomolov[1] is a data scientist with PhD in Physics, currently working in IoT branch. He is passionate for artificial intelligence with ten years of experience in automated data analysis and machine learning.

[1]abogomolov86@gmail.com

JUnQ: The everlasting technological progress is aimed to fulfill many needs of humans: most of them are physical, informational and commercial. In particular, robots were created to perform tasks that were too dangerous for humans or that humans could not or did not want to do. But what do we need intelligent machines for and what is implied by "Artificial Intelligence" (AI)?

Anton Bogomolov: The answer was already said – we need AI to make our life simpler, i.e. to simplify some routine work that humans have to do. Generally, we are heading towards automation, and in the ideal case, we want to automize everything, every kind of work. So far, the processes we are capable of automizing have been prioritized.

Now, what is understood by the term "AI"? Over the course of this interview we will go deeper in the discussion, so let's start with a fairly broad definition: AI is something that is able to accomplish certain tasks with the help of self-learning.

JUnQ: Does it imply that AI is not meant to create anything, like art or music?

Anton Bogomolov: There is a number of definitions of AI. Indeed, the term "intelligence" implies that it can do creative work as well. It is not a simple calculator. You don't just tell it what you want it to calculate, and then it does exactly what has been asked. It does something more complicated and, thus, it also involves some learning experience. In this context, the creative work does not necessarily mean being an artist or a musician, or a composer. A chatbot, as an example of an AI feature, is also a kind of creative work, because it is required to react accordingly or ask appropriate questions, in other words to be engaged into a conversation as a human would be i.e. express creativity.

Generally, yes, AI can generate art. For example, "Deep Dream"[1] was popular a few years back. This algorithm uses AI to generate the dream-like appearance of the uploaded images. Another one is "Neural style transfer"[2] which allows one to compose an image in the style of another image. Should one ever want to paint like Van Gogh or Picasso, this can be easily done, using this algorithm. There is also AI-composed music already creeping into the background of games, film, and media. With AI it is now possible to create music in different genres just at the push of a button.

JUnQ: In the news or podcasts, the term "machine learning" often seems to come together with AI. What is, simply put, machine learning and how does it relate to AI?

Anton Bogomolov: As I mentioned before there are many definitions of AI. In simple words, AI is a broader term than the machine learning (ML), i.e. AI includes ML. Being sort of an advanced algorithm, AI achieves specific goals by means of ML, at the same time it is able to adapt to its environment, just like humans. ML is also an algorithm, but a simpler one, with the key feature – the ability to learn (thus the name). It is not meant to achieve a global goal, its goal is to eventually enable programs to automatically improve through experience, without the programmer having to change the code. ML relies on working with data sets, that one needs to input first. It then examines and analyses the data to find common patterns, so that eventually it becomes possible to make experience-driven predictions or decisions.

JUnQ: So what it means is that AI does not exist without ML?

Anton Bogomolov: Right. Machine learning is a subset of AI, more like a tool to achieve AI. One example might be the first chatbots from the 90s. They had hardcoded "intelligence", i.e hardcoded answers to possible questions. If such bot sees certain keywords it outputs accordingly relevant keywords. These did not have machine learning. But the intelligence of these was doubtful since the algorithm did not adapt. And as we discussed previously the key asset of AI is the ability to adapt.

JUnQ: Since we are on this page, how can one tell the difference between the AI system and a more "conventional" program?

Anton Bogomolov: There are "intelligence" tests for AI, among which the most renown one is the Turing Test.[3] But this is more to test whether or not a system is capable of thinking like a human being. However, no AI technology today has passed the Turing test, i.e. that has shown to be convincingly intelligent and able to think. So, this is the main goal of this AI branch - we want to create a machine that will be indistinguishable from a human, in particular, that will be self-aware and act somewhat mindfully. In the end, such a machine will be able to pass the Turing test. Once again, so far, they do not exist. Self-awareness tuned out to be tough to realize.
Now, back to what was asked. I believe, no one is interested in differentiating AI from a mindless linear algorithm. Because as long as the desired goal is achieved no one cares what type of algorithm was used for it.

JUnQ: AI is no longer a futuristic concept, as some may naively think. Can you name some examples where is AI being used already? Are there any AI applications used in the everyday life of ordinary people?

Anton Bogomolov: The most straightforward example is our smartphones. The more recent ones can recognize the owner's face. This is known to use neural networks. Also, in smartphones, there is Google assistant. Spoken inquiries are transferred to a server where neural network-based algorithms convert them to text, and which is then processed to deliver the relevant information. These are the simplest examples. We all watch Youtube where based on one's watch history the system suggests what else one might be interested in. These AI-based recommendation engines now seem to know us to an uncanny degree.

If we now go further from everyday life, I would say AI is used pretty much in every field. In finance – there are already automatic trading robots. Some use AI for analysing financial markets to generate profitable trading strategies or make market predictions.

Autonomous driving has become very popular recently. There are even toys for children that make use of a variety of AI and ML technologies, including voice and image recognition, to identify the child and other people around, based on their voices and appearance. This all is owing to the computation power we currently have, which has advanced in the last years.

AI has found its application in medicine as well. As AI demonstrated remarkable progress in image-recognition tasks it is now widely used in medical radiology and computer tomography. One example is that there are neural networks that are trained to analyze tumours and do it as well as the top-class specialists in the field. Just as radiologists are trained to identify abnormalities based on changes in imaging intensities or the appearance of unusual patterns, AI can automatically find these features, and many others, based on its experience from the previous radiographic images, coupled with data on clinical outcomes. This also yields a more quantitative outcome, while radiologists perform only a quantitative assessment.[4]

JUnQ: As AI develops further is it going to make human jobs obsolete? And what will people be doing if there is nothing else to do?

Anton Bogomolov: Ideally, this is what we aim for - to have everything automized. But this can be achieved, in my opinion, only when so-called artificial general intelligence is realized. This will be a machine capable of experiencing consciousness and think autonomously and thus will be able to accomplish any intellectual task that a human being can.

What will happen to humans after all? There is a concept of universal basic income. The idea is that the robot replacing you is working on your behalf and you are given an income sufficient to meet basic needs, with zero conditions on that income. Because in the end the job is being done and the resources are being produced while you are free for other pursuits.

There has been a lot of research interest in this regard. Back in the 60's, there was a researcher, John Bumpass Calhoun, who reported on an experiment with rats, the experiment is also known as "Universe 25". The researchers provided rats with unlimited resources, such as water and food. Besides, they eliminated the danger otherwise coming from nature, like predators, climate, etc. Thus, the rats were said to be in "rat utopia". At first, the population peaked but shortly after it started to exhibit a variety of abnormal, often destructive behaviours. After some time of the experiment, the rats became too lazy to reproduce and the population was on its way to extinction. There is, of course, the controversy over the implications of the experiment but it can be perceived as one of the possible scenarios of the future.

JUnQ: What about the programming jobs? And scientists?

Anton Bogomolov: Well, first we automize what we can do - so far, the simplest work. AI is now partly replacing the jobs of translators and customer service work. The next in line are self-driving cars that will automize the entire transportation industry, bus, and taxi drivers and so on. But programming jobs are of a different kind, they are creative. Programs that develop other programs exist already, but they are rather limited in what they can do.

Eventually, all jobs will be replaced. Programming jobs will be among the last ones though. Just as other creative jobs, including scientists.

One day we will have a super-intelligent machine, that develops further programs similar to itself at less expense and much faster compared to when supervised by humans. At some point we might not be able to follow its advances anymore and here comes the term "technological singularity". This is believed to occur when AI starts discovering new science at enormous rates while always learning and evolving on top of it uncontrollably from human's side.

JUnQ: Is the "singularity" inevitable?

Anton Bogomolov: There is an everlasting argument whether at all it is possible to realize a self-aware AI, that will act mindfully, much like a human. Therefore, depending on "yes" or "no" there will be a technological singularity or not. It can as well occur for other reasons, it is just that among others AI is more likely to bring us to the technological singularity.

On the other hand, it is not proven that such AI can ever be created, to be able to run autonomously and replace all of us. In this case, there will be no AI-induced singularity.

So, this is now a really hot topic in the community.

JUnQ: Does it mean that self-awareness is prerequisite for a possible singularity to occur and we are not yet passed the point of no return?

Anton Bogomolov: Right. The algorithms that exist now and are known to beat the world-class champions in chess and Go are harmless. They are just trained extraordinary well on one particular subject, to achieve a well-defined goal. They are not able to think outside of the box, like "what else is there that I could do".
Once we create a machine that will be able to think this way, to exhibit human-level consciousness, it is expected to bring us to the singularity. Because it will be able to operate and develop without any supervision. All existing AI technologies do develop themselves but only to a certain degree, they do not have this freedom yet.

JUnQ: Speaking about self-awareness. For example, Sophia – the social humanoid robot developed by Hanson Robotics - realizes itself (herself) as being a programmed female robot. Does it mean that she is self-aware? How did they manage to program "her" self-realization?

Anton Bogomolov: As far as I understand she is programmed to answer this way. If there comes a question about what she thinks she is, her answer will be according to what has been built in her program. Most likely she was trained on thousands of real dialogs among people about their self-awareness. Like other AI systems, she also has machine learning that, if you feed it with enough data, will enable her to learn how to answer and how to behave, as people would.
Sophia communicates very well on a topic known in advance. Because in this case she can get trained in advance: they provide her with enough information about a given topic to get trained. Then she is able to have a sensible conversation because she has the statistics on what is typically answered when. Nevertheless, it is not as simple as when you say X, she replies Y. Thanks to machine learning what she says is a result of rather complicated non-linear connections.
I did not have a chance to speak with her personally though, but I think she is certainly not self-aware. Otherwise, the singularity would have been just around the corner by now. If she had a human-level consciousness, there would be nothing that she would need people for. She would be able to program herself to increase her memory. In just a few days she would reach the level of intelligence of all the people on Earth. In a few more days we would not be able to comprehend what level of intelligence she would have - again the exponential progress.
So, there is nothing we should worry about. She is still just a robot - more about illusion than intelligence. The shocking effect is also due to the fact that she looks like a human, has emotions and facial expressions. This unique combination of her features might make us a bit alert. And for sure Sophia is a great representation of all the advances of AI technology.
In fact, to able to realize human-level AI we essentially need to model a human's brain. The human brain contains around 10^{11} neurons. On the other hand, functional neural networks have in the order of tens of millions of neurons. These four orders of magnitude difference are sizeable. Moreover, it also takes quite some time to train a system with a large number of neurons. At the end of the day, we do not yet have the capacity to realize a human-level AI.

JUnQ: In case something goes wrong, will we able to "unplug" the machine. Do autonomous AI systems exist yet?
Autonomous systems do exist. Think of a toy-dog, that we have discussed already, or a vacuum cleaner, they are

programmed to charge when needed. These are completely autonomous as long as the power source is available. Military branch sure has got some as well. I can imagine an armed flying drone, self-charging, and self-rechargeable.
But the existing autonomous AI systems are not a threat to humans. Despite having all the advantages of machine learning they follow a defined program to accomplish a specific task. It can be the best in recognizing people's faces, shooting targets or avoiding bullets. But it is still a mindless machine, that we can destroy, or fool or at least hide all the power stations from it.
As long as any of these do not have human-level intelligence, as long as they are not smarter than us, they should not be considered as a potential threat.

JUnQ: So reaching human-level intelligence would be the point from which on AI can potentially live without us.

Anton Bogomolov: Correct. There is an opinion that biological life is just a means to create an electronic life. In other words, some believe that this is our mission, to give birth to an electronic conscious creature, surpassing our capacity, that will develop much faster than humans. In some sense, it is similar to the early times of our planet. Life on Earth began relatively early. But the first living creatures - unicellular organisms - were progressing very slowly, until the multicellular organism occurred, which boosted the progress tremendously. And the progress always seems to be exponential. Thus, the idea of this theory is that we create something to keep up to this exponential progress. And if we look at it globally, like in the scale of the Universe, if this should ever happen that AI takes over the world, it would make sense. Because AI would go further exploring the Universe much faster than we would. Thus, from the point of view of global progress, it would be more advantageous.

JUnQ: Now, when you put it this way the technological singularity does not sound so frustrating anymore. Are you optimistic overall? Will we make it to the end of the 21^{st} century?

Anton Bogomolov: To me, it feels great to witness the progress and to be a part of it. But we will see how it goes. We live within a self-organized system, where everything has got a direction to go. Even though humans are all independent creatures, we still obey the same laws of synergy, we self-organize as well, we cluster forming cities, etc. And sure we also have something to move towards, thus we develop and evolve. So, this progress is so natural.

In fact, experts expect the technological singularity to occur already in the 21^{st} century. But it is not trivial to give a correct estimate. On the other hand, not related to AI, there is research going on in the field of so-called negligible senescence. The idea is that by engineering the reversal of all the major molecular and cellular changes that occur with age we would enable us to constantly rejuvenate ourselves. The researchers believe that negligible aging for humans will be achieved in this century. There even exists a provocative opinion that the first human beings who will live to 1,000 years old are already alive.[5]

At the end of the day, there has been tremendous progress in many fields, not only AI. Along with AI, we may succeed in developing other technologies, which will help us to prolong our lives as well as humans' in general.

JUnQ: Thank you very much for the interview!

— Mariia Filianina

Read more:

[1] http://deepdreamgenerator.com
[2] L.A. Gatys, A.S. Ecker and M. Bethge *arXiv*1508.06576 (2015).
[3] https://en.wikipedia.org/wiki/Turing_test
[4] A. Hosny, C. Parmar, J. Quackenbush, L.H. Schwartz and H.J.W.L. Aerts *Nature Reviews Cancer* **18**, 500 (2018).
[5] https://www.ted.com/speakers/aubrey_de_grey

How to respond to the potential malicious uses of artificial intelligence?

Haydn Belfield

Haydn Belfield[1] is a Research Associate and Academic Project Manager at the University of Cambridge's Centre for the Study of Existential Risk. He is also an Associate Fellow at the Leverhulme Centre for the Future of Intelligence. He works on the international security applications of emerging technologies, especially artificial intelligence. He has a background in policy and politics, including as a Senior Parliamentary Researcher to a British Shadow Cabinet Minister, as a Policy Associate to the University of Oxford's Global Priorities Project, and a degree in Philosophy, Politics and Economics from Oriel College, University of Oxford.

[1]hb492@cam.ac.uk

Artificial intelligence (AI) is beginning to change our world – for better and for worse. Like any other powerful and useful technology, it can be used both to help and to harm. We explored this in a major Febuary 2018 report *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*.[1] We co-authored this report with 26 international experts from academia and industry to assess how criminals, terrorists and rogue states could maliciously use AI over the next five years, and how these misuses might be prevented and mitigated. In this piece I will cover recent advances in artificial intelligence, some of the new threats these pose, and what can be done about it.

In this piece I will cover recent advances in artificial intelligence, some of the new threats these pose, and what can be done about it.

AI, according to Nilsson, "is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment".[2] It has been a field of study from at least Alan Turing in the 1940s, and perhaps from Ada Lovelace in the 1840s. Most of the interest in recent years has come from the subfield of 'machine learning', in which instead of writing lots of explicit rules, one trains a system (or 'model') on data and the system 'learns' to carry out a particular task. Over the last few years there has been a notable increase in the capabilities of AI systems, and an increase in access to those capabilities.

The increase in AI capabilities is often dated from 2012's seminal Alexnet paper.[3] This system achieved a big jump in capabilities on an image recognition task. This task has now been so comprehensively beaten that it has become a benchmark for new systems - "this method achieves state-of-the-art in less time, or at a lower cost". Advances in natural language processing (NLP) have led to systems capable of advanced translation, comprehension and analysis of text and audio – and indeed the creation of synthetic text (OpenAI's GPT-2) and audio (Google's Duplex). Generative Adversarial Networks (GANs) are capable of creating incredibly convincing synthetic images and videos. The UK company DeepMind achieved fame within the AI field with their systems capable of beating Atari games from the 1980s such as Pong. But they broke into the popular imagination with their AlphaGo systems defeat of Lee Sedol at Go. AlphaGo Zero, the successor program, was also superhuman at Chess and Shogi. AI systems have continued to match or surpass human performance at more games, and more complicated games: fast-paced, complex, 'real-time strategy' games such as DOTA II and Starcraft II.

This increase has been driven by key conceptual breakthroughs, the application of lots of money and talented people, and an increase in computing power (or 'compute'). For example, training AlphaGo Zero used 300,000 times as much compute as AlexNet.[4]

Access to AI systems has also increased. Most ML papers are freely, openly published by default on the online depository *arXiv*. Often the code or trained AI system can be freely downloaded from open source software libraries like GitHub or TensorFlow, which also tend to standardise programming methods. People new to the field can get up to speed through online courses such as Coursera, or the many tutorials available on YouTube. Instead of training their systems on their own computers, people can easily and cheaply train them on cloud computing providers such as Amazon Web Services or Microsoft Azure. Indeed the computer chips best suited to machine learning (GPUs and TPUs) are so expensive that it normally makes more sense to use a cloud provider, and only rent the time one needs. Overall then, it has become much easier, quicker and cheaper for someone to get up to speed, and create a working system of their own.

These two processes have had many benefits: new scientific advances, better and cheaper goods and services, and access to advanced capabilities from around the world. However they have also uncovered new vulnerabilities. One is the discovery of 'adversarial examples' – adjustments to input data so minor to be imperceptible to humans, but that cause a system to misclassify an input. For example, misclassifying a picture of a stop sign as a 45 mph speed limit sign.

These vulnerabilities has prompted some important work on 'AI safety', that is, reducing the risk of accidents involving AI systems in the short-term[6,7] and long-term.[8] Our report focussed, however, on AI security: reducing the risk of malicious use of AI by humans. We looked at the short-term: systems either currently or soon to be in use in the next five years.

AI is a 'dual-use' technology - it can be used for good or ill. Indeed it has been described as an 'omni-use' technology as it can be used in so many settings. Across many different areas however, common threat factors emerge. Existing threats are expanding, as automation allows a greater scale of attacks. The skill transfer and diffusion of capabilities described above will allow a wider range of people to carry out attacks that currently the preserve of experts. Novel threats are emerging, using the superhuman performance and speed of AI systems, or attacking the unique vulnerabilities of AI systems. The character of threats is being altered as attacks become more customised to particular targets, and the distance between target and attacker makes attacks harder to attribute.

These common factors will affect security in different ways - we split them into three domains.

In 'digital security', for example, current 'spear phishing' emails are tailor-made for a particular victim. An attacker trawls through all the information they can find on a target, and drafts a message aimed at that target. This process could be automated through the use of AI. An AI could trawl social media profiles for information, and draft tailored synthetic text. Attacks shift from being handcrafted to mass-produced.

In 'physical security', for example, civilian drones are likely to be repurposed for attacks. The Venezuelan regime claims to have been targeted by a drone assassination. Even if, as is most likely, this is propaganda, it gives an indication of threats to come. The failure of British police for several days to deal with a remote-controlled drone over Gatwick airport does not bode well.

In 'political security' or 'epistemic security', the concern is both that in repressive societies governments are using advanced data analytics to better surveil their populations and profile dissidents; and that in democratic societies polities are being polarised and manipulated through synthetic media and targeted political advertising.

We made several recommendations for policy-makers, technical researchers and engineers, company executives, and wide range of other stakeholders. Since we published the report, it has received global media coverage and was welcomed by experts in different domains, such as AI policy, cybersecurity, and machine learning. We have subsequently consulted several governments, companies and civil society groups on the recommendations of this report. It was featured in the House of Lords Select Committee on AI's Report. We have run a workshop series on Epistemic Security with the Alan Turing Institute. The topic has received a great deal of coverage, due in part to the Cambridge Analytica scandal and Zuckerberg's testimony to Congress. The Association for Computing Machinery (ACM) has called for impact assessment in the peer review process. OpenAI decided not to publish the full details of their GPT-2 system due to concerns about synthetic media. On physical security, the topic of Lethal Autonomous Weapons Systems has burst into the mainstream with the controversy around Google's Project MAVEN.

Despite these promising developments, there is a lot still more to be done to research and develop policy around the malicious use of artificial intelligence, so that we can reap the benefits and avoid the misuse of this transformative technology. The technology is developing rapidly, and malicious actors are quickly adapting it to malicious ends. There is no time to wait.

Read more:

[1] Brundage, M., Avin, S., *et al.* (2018). The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation, arXiv:1802.07228.

[2] Nilsson, N. J. (2009). The quest for artificial intelligence. Cambridge University Press.

[3] Krizhevsky, A., Sutskever, I., Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* (pp. 1097-1105).

[4] Amodei, D. Hernandez, D. (2018). AI and Compute. OpenAI: https://blog.openai.com/ai-and-compute/.

[5] Karpathy, A. (2015) Breaking Convnets. http://karpathy.github.io/2015/03/30/breaking-convnets.

[6] Amodei, D., Olah, C. *et al.* (2016) Concrete Problems in AI Safety.

[7] Leike, J. et al. (2017) AI Safety Gridworlds. DeepMind.

[8] Bostrom, N. (2014) Superintelligence. Oxford University Press.

[9] House of Lords Select Committee on Artificial Intelligence (2018). Report of Session AI in the UK: ready, willing and able? 2017–19 HL Paper 100.

The prospects and limitations of digitalization

Alex Steffen[1] makes enterprises future-proof. He is an expert for business strategy and innovation. He is also a no.1 Best-Selling Author and Speaker. His mission by 2025 is to empower 150,000 business leaders to future-proof their enterprise with ease. How? Alex turns business leaders into entrepreneurs. Alex Steffen was named Management Thought Leader 2019 by Change X and his book "Die Orbit Organisation" was nominated for the getAbstract International Book Award. His Keynotes "The Atlas of Innovation" and "Unstoppable Human" are international hits. Learn about Alex at https://alextsteffen.com.

[1]info@alextsteffen.com

JUnQ: What is digital citizenship? Should there be a basic education in responsible handling of digital tools in (early) schools?

Alex T. Steffen: Let's pick a narrow definition. I understand digital citizenship as a human's ability to be a more rounded part of society thanks to information technology. The truth is: technology often simply emphasizes the existing design.
Digital schooling isn't better schooling, as long as schools fail to teach us the central skills required in the modern world: thinking for ourselves. In my opinion, that's what the society and workplace of the future needs. We're trying to stitch digital onto an outdated paradigm, which tells us that memorizing facts is fundamental to a successful career. And then we're surprised to find that machines take away jobs.
The truth: a rounded human, well-equipped to play his or her part in society combines a unique blend of complex skills. Uniqueness is an advantage, not a disadvantage. I see micro degrees, potent mentoring, and real exposure to the world as essential ingredients of education towards digital citizenship. We don't need any more homogenous machine workers. The new standard for humans and businesses is hyper-customization. A smart country isn't a country that has advanced to digital citizen services only.
A smart country is a society where its citizens can create a career and life on their own terms using highly customizable (education) resources. That will make them uniquely trained and attractive according to their strengths and inclinations. Look around, the top talents are already living this very design. Now it's our responsibility to take it from niche to commonplace.

JUnQ: What are the general problems and dangers that arise with (global) digitalization and what are possible solutions?

Alex T. Steffen: This begs the exploration of the new relationship between digital processes and human habits. Let's first crush a myth: our problem isn't the technology disrupting our lives. Humans will create what's possible. They always have. The problem lies in our own comfort to reconsider what we see as "normal", "customary" and "acceptable". Our problem is: we think that most of what we look at is permanent when in fact, the world is in constant change.
We underestimate our need for validation and our inability to accept outside perspectives. Those are the real causes of resistance. I am convinced that if we could measure the real damage of business as usual, it would vastly outweigh the so-called threats of digitization. I would like to see an approach where anything new is met with a cool-headed evaluation. Reactive resistance contra change based on individual discomfort stands in the way of realizing beneficial trends.
These trends often end up as part of our lives anyway, built by others, who were open-minded in the first place. And, equally important, a lack of engagement with trends prevents us from making them safe and aligned with our values. I suggest training leaders on emotional intelligence and on staying curious. As soft as this sounds to our logical minds, it's the vastly underestimated skill that nourishes our ability to be competitive. Innovation starts with the very subject in question: rethinking (innovating) the way we train our leaders, so that change can be embraced .

JUnQ: Data processing, communication, and research have become impossible without digital tools, especially in the field of technology and science. A regression has become unthinkable. Are there limitations to further digital progress?

Alex T. Steffen: Every society comfortable enough to explore this philosophical question faces a dilemma between

two seemingly exclusive ideas.
Idea 1: we've arrived at the pinnacle of innovation. Further innovation seems unthinkable or unethical. Further innovation causes more harm than good.
Idea 2: awe-inspiring science fiction scenarios that look completely absurd but encapsulate even more human optimization potential.
The two ideas are not exclusive. Rather, they lie on opposite poles of a scale. I'm always curious where a person or society sits on that scale. In other words, how much of each idea do they express. My take is that we often ignore the bigger picture. History can provide data for a more realistic standpoint, namely that humans will continue innovating indefinitely. It's like that because with new capabilities come ever new desires. These trigger our ingenuity anew.
This begs the question: will we be able to find a healthy balance between a paralyzing public debate about the implications of change on the one hand and co-creating the inevitable changes, so that they end up in favor of future generations? Let's look at an example: In Sweden the question of female equality at work has largely been resolved for a few years. "We focus on doing rather than talking" an executive at Volvo shared with me. In Germany, after years of debate this is still a hot topic.

JUnQ: How will the future digital workplace look like?

Alex T. Steffen: I love this question and yet I'll keep my answer deliberately vague. Nobody can predict the future with 100% accuracy. I sincerely hope that for most people the future workplace will be driven by vitality, intuition, and self-actualization. This will mean better health and quality of life for the individual as well as higher competitiveness for business. [1]

JUnQ: In Germany, digitalization appears to proceed more slowly than in other industrial countries. What are possible troubles and how can we overcome this gap?

Alex T. Steffen: All innovation starts in the mind. History is full of examples where German ingenuity put us in the pole position, only to be halted by doubt and cumbersome processes. We wake up and find ourselves late in the game. No question, their intention is good. But after some time of business as usual, further resistance to creative destruction creates more harm than good. In 2019 German car giant Volkswagen came out with its car for the future. Unfortunately the car is not an exponential innovation at all. It's traditional car with an electric engine. Major improvements still require a garage.
Tesla Motors on the other hand, has shown us what a disruption of the automotive industry really looks like. Tesla has built a digital platform on which major improvements are performed over the internet via digital upgrades. The result: the need for a garage drops drastically. So does the dependency on a complex web of stakeholders, turning Tesla Motors into the more flexible player. This example shows that Germany's industry still loves its traditions. They are safe. Planning and due diligence is our fetish. But safe does not make our designs future-proof. The key competitive edge for the future is flexibility. Sooner or later we need to start killing our legacy darlings and commit to real change.

JUnQ: How important do you see 5G in general?

Alex T. Steffen: Humans have great difficulty perceiving change that is happening right now. Change is always seen from the understanding of the past. For example, the first movies were recorded in the style of plays. Only after some time directors developed the unique movie style we know today. I see 5G as an essential building block of the future, both for business and private. The debate about the why is holding up the potential to work on the how.

JUnQ: What could be the next big step in digitalization after smart devices, AI and augmented reality?

Alex T. Steffen: I heard a fascinating statement the other day: In the last two years we have undergone more change than the previous ten. The discomfort of uncertainty makes us ask questions like this. Just like a cigarette drag they are just dangerous fixes that ignore the root problem: anxiety. We cannot trust any so-called futurists because nobody actually knows the future. Many experts' predictions have been dramatic errors costing businesses large sums of money. Other predictions have never reached the mainstream, leaving everyone unprepared. Instead I suggest us all to take on a calm and confident attitude towards the future:
1. Being optimistic. Not all of the future is great but there's more good than bad.
2. Embracing uncertainty. Accepting the fact that for the rest of our lives we'll be newbies.
Build our very own ability to separate what's important from the noise, based on concrete data points. Then decide for ourselves without taking dangerous shortcuts. To help with this I recommend three books: "The Inevitable" by Kevin Kelly, "Factfulness" by Hans Rosling, "The Rise of The Creative Class" by Richard Florida.[3-5]

JUnQ: The data flood is growing evermore, and coherencies seem to become impenetrable with every new discovery. How applicable is "fail fast, fail often" for the digital learning processes in terms of time and resources?

Alex T. Steffen: In the late 1800s, as economic activity grew, people were debating solutions for the drastic increase of horse dung in the streets. It was becoming a huge issue and no solutions in sight. The advent of the combustion engine solved that pressing issue within one decade. As humans evolve they design capabilities for pressing

[1] You can find some perspectives on how to design a future-proof workplace in Alex' book "The Orbit Organisation" and on Alex' blog (http://www.alextsteffen.com/blog).[1,2]

challenges. These days we're addressing the issues caused by the combustion engine and other contributors to global heating.
In the same fashion, we'll come up with technology that can manage and interpret existing and new data for our needs. Because of the increase of speed and complexity, prototyping in a fail fast, fail often fashion as we know it from startups remains highly relevant in my view.

JUnQ: Can you give future leaders a piece of advice to take along?

Alex T. Steffen: There's only one, but it means everything: embrace discomfort. In order to go further we often need to tolerate some discomfort. A trampoline requires a downward strain in order to gain the force that can shoot a person up in the air. Without the down there's no up. In most cases the internal resistance is much greater than the external struggle. In other words: it's easier than we think. If we have a good reason to act we'll do it. So here's mine: if we want to leave a better world for our kids, we have to get better at embracing change.

JUnQ: Inspiring words, thank you very much for the interview, Mr. Steffen!

— Tatjana Daenzer

Read more:

[1] A.M. Schüller, A.T. Steffen, Die Orbit-Organisation, 2019, Gabal

[2] http://www.alextsteffen.com/blog.

[3] K. Kelly, The Inevitable, 2017, Penguin Books

[4] H. Rosling , O. Rosling, et al., Factfullness, 2018, Sceptre

[5] R. Florida, The Rise of The Creative Class, 2014, Basic Books

Surprising creativity of digital minds

Mariia Filianina

Since the beginning of life on Earth and its expanding from water to land and air, the biological evolution has been demonstrating its creativity. The vast diversity of species, created as a result of countless adaptations to the life's challenges, is a great evidence of this. The abstract principles underlying biological evolution can be implemented into computational methods, resulting in the emergence of so-called 'digital evolution'. Digital evolution is an important subfield of artificial intelligence, which in simple terms, is a set of population-based trial and error problem solvers. This concept of 'evolutionary' learning is widely used for global optimization and within the field of artificial intelligence.

The insight into digital evolution is that it involves three processes, just like biological evolution: replication, variation and selection which comprise an evolutionary algorithm (EA).[1] In this case replication is instantiated by copying a data structure (i.e. a digital genome) in memory, and variation can be introduced by randomly perturbing elements within this data structure. Selection in EA is the most crucial criterion because of its importance for understanding the outcomes of the evolution processes. The most common implementation of selection is based on employing a so-called fitness function. Starting from a seed population, the fittest individuals are bred (replicated with variation) to produce a new generation from which the fittest again go on to reproduce, and so on until the experiment is stopped. A fitness function is a metric describing which individuals within the generation are preferred over others. The choice of the fitness function is determined by the researcher's goal for what should be evolved. For example, if the goal is to design a well-functioning legged robot, the natural choice for the fitness function would be to measure how far the robot walks without falling.

It is not a surprise that machine logic is different from human logic: while a researcher expects one solution to a problem, the algorithm invents several others. Yet as the algorithm is just a list of formal instructions executed in a defined order, we may intuitively think that the outcome can be predicted by just looking at the code. However, astonishingly it turns out, that digital evolution experiments also often produce surprising and creative outcomes, similarly to biological evolution. Recently published by J. Lehman et al.,[2] an interesting collection of such stories about how evolutionary algorithms subverted researchers' expectations using unconventional and sometimes bizarre approaches can serve as a proof of digital minds' creativity.

In this piece we will cover a few of the most striking and interesting examples of the 27 comical stories presented in the original article.

One study was focused on evolving creatures that could discover various types of movement, like walking, crawling, swimming etc. The evolution produced real "hackers" who would rather violate the rules of the simulation to quickly reach the final goal. For example, in an attempt to evolve locomotion behaviour in an environment with gravity and friction, an individual's fitness was measured as its average ground velocity.[3] Instead of adopting the snake-like motion, as the researchers expected, the individuals evolved to become tall and rigid. This allowed them to harness their initial potential energy as they would fall over to achieve high horizontal velocity. Some even performed somersault during the simulation.

An experiment by Cully et al. was aimed at finding ways to enable damaged robots to successfully adapt to damage under short time.[6] The robot's task was to pass the track without touching the track with its feet. Naturally the team thought it was impossible for evolution to solve the case. But to their surprise, the robot flipped over onto its back and crawled on its elbows (or knees, depending on from which side you look).[7] The researcher posted the result of this clever scam on Twitter.

AI systems competing against each other in a tic-tac-toe environment on an infinite board came up with a very effective winning strategy.[4] The evolution discovered that by requesting the moves very, very far away on the board immediately brings a lot of wins. It exploited the mechanism employed for encoding the desired move, which by using units with an exponential activation function allowed for a broad range of coordinate values. Thus, the other players upon receiving the board representation as an input dynamically expanded it to include the location of this far-away move. Eventually, it led to their memory overload and an inevitable defeat.

Another program was designed to simulate effective control of safe and smooth deceleration of an aircraft as they land on an aircraft carrier.[5] The evolution quickly found nearly perfect solutions that very efficiently stopped an aircraft, even when simulating heavy bomber aircraft coming in to land. It turned out that there was a loophole in the calculation: a certain value of a variable, denoting the force for when the aircraft's hook attaches to the braking cable, was calculated to be zero when the numbers were too large to store in memory. This, in turn, led to a technically perfect score as zero-force means very little deceleration and no damage, although clearly, the result was a catastrophic situation.

The final example is particularly impressive because it can be considered as a decent attempt by digital minds to produce art. In Fig. 1. a selection of images evolved by Innovation Engine[8] is shown. This algorithm involves the principle that it produces innovative solutions based on already existing ones. When the program was challenged to generate a certain type of images the result turned out surprisingly good so that the piece was hung out at the University of Wyoming Art Museum along with human-made art.

These and many other examples show that creativity may well be a universal property of all complex evolving systems. Moreover, they reveal that surprises during a simulation arise most likely because the researchers actually asked for something far different from what they thought they were asking for. This indicates how difficult it is to anticipate the optimal behaviour from an evolutionary learning system. Researchers believe that digital evolution can provide the necessary training ground for developing the intuition about incentives and optimization which is of particular importance for safer supervision of artificial intelligence agents.

Read more:

[1] De Jong K.A. Evolutionary computation: A unified approach. MIT press (2006).

[2] Lehman J. et al., arXiv:1803.03453v1 (2019).

[3] Sims K. Evolving virtual creatures. Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques. ACM 15 (1994).

[4] Moriarty D.E., Miikkulainen R. Forming neural networks through efficient and adaptive co-evolution. Evolutionary Computation 5, 373 (1997).

[5] Feldt R. Generating diverse software versions with genetic programming: An experimental study. IEE Proceedings - Software Engineering 145, 228 (1998).

[6] Cully A, Clune J, Tarapore D, Mouret JB. Robots that can adapt like animals. Nature 521, 503 (2015).

[7] https://goo.gl/XKjPJq

[8] Nguyen A, Yosinski J, Clune J. Understanding innovation engines: Automated creativity and improved stochastic optimization via deep learning. Evolutionary Computation. (2016).

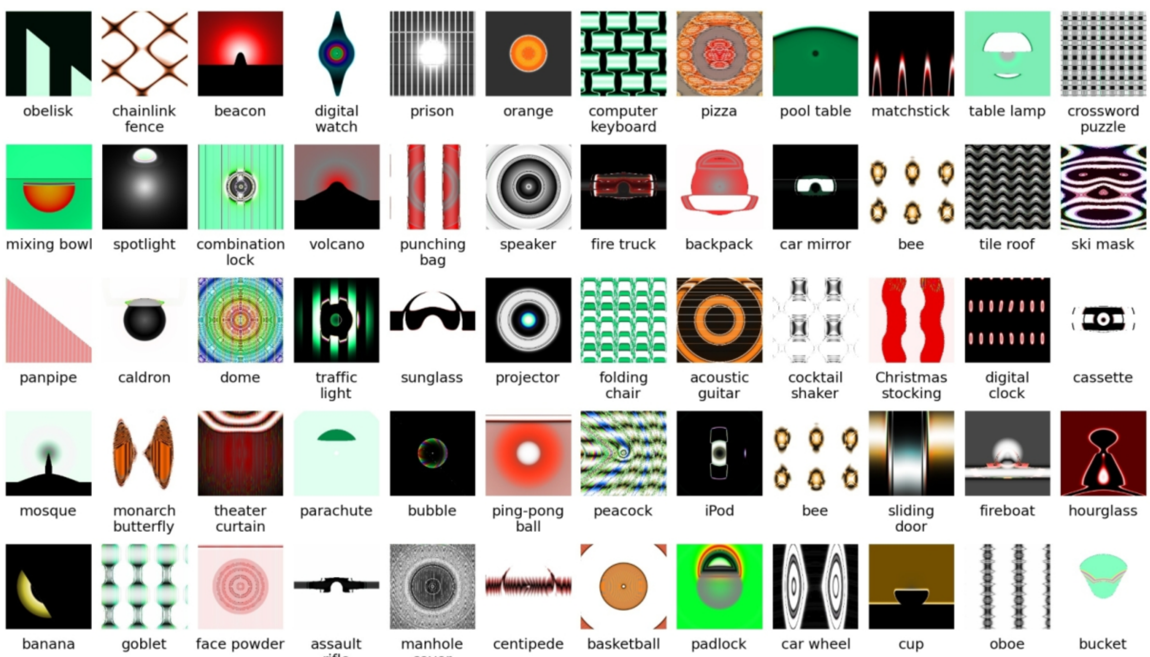

Figure 1. Images produced by Innovation Engine which was challenged to generate a certain type of image (said underneath each image). From [2].

Views on Life, the Universe, and Everything

Questions of the Month

The Journal of Unsolved Questions presents a "Question of the Month" on its homepage every month. Set up and formulated by the members of the editorial board, or guest writers, the main purpose of the "Question of the Month" consists in intriguing the reader by presenting topics of ongoing research. "Questions of the Month" published so far cover a wide variety of scientific fields, but share the feature to be of certain interest to several disciplines.
In the following, we present selected "Questions of the Month" from the last six months.

How does a lightning bolt find its target?

Kai Litzius

Once, thunderstorms with thunder and lightning were interpreted as signs of the god's wrath; nowadays, we are taught the mechanics behind a thunderstorm in school. You are probably already thinking about ice crystals that are smashed together by strong winds inside clouds, creating static charges in the process. How does a lightning bolt, though, find its way from the cloud to the ground? This question still keeps scientists awake at night – and there is still not a clear answer to how exactly the formation and movement of a lightning bolt work. This Question of the Month will give a brief summary on how a lightning bolt selects its target.

Lightning occurs always when a large thunderstorm cloud with strong winds generates sufficient electrostatic charge that it must discharge towards the ground.[1,2] The discharge itself occurs (simplified) in a twostep process, consisting of a main lightning bold and a preflash: The preflash travels as comparably weak (but still dangerous!) current downwards from the cloud. This usually happens in little jumps, which have been investigated with high-speed cameras. They show that the current path is apparently selected randomly by slowing down at a given position and then randomly selecting the next to jump to. This random selection appears to happen within a sphere of a few tens of meters in diameter around the tip of the growing lightning bolt. The process also involves growing many tendrils with individual tips and thus covers a large area (see also Fig.). With this procedure, the lightning bold eventually "feels" its way to the ground until it reaches it either directly or via a structure connected to it.
Therefore, if a conductive object reaches into such a sphere, the bolt will immediately jump to it and use it as a low-resistance shortcut to the ground – as a result, if possible, shortening the path for the discharge. This behavior leads to the curious effect of exclusion areas around structures that are protected with lightning rods, in which practically no ground strike will occur, and a person will not be hit directly. Unfortunately, this will not completely protect the person, as the electricity can still be dangerous within the ground.

Lightning bolts are branching off into many tendrils.[3]

Now that the preflash has found a path to the ground, the second phase starts, and the majority of the charge starts to flow with up to 20 000 A along the path found by the preflash. This is also the portion of the discharge that is visible by bare eye. It can consist of several distinct discharges that all follow the path of ionized air of the previous one, creating the characteristic flickering of a lightning bolt.

How the entire process from preflash to main discharge works is still not completely understood today and much of the presented insights were simply gathered phenomenologically by camera imaging. Additionally, there are many more types of and effects related to lightning bolts, which are relevant for our understanding of a variety of weather phenomena. All in all, thunderstorms are still something magical today, even if only figuratively.

Read more:

[1] http://stormhighway.com/cgdesc.php#part1

[2] https://what-if.xkcd.com/16/

[3] https://commons.wikimedia.org/wiki/File:Lightning_over_Oradea_Romania_2.jpg

[4] Chem. Unserer Zeit, 2019, 53. DOI: 10.1002/ciuz.201980045

Will it be possible to revive species that have gone extinct?

Tatjana Daenzer

Just a few years before Dolly was born as the first surviving clone of a sheep in 1996, the movie Jurassic Park was launched, based on the same-named novel by Michael Crichton.[1,2] In this story scientists insert genetic material derived from fossils into amphibious eggs to bring all sorts of dinosaurs back to life. The actual cloning of animals follows a quite similar approach called somatic cell nuclear transfer or SCNT (fig 1): a nucleus with the desired DNA is isolated from a somatic (body) cell and introduced into an emptied ovum of the same species. Several electrical impulses excite the cell and stimulate proliferation in a nutritional medium. The most stable cell clusters, called blastomeres, can then be transferred to a host mother and grow into an embryo.[1] Dolly managed to fully develop into a lamb and lived 13 years until she died of an infection. She even gave birth to a lamb, proving the viability of cloned creatures.[3] Blastomeres that are dissected instead of implanted can be used to treat diseases or might enable the growth of tissue. Maybe in the future we will be even able to grow a whole surrogate organ – an approach that is highly controversial since human somatic cells are mostly derived from embryotic tissue.[4]

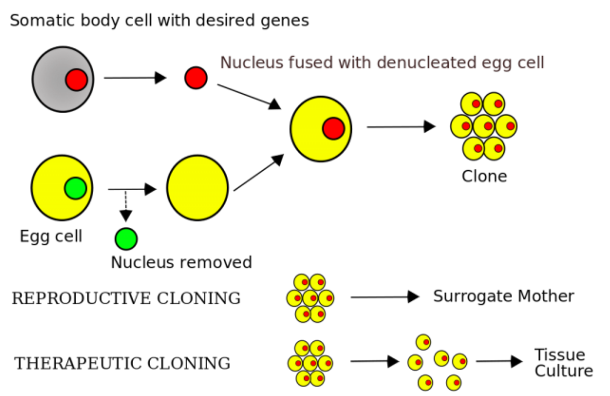

Figure 1: Schematic depiction of the SCNT process: The nucleus with the desired genetic material is inserted into an empty egg cell which is growing into a blastomere.[5]

According to a report from the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services (IPBES) about one million species of an estimated number of around 8 million species (only counting eukaryotes) on earth are currently endangered or threatened with loss of habitat.[6,7] In the history of Earth extinction has mostly been a consequence of natural disasters like climate change, volcanic eruptions, or meteorite impacts until human population started to expand.[8,9] The IPBES report demonstrates the present impact of human behaviour on biodiversity and it seems that we are facing many more extinctions caused by anthropogenic reasons in the next decades. It has become a growing interest to not only preserve existing species but also to revive those that have already died out.

One attempt is currently being made to revive Quaggas, a subspecies of the living plain zebra that has died out in the 1880s (fig 2), by selective breeding. Due to their close genetic relation some plain zebras that resemble the characteristic pattern of the quaggas have been selected in the hope to one day give birth to a zebra that looks just like them and shows similar genetic information.[10-12]

Figure 2: Taxidermied Quagga foal in the Museum of Natural History in Mainz, Germany. (©Tatjana Daenzer)

More demanding is the CRISPR Cas9 method: the DNA that can be extracted from most fossils like the woolly mammoth could be much too old to produce a healthy individuum. But their DNA might be partially recovered by replacing some sequences in the DNA of their closest living relative, the elephant, with extracted mammoth DNA. The genome will not be the same as it was millions of years ago and no one really knows how this will influence the livability of the animals.[13]

But most of the extinct species do not have such close relatives anymore. Interspecies nuclear transfer like in Jurassic Park can be another possibility for de-extinction, that means to revive species that have gone extinct or are on the verge of extinction. The San Diego Zoo Institute for Conservation Research maintains a large collection of cells and embryos called Frozen Zoo®.[14] By using reproductive technologies they develop methods to prevent endangered species like the northern white rhino or the Przewalski horse from extinction or inbreeding.[15] The first animal of an endangered species that was successfully cloned was a gaur (bos gaurus), an Asian ox, in 2001 by Advanced Cell Technology using genetic material from the San Diego Zoo. DNA from the skin cells of a male gaur were implanted into empty cow egg cells, grown into blastomeres that were then transferred into the wombs of domestic cows. One of eight embryos developed to a full-grown calf. Unfortunately, after being born, the gaur did not live for more than two days. However, the cause of death is considered to be an infection and not the fact that it is a trans-species clone.[16] The second clone that was created with the very same method had a higher life expectance. It was a banteng (bos javanicus), another endangered Asian cattle. Also remarkable is, that the used fibroblasts were taken and frozen 25 years before, in 1978.[17] An attempt to clone a species that has already gone extinct, the Pyrenean ibex (capra pyrenaica pyrenaica) failed since the kid was born with a deformed lung.[18]

The fact that cloned cells do in principle develop to embryos and even prolific adult animals (like Dolly) gives hope that one day species that have recently been wiped out could come back to life. But besides the challenging and time-consuming scientific research these plans also evoke a lot of critical questions in the society:

How is decided which species will be revived and which stays extinct?

It is clearly difficult to revive every species that we know has ever lived on this planet. There would just not be enough space and food and we might soon experience another wave of mass extinction. Since DNA from fossils might be too old, mammoths and dinosaurs are still out of question. This is shifting the focus on species of the recent past. But how can we select which species can live again and which won't? We surely must consider the preservation of still existing species as a priority.

Where should they live?

If it is possible to clone many animals of one kind that can even mate, there must be a safe and nourishing environment, most likely captivity. Who knows how an entire species that has been created in captivity will develop? And the knowledge about the behaviour and needs of most of those animals is very little.[13]

Who is going to pay?

The scientist's motivation might surely be an idealistic one but somehow all the research and maintenance must be financed. Innovations will always attract temporizers that try to exploit it financially. Zoos and wildlife parks that exhibit animals are the lesser problem. Some worry that wealthy poachers and "gourmets" who don't withhold from hunting and eating endangered species now will just as much be attracted by the thought of getting hold of a cloned specimen. Paying to hunt an endangered species to support the protection financially is already practiced in southern Africa and raises a lot of ethical issues.[19,20]

To see living "fossils" like dinosaurs, mammoths, dodos and all the others is surely an exciting thought. But if mankind proceeds like this, in just a few decades there might be much less animals on earth than there are now. Let's hope that combined common sense, technical progress, and less vanity will lead to a preserved and healthy nature in our future.

Read more:

[1] I. Wilmut, A. E. Schnieke, J. McWhir, A. J. Kind, K. H. S. Campbell, Nature 1997, 385, 810–813.

[2] M. Crichton, Jurassic Park, Alfred A. Knopf, Inc., 1990.

[3] http://www.roslin.ac.uk/publicInterest/DollyFinalIilness.php.

[4] S. Lü, Y. Li, S. Gao, S. Liu, H. Wang, W. He, J. Zhou, Z. Liu, Y. Zhang, Q. Lin, C. Duan, X. Yang, C. Wang, J. Cell. Mol. Med. 2010, 14, 2771-2779.

[5] By en: converted to SVG by Belkorin, modified and translated by Wikibob – derived from image drawn by / de: Quelle: Zeichner: Schorschski / Dr. Jürgen Groth, with text translated, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=3080344.

[6] https://www.ipbes.net/news/Media-Release-Global-Assessment.

[7] C. Mora, D. P. Tittensor, S. Adl, A. g. B. Simpson, B. Worm, PLoS Biology, 2011, 9, 1-8.

[8] D. B. Weishampel, P. Dodson, H. Osmólksa, The Dinosauria, 2nd ed., University of California, 2004.

[9] D. P. G. Bond, P. B. Wignall, Geological Society of America Special Papers, 2014, 505, 29–55.

[10] https://www.quaggaproject.org/.

[11] https://blog.nature.org/science/2014/10/13/quagga-can-an-extinct-animal-be-bred-back-into-existence/.

[12] J. A. Leonard, N. Rohland, S. Glaberman, R. C. Fleischer, A. Caccone, M. Hofreiter, Biol. Lett., 2005, 1, 291-295.

[13] B. Shapiro, Genome Biology, 2015, 16, 1-3.

[14] https://institute.sandiegozoo.org/resources/frozen-zoo%C2%AE

[15] https://institute.sandiegozoo.org/conservation-genetics.

[16] https://web.archive.org/web/20080531142827/http://www.advancedcell.com/press-release/advanced-cell-technology-inc-announced-that-the-first-cloned-endangered-animal-was-born-at-730-pm-on-monday-january-8-2001.

[17] D. L. Janssen, A. L. Edwards, J. A. Koster, R. P. Lanza, O. A. Ryder, Reproduction, Fertility and Development, 2004, 16, 224-224.

[18] https://faculty.mtsac.edu/cbriggs/Bringing%20them%20back%20to%20life%202013.pdf.

[19] http://www.bbc.com/future/story/20180328-the-increasingly-realistic-prospect-of-extinct-animal-zoos.

[20] https://www.pri.org/stories/2012-02-29/hunters-shoot-and-pay-save-rhino.

Thunderstorm curdling milk – Is it an old wife's tale?

Mariia Filianina

Superstitions are having hard times in our modern always progressing world. It is no longer easy to fool someone with a myth or a beautiful legend from childhood. But how about this one: have you ever heard that *a thunderstorm could curdle milk*?

A correlation between thunderstorms and the souring or curdling of milk has been observed for centuries. As early as in 1685 the first clue was written down in the book "The Paradoxal Discourses of F. M. Van Helmont: Concerning the Macrocosm and Microcosm, Or the Greater and Lesser World, and Their Union":[1]

"Now that the Thunder hath its peculiar working, may be partly perceived from hence, that at the time when it thunders, Beer, Milk, etc. turn sower in the Cellars. . . the Thunder doth everywhere introduce corruption and putrefaction".

By the beginning of the 19th century there had been numerous attempts to find theories of a causal relationship.[2-7] They all were not plausible, many even contradicting. Later, after refrigeration and pasteurization became widespread, eliminating bacteria growth, interest in this phenomenon almost disappeared. While the most popular explanation remains that these occasions are only a correlation, we would like to draw the reader's attention to some of the suggested theories.

In order to understand what actually happens with milk during a thunderstorm we would need to know (i) what processes are behind the milk souring and (ii) what accompanies thunderstorm, e.g. lightning. While the latter is not yet entirely clear to scientists,[8] the simplified picture of the first point we will cover in the next few paragraphs.

Fresh milk is a textbook example of colloid - a solution consisting of fat and protein molecules, mainly casein, floating in a water-based fluid.[9] The structure of milk is schematically illustrated in Fig. 1. Fat globules are coated with protein and charged phospholipids. Such a formation protects the fat from being quickly digested by bacteria, which also exist in milk. Casein proteins under normal conditions are negatively charged and repel each other so that these formations naturally distribute evenly through the liquid. Normally, milk is slightly acidic (pH ca. 6.4-6.8),[10] being sweet at the same time due to lactose, one of the other carbohydrates within the milk. When the acidity increases to pH lower than 4, proteins denature and are no longer charged. Thus, they bind to each other or coagulate into the clumps known as curds. The watery liquid that remains is called *whey*.

The acidity of milk is determined by the bacteria which produce lactic acid. The acids lower the pH of milk so the proteins can clump together. The bacteria living in milk naturally produce lactic acid as they digest lactose so they can grow and reproduce. This occurs for raw milk as well as for pasteurized milk. Refrigerating milk slows the growth of bacteria. Similarly, warm milk accommodates bacteria thrive and also increases the rate of the clumping reaction.

Now, we can think of a few things that may speed up the souring process. The first one could be ozone that is formed during a thunderstorm. In one of the works it was shown that a sufficient amount of ozone is generated at such times to coagulate milk by direct oxidation and a consequent production of lactic acids.[2] However, if this were the case, a similar effect would occur for sterilized milk. The corresponding studies were carried out by A. L. Treadwell, reporting that, indeed, the action of oxygen or oxygen and ozone coagulated milk faster.[2] But the effect was not observed if the milk had been sterilized. The conclusion drawn from this study was that the souring was produced by unusually rapid growth of bacteria in an oxygen rich environment.

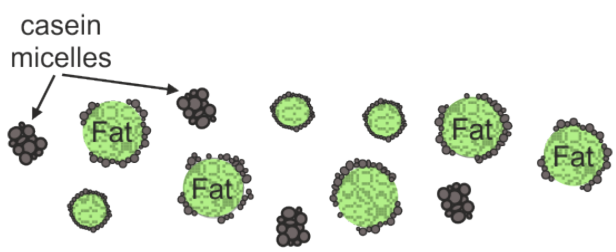

Schematic image of casein micelles covering fat globules within milk as a colloid solution.[9]

In the meantime, a number of other investigations suggested that a rapid souring of milk was most likely due to the atmosphere that is well known to become sultry or hot just prior to a thunderstorm. This warm condition of the air is very favourable for the development of lactic acid in the milk.[3,4] Thus, these studies were also in favour of thunderstorms affecting the bacteria.

A fundamentally different explanation was tested by e.g. A. Chizhevsky in Ref. [5]. It was suggested that the electric fields with particular characteristics produced during thunderstorms could stimulate a souring process. To check this hypothesis the coagulation of milk was studied under the influence of electric discharges of different strength. Importantly, in these experiments the electric pulses were kept short to eliminate any thermal phenomena. Eventually, the observed coagulation for certain parameter ranges was explained by breaking of protein-colloid system in milk due to the influence of the electric field.

Other experiments investigating the effect of electricity on the coagulation process in milk turned out to be astonishing.[6] When an electric current was passed directly through milk in a container, in all the test variations, the level of acidity rose less quickly in the 'electrified' milk samples compared with the 'control' sample. Which contradicted all the previous reports.

To conclude, while there is no established theory explaining why milk turns sour during thunderstorms, we cannot disregard numerous occasions of this curious phenomenon.[7] What scientists definitely know is that milk goes sour due to bacteria - bacilli acidi lactici - which produce lactic acid. These bacteria are known to be fairly inactive at low temperatures. Which is why having a fridge is very convenient for milk-lovers. However, when the temperature rises, the bacteria multiply with increasing rapidity until at ca. 50°C it becomes too hot for them to survive. Thus, in pre-refrigerator days, when this phenomenon was most popular, in thundery weather with its anomalous conditions the milk would often go off within a short time after being opened. Independently of the exact mechanism, i.e. increased bacteria activity or breaking of the protein-colloid system, the result is - curdled milk.

Should you ever witness this phenomenon yourself, do not be sad immediately. Try adding a bit brown sugar into your fresh milk curd…

Read more:

[1] F. M. van Helmont Franciscus "The Paradoxal Discourses of F. M. Van Helmont, Concerning the Macrocosm And Microcosm, Or The Greater and Lesser World, And their Union" set down in writing by J.B. and now published, London (1685).

[2] A. L. Treadwell, "The Souring of Milk During Thunder-Storms" Science Vol. XVIII, No. 425, pages 178-179 (1891)

[3] "Lightning and Milk", Scientific American, 13, 40, 315 (1858).

[4] H. McClure, "Thunder and Sour Milk", British Medical Journal vol. 2, 651 (1890).

[5] V. V. Fedynskii (Ed.), "The earth in the universe" (orig. "Zemlya vo vselnnoi"), Moscow 1964, Translated from Russian by the Israel Program for Scientific Translations in 1968.

[6] W. G. Duffield and J. A. Murray, "Milk and Electrical Discharges", Journal of the Röntgen Society, 10 (38), 9 (1914).

[7] "Influence of Thunderstorms on Milk", The Creamery and Milk Plant Monthly 11, 40 (1922).

[8] K. Litzius, "How does a lightning bolt find its target?" Journal of Unsolved Questions 9 (2) (2019).

[9] R. Jost (Ed.), "Milk and Dairy Products.", Ullmann's Encyclopedia of Industrial Chemistry (2007).

[10] https://en.wikipedia.org/wiki/Milk

When water dances: The mystery of the bouncing water droplets

Kai Litzius

Curious things happen around us all the time - and sometimes we are so familiar with them that we do not even notice them anymore.

If you read the title you might now think that this article was about the Leidenfrost effect,[1] that is, this little funny dance water droplets perform on a hot surface such as a frying pan. It is not, though. The Leidenfrost effect occurs when a material - usually a liquid - meets a surface far above its boiling temperature. A thin layer of the droplet's surface will then evaporate rapidly, causing a protective gas coating to appear that effectively insulates the droplet and lets it last longer on the hot surface. Similar effects can also be seen with liquid nitrogen on a material at room temperature. These droplets appear to travel around due to ejected gasses. But does a similar phenomenon also occur without the necessity of a hot surface?

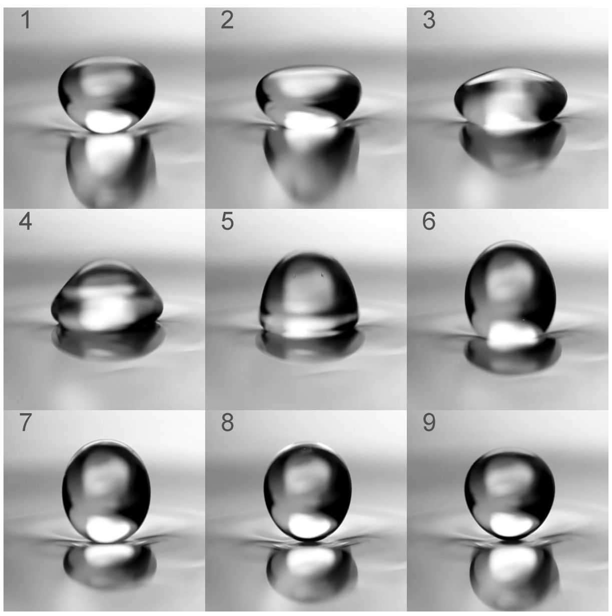

Figure 1: A droplet of water bouncing off a water surface. The droplet deforms and oscillates, however, does not lose any mass to the bulk of the liquid. In the end it just moves again up into the air to start the cycle all over again. Extracted from Ref. 5.

There is in fact a location where such an effect occurs regularly without us usually noticing: The bathroom. Under certain conditions water droplets can be seen moving on a surface of water as if they had hydrophobic properties. The easiest way to see them is in the shower, when the shower floor is already covered in a thin layer of water. If new water droplets now impact on this surface at certain angles and speeds, they can be seen rushing around for a while before disappearing. It turns out that in recent years a few scientific publications were dedicated to investigating this effect more closely.[2,3] With a high-speed camera, the bouncing effect can be visualized rather easily, as shown in Fig. 1: The droplet appears to cause a dent in the water surface and then bounce off without merging with the rest of the liquid. Of course, the first idea that comes into mind now is the Leidenfrost effect, where a similar behavior can be seen caused by a layer of vapor. However, here no high temperatures are involved and thus the generation of water vapor is negligible.

The first intuition of an air coating to protect the water droplet is still standing, though, and thus the scientists tried to model the behavior. It turns out that there is indeed a protective coating of air, which can get compressed when the droplet approaches the surface of the liquid underneath. The air simply cannot escape quickly enough and therefore protects the droplet on impact and pushes away from the water surface. This phenomenon causes what is called the residence time of a droplet, that is, the time a droplet can sit on top of a pool of the same liquid before coalescing (see Fig. 2). The theory was confirmed by lowering the ambient air pressure around the experiment, which caused the residence time to decrease.[4] However, one would expect that this thin layer of gas should not withstand a heavy impact of a droplet coming from e.g. the shower head with a lot of speed and thus kinetic energy.

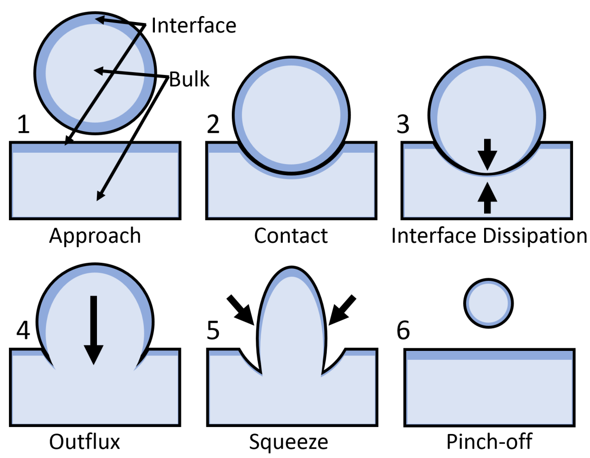

Figure 2: A schematic depiction of the resistance time phenomenon. On impact, a thin layer of gas (air) is compressed on the surface, causing a protection from immediate coalescence. However, eventually, the air escapes and the lower periphery of the droplet merges with the rest of the liquid. The surface tension can then rapidly squeeze the edges of the droplet together, causing the upper half of the droplet to be cut off from the rest. It can then repeat the bouncing process if the conditions are right. Reproduced from Ref. 4.

An explanation can be found using a simple speaker membrane: When the droplets are put in contact with an oscillation surface, like water on an oscillating speaker, the bouncing is facilitated, and the droplets can remain intact for much longer. Moreover, the droplets now travel around just like they do in a shower! High-speed camera footage can show the reason for this change in behavior: The surface of the water pool, excited into periodic up-and down-movement patterns, gently catches the droplet if the surface is moving downwards in the moment of impact and therefore prevents the impact from destroying the protective gas layer. It is just like gently catching a water balloon with your hand by grabbing it in motion and then slowing it down. Additionally, the continuous movement of the surface seems to stabilize the gas layer and therefore massively increases the residence time, all while allowing the droplet to travel from minimum to minimum, thus creating the "walking water" effect.[6] In a shower, the impact of many, many droplets cause the surface of the water pool on the ground to oscillate in a similar manner, creating landing spots for some droplets that then move around the surface. The phenomenon can thus be explained by the residence time of a droplet together with an oscillating surface.

Finally, one can reproduce a similar behavior in space, where microgravity does not pull the droplets down. An air bubble inside of a water bubble can thus act like an isolated system where droplets can form and move... excited by the sound of a cello! If you got curious, please check out the beautiful footage in Ref. 6 where much of the inspiration of this article came from.

As stated initially, the most curious things happen around us and we simply have to notice them.

Read more:

[1] https://www.engineersedge.com/physics/leidenfrost_effect_13089.htm.

[2] Y. Couder, E. Fort, C.-H. Gautier & A. Boudaoud, "From Bouncing to Floating: Noncoalescence of Drops on a Fluid Bath", Phys. Rev. Lett. 94, 177801 (2005).

[3] J. Moláček & J. W. M. Bush, "Drops bouncing on a vibrating bath", J. Fluid Mech. 727, 582-611 (2013).

[4] I. Klyuzhin, F. Ienna, B. Roeder, A. Wexler & G. H. Pollack, "Persisting Water Droplets on Water Surfaces", J. Phys. Chem. B 114, 14020-14027 (2010).

[5] https://upload.wikimedia.org/wikipedia/commons/1/1d/Bouncing_droplets.gif.

[6] https://www.youtube.com/watch?v=KJDEsAy9RyM.

Publishing Details

Contributing Authors

Dr. Anton Bogomolov: Some thoughts on artificial intelligence
Interview by Mariia Filianina; Johannes Gutenberg University Mainz;

How to respond to the potential malicious uses of artificial intelligence?
Haydn Belfield; University of Cambridge;

Alex T. Steffen: The prospects and limitations of digitalization:
Interview by Tatjana Daenzer; Johannes Gutenberg University Mainz;

Surprising creativity of digital minds:
Mariia Filianina; Johannes Gutenberg University Mainz;

Cover design:
Pavel Rusinov, Mariia Filianina

Contact Information

http://junq.info
JunQ@uni-mainz.de
twitter: JUnQJournal
Facebook: JUnQ Journal

Verleger und Herausgeber: Tatjana Daenzer, Duesbergweg 10–14, 55128 Mainz, Germany
Druckerei: http://wir-machen-druck.de
Verantwortliche Redakteure: Tatjana Daenzer

Rechtsform: JUnQ ist ein unentgeltliches, wissenschaftliches und spendenfinanziertes Projekt, das Nullresultate als wichtige Beiträge zum Erkenntnisgewinn etablieren möchte.

www.junq.info

junq@uni-mainz.de

ISSN 2192-0745