



Benelux AI Newsletter

Would you like to announce your own events on the BNVKI website and newsletter? Send an email to events@bnvki.org! Also reports of past events are welcome. If you have other news or contributions for the BNVKI newsletter, these can be sent to editor@bnvki.org.

BNAIC/BENELEARN 2021 will be held in Luxembourg

The 2021 edition of BNAIC and BENELERAN will be organized by the University of Luxembourg and the Luxembourg Institute of Science and Technology on the 10-12th of November. The event will take place at the new Belval Campus, called Cité des Sciences, in the southern part of the country. The general chairs of the conference are Thibaud Latour (LIST) and Leon van der Torre (UniLu), the chair of the program committee is Luis Leiva (UniLu). Next to the general topics of AI and ML, this year the organizers will invite submission for four special tracks: AI&Art (chaired by Christoph Schommer), AI&Ethics (chaired by Amro Najjar), AI&Law (chaired by Réka Markovich), and AI&Systems (chaired by Cédric Pruski).



Belval Campus (picture by Giovanni Casini)

[Read more](#)

AiNed investment programme

The AiNed investment programme for artificial intelligence run by the Netherlands AI Coalition (NL AIC) is getting 276 million euros as one of the innovation proposals that the National Growth Fund will be backing. The decision that the advisory committee chaired by Jeroen Dijsselbloem has made about the first round of the National Growth Fund was adopted by the council of ministers today. It includes honouring the 5 proposals for research and innovation that Mona Keijzer, the State Secretary for Economic Affairs and Climate Change, submitted on behalf of the collaborating companies, centres of expertise and authority bodies, including the NL AIC. Visit the [NL AIC website](#) for more information.

[Read more](#)

BNVKI Interview with Prof. dr. Bert Kappen

By Emil Rijcken

In 1987 Professor Kappen received his PhD in theoretical physics from Rockefeller University in New York. After having conducted research at Philips for two years, he returned to academia in 1989. In a quest to understand human-kind better, he changed his research area to neural networks and machine learning.

Professor Kappen's research interests lie at the interface between statistical physics, computer science, computational biology, control theory and artificial intelligence. For years, he has argued that computing, memory, and energy consumption will be the main challenges for AI's future growth. Instead of limiting his research by these challenges, professor Kappen recently published a revolutionary article in Nature titled 'An atomic Boltzmann machine capable of self-adaption'. In this work, he has demonstrated the possibility of training a neural network at the nuclear level. This finding may be the precursor of a new generation of algorithms consuming only a fraction of current models' energy absorption. During this interview, we will discuss what professor Kappen found, the potential impact of his findings, and whether we can draw links between his findings and humans.

Professor Kappen, do you primarily consider yourself to be a physicist or a computer scientist?

That is a difficult question. Most of my publications are in machine learning journals and conferences. In that sense, I am a computer scientist. However, the methods and methodologies that I use mostly, stem from physics. So, I consider myself primarily a physicist in the end.

For instance, much of my work has been on approximate inference, which efficiently approximates computations for low-order statistics in very large probability models. This problem is dominated by various techniques such as variational methods and Monte Carlo sampling, all rooted in physics. In much of my work, I transfer these technologies from the physics community to machine learning and AI.

Why did you decide to transition from physics to computer science?

I did my PhD in theoretical high energy particle physics and was always fascinated by the sophistication of the field's mathematics. It has incredible beauty, we have had incredible achievements, and it has been a privilege to learn about it.

At the same time, I felt that it was very remote from my personal daily life.

While pursuing my PhD, I was always very interested in how my brain works. The notion of the 'I' and the question '*who am I?*' has always fascinated me. This question can be answered by knowledge of the brain, neuroscience, AI and philosophy.

My move to artificial neural networks, away from particle physics, was mainly motivated by this fascination. I find this a beautiful field, as you are modelling systems that mimic humans in a certain, limited way. And therefore, it is an attempt to understand yourself.

Does the future of AI lie in the intersection between physics and computer science?

It is hard to predict the future. For instance, I did not see the deep learning success coming, although these techniques have been around since the 80s already. It came as a surprise that we got this spectacular performance using massive amounts of data and computing power. That said, I think there is vast potential in the intersection between these fields. But profits can arise from a variety of other disciplines as well.

Many people claim that we should employ the knowledge of the brain to help AI advance. This is a compelling and persuasive idea but very hard to do in practice. The number of novel brain mechanisms contributing to AI and new technologies is very few, if not none. Still, neuroscience will inspire us to do this research and therefore remains relevant for AI's future, and so will engineering and new insights from physics.

In conclusion, I think there are many roads to success and many significant contributions. But my particular interest lies, of course, in the interface between machine learning and physics.

My recent work is an example of this interface where we let physics tackle massive computational problems. We can identify properties that you essentially get for free, using physics, and use them for our purpose.

Let's talk about your recent work. What were the findings in your Nature paper?

This is a collaboration with Professor Alexander A. Khajetoorians from the Condensed Matter Physics research group in Nijmegen (IMN).

He is one of the world's leaders in manipulating individual atoms on surfaces. Do you know the old pictures in which IBM wrote its name with 35 atoms? He can also do such things. Recently, he has found that certain atoms (cobalt in this case) can be in two states while being on the surface of a substrate of black phosphorus. Essentially, that means that these atoms interact with their substrate; the electron cloud around the atom can be in two configurations as it displaces itself on the atom's surface mechanically. You can think of these atoms as implementing a single stochastic bit in two states, and this '*dance of the atoms*' can be measured as a voltage.

Earlier, IMN has found that the atoms dance with different timescales if you put many of them on the surface, some dance very slow, while others very fast and in a complex pattern. That finding inspired us to make a connection to learning.

The brain can be modelled as a neural network consisting of neurons and connections, which you call synapses. For a given synaptic connectivity in the human brain, the neurons implement certain functionalities such as perception or motor control. Essentially, learning happens when synapses change to other states. While neurons change on a very short timescale, milliseconds, the synapses change much slower. The latter can take minutes, hours, days or much longer (think about the time it takes a baby to walk).

Just like the dancing atoms, you can think of a neural network as a network of two types of variables: some moving very fast and some moving very slow.

In our work, we demonstrate these two types of timescales on this material. The separation of timescales is massive: fast variables change in milliseconds, whereas slow variables may take seconds, minutes or even more to change, just like the brain. This timescale separation can be thought of as implementing a form of memory: the slow variables change over time and change the connectivity between the fast variables in a spin array.

We built this learning process in a simple setup in which we arranged seven atoms.

Three atoms are our fast variables, the neurons, and the four slow variables control the connectivity and thresholds, just like in classical neural networks. Thus, this system can be formalized as a Boltzmann machine.

Why is this finding important?

It shows how you can make physics work for you, and there is a clear difference with traditional neuromorphic computing. In neuromorphic computing, people build a material that acts as a neuron and implement it physically on a CMOS device; the whole network is engineered to behave in the way you want. In contrast, we do not engineer a network with our approach but look at the materials' properties as they come and use them. Our approach has a massive potential to be scaled up.

We only experimented with seven atoms. However, this can be scaled up to larger systems in the semiconductor industry. By doping materials up, you can scale up to millions or hundreds of millions of spins in a tiny piece of material and get a similar '*two-timescale behaviour*'.

However, there are still massive problems ahead of us. We have not solved the I/O problem yet; how to read the signals from the atoms. While there are still many things to do, in principle, this shows that you can use a material that learns by itself and is not engineered to do so. It also enables large-scale applications and is very energy efficient; we perform computations at the atomic scale, approaching the lower limit of energy consumption.

How much energy could be saved with an atomic network in comparison to our current networks?

It is hard to say, and I cannot quote you any numbers. The main bottleneck is the peripheral hardware to convert analogue signals to digital information. As an indication: it takes a thousand computers (each operating on the scale of kilowatts) to solve AlphaGo, which totals a scale of megawatts. On the other side, the brain operates on 25 watts approximately (similar to a lightbulb), which is a massive difference. Our approach's consumption would be lower than the brain absorbs, as the energy used for flipping an atom is very low. But again, the consumption of the peripheral hardware needs to be added to the consumption as well

So it saves massive amounts of energy. Is it also faster than current computing?

No, but the (real) brain is not very fast either. The creation of an action potential in a biological cell has a typical timescale of milliseconds (a thousandth of a second). In modern computer hardware this would not be very fast, as there you have clock

instances going into gigahertz (a billionth of a second). In that sense, the brain is very slow, and our materials are not very fast either. However, our approach will be highly parallelizable, instead of a sequential process to boost the clock cycle. Also, we could vary the timescale for faster computations probably. However, that is currently not our focus as we are working on the separation of timescales instead.

What are other problems that still need to be solved?

So far, we obtained the separations of timescales by rearranging the atoms in a specific way. Our goal is to get timescale separations without having such a particular geometric arrangement. If we find the suitable atoms to do this, we will have much more flexibility as we could just sprinkle the atoms on the substrate. In that case, to get interesting results, it would not matter where they are.

Two other challenges ahead of us are:

1. Scaling up – at some point, we need to embed this in semiconductor devices.
2. I/O problem – how to transfer an input signal as a high-dimensional array to some Fourier basis and interact in the Fourier mode.

Both are still open problems, and I expect it will take a decade before seeing practical device of this kind.

Will more research groups around the globe focus on this subject and what changes can we expect?

We have been unique in what we do so far. I cannot say how many groups will latch on to this, but the initial reactions are very positive. Many people are intrigued, and in the condensed matter community it is a novelty as well.

We are likely to see a change in the interaction between physics and algorithms based on our work.

With neuromorphic computing, the idea is that whatever software idea one has, it will be built into the hardware by others afterwards. The hardware community then focuses on getting the same functionality faster and or cheaper. But because the hardware is adjusted to software requirements, there is not much interest from the theoretical/algorithmic community.

Our work follows a different dynamic. In our case, we assess what properties from physics are at our disposal and how they can serve algorithmic problems. Then, we create new algorithms accordingly. From this change of dynamic, I expect the creation of a whole new set of different algorithms. A similar dynamic will also be encountered in the quantum regime, where we will design algorithms based on the hardware's quantum features.

Would you say this is a transition to a bottom-up approach instead of a top-down one? The algorithms are not leading anymore, but the hardware is leading instead?

Yes, in a way, but it is still top-down as well. To see the limitations of the top-down-approach, let's take the brain as an example. The brain consists of 90% water, so a full top-down approach would require our silicon devices to be made with water as well, which they are not. So clearly, the brain principle cannot be implemented fully, and we need different materials. However, using different hardware may lead to suboptimalities as the brain is optimized for its own type of hardware.

At the same time, some principles from the brain are robust and generalizable. An example is local computing; it is more efficient to perform local computations in which the required variables are there instead of shipping data back and forward. Algorithmically, that is a challenging constraint. Also, stochasticity plays a role as well in practice and there is a theoretical argument that stochastic systems learn better than systems without stochasticity.

In conclusion, top-down principles can help to guide you with the selection of bottom-up approach.

Boltzmann machines are a form of unsupervised learning. Can you also create supervised algorithms at the atomic level?

Yes. Trivially, unsupervised learning can implement supervised learning. You can take $n-1$ of your bits as input, and then the n th bit is the output based on which you train. After training, you clamp the $n-1$ inputs and look at the value of the n th bit.

In our recent work on the *quantum Boltzmann machine*, we demonstrate that this works. Additionally, while working with the quantum variant, you get a functionality that

can learn many more problems than with a traditional Boltzmann machine. For instance, suppose you would like to learn an XOR- or a parity problem (non-linear) and train with $n-1$ inputs, and the n th variable is the output. For these problems, the classical Boltzmann machine only learns a linear separator and therefore, an unsupervised scheme without hidden variables cannot learn the correct parameters. However, when using a quantum Boltzmann machine with only visible units, you can already learn an XOR problem. This possibility illustrates that quantum devices could already perform sophisticated supervised learning with 'only' an unsupervised algorithm.

Let's link humans to computers; can we understand the human brain better from these algorithms?

At this moment, it is still a far reach to understand the human brain from algorithms. During the last thirty years, I have been seeking to answer how human thoughts and emotions can be explained in algorithmic terms and why we love music, art and beauty.

For me, no mathematical description has been satisfying yet. To model music perception for instance, you could describe it as a Hidden Markov model and may obtain a lousy- or a good performance. But regardless of the outcome, it does not address how I find beauty in music as it does not connect to that in any way. These approaches are just a bunch of formulas, and they do not help me to understand what I experience when I see great beauty. We can understand perception in algorithmic terms, but those formulas do not explain our consciousness's personal experience. This is a complex topic, and my conclusion is that no algorithmic method suffices to understand consciousness.

I was greatly inspired by the book *'Zen and the art of motorcycle maintenance'* by Robert M Pirsig. The main character is an English teacher and teaches English writing to a class of students. He is frustrated by their mediocracy and questions the possibility of devising a set of rules to produce a high-quality text if the students follow these rules. He finds that no such rule system exists and cannot be written down. Therefore, he concludes the existence of high-quality phenomena which cannot be captured in rules. As rules are algorithms, certain concepts in nature exist that we cannot capture in algorithmic terms.

One of the problems we see with AI is that it is very mechanical. And logically, if you keep on following rational thinking, you stay within the realm of mechanical thinking, algorithms, and rule-based systems. In that way, you will never understand what

consciousness is (just like how the teacher could not describe quality).

The philosopher Daniel Dennett (and many others) claims that consciousness does not exist and is epiphenomenal. Similarly, many scientists struggle with the phenomenon of consciousness. They work on the premise that whatever they study needs to be described in rational terms; if it cannot be studied in rational terms, it cannot be the subject of a scientific study. And as a result, you could say it does not exist, in the scientific sense.

But obviously, consciousness does exist, and therefore we need to approach it differently.

My interest in quantum computing arose from my efforts to address the question of consciousness. I have no evidence that quantum computing will provide answers. However, algorithmic thinking has its limits, and nothing else will cut it.

I take an open view of whether this is possible. But the world is quantum, so why should I not try?

In contrast to the scientists that I mentioned before, some people link consciousness to quantum processes. These links are fascinating and exciting but also easy to criticize and full of problems. For instance, our quantum computers must operate very close to absolute zero temperature; otherwise, all the quantum effects are destroyed. So how can a quantum process ever work in the brain? The brain is 37 degrees Celsius, wet, and noisy; any quantum process is likely to be killed and will not be measurable.

At the same time, there is some evidence that quantum effects may take a role in biology. Various papers propose the idea that photosynthesis within plants is based on quantum effects. Others suggest that the magnetic compass of migrating birds are associated with quantum effects as well. Therefore, it is not entirely outrageous that quantum effects can be operational at room temperature in biological systems. But clearly, we do not know how that works, and I am very interested in exploring these directions.

How would our understanding of the brain change if we find that quantum processes occur in the brain?

Quantum processes have a distinguishing feature that deviates from classical thinking; the quantum world is holistic and not reductionistic. Consciousness is not reductionistic

either. However, humans are very acquainted with the reductionistic view. To understand how something works, we break it up into components, study the components and understand the interactions. Based on this process, we believe that we understand what it is and what it consists of.

This is a very pervasive approach, and it has shown to be very successful in science. All classical physics is also reductionistic; you have particles, properties, and interactions, which aligns with our classical understanding of the world. But in the quantum world, things are different. You cannot assign specific properties to an individual object because the objects' properties are shared between the objects they interact with. So, if you have two objects, you cannot assign properties to any of them as both are a function of the environment. Therefore, the quantum world cannot be approached through a reductionistic scheme.

There is a certain holism about that, which I find very appealing.

One of the main contributors to quantum mechanics was David Bohm. In his textbook about quantum mechanics, he dedicates an entire section to the analogy between quantum processes and thought processes. He describes that according to quantum theory, you interact with something as soon as you measure it. And as soon as you interact with it, you change the system already. Therefore, you cannot solely observe a quantum system because you change it as soon as you observe it. It changes to a particular state, and after measurement, all you know is the post-measurement state, but it will not tell you what the system's state was before the measurements.

Bohm then draws the analogy with thought processes. If someone asks what you are thinking about, your thoughts change as they shift towards the question. Until then, you were not thinking about the question, and now your thought process is with the post-question thoughts. The book elaborates on how classical thinking relates to logic, and that quantum logic could be a brain process. This section is a fascinating read and is only five pages. To answer your question, these quantum features would be a way to describe our thought processes.

So now we have to demonstrate that quantum processes can take place in the brain?

One way would be to propose quantum mechanics as a function of the brain. Then you need to make it work at 37 degrees in a noisy environment. But instead, you can also do quantum-like processing and look at the probabilistic logic proposed by quantum

mechanics. This logic is different from classical logic, and therefore you get deviations. For example, Bayes' rule does not hold in quantum logic as you get additional terms there. Another approach is to take a pragmatic view and assume the quantum processes are in the brain. In that case, we can take the consequences of that view without focusing on the implementation. By doing so, we can take the quantum probabilistic framework to describe a variety of situations.

One particular example of this is '*quantum cognition*', where people try to use this quantum formalism just as a mathematical description of human beings' irrational behaviour. It is well known that people do not follow Bayes' rule. Maybe quantum correction terms can help to explain the deviations from Bayes' rule in people's behaviour.

If we find ourselves able to quantify consciousness in the future, would that be the last puzzle piece in your quest to understand the human brain?

I think so, but we would have to show that this is true first. This holistic picture of the non-reductionistic view is related to someone's perception of beauty and art. I reject the idea that a classical device can be built that explains how I perceive beauty. But I can imagine that a quantum device could do that with its non-reductionistic, holistic feature.

If we find ourselves doing so successfully, that would be amazing and a real breakthrough. But for now, we do not know whether that is possible, and I think it is worth researching.

We have come to the end of the interview; do you have any last remarks?

Our world, society, and people are highly influenced by the world's scientific view. The western scientific view of the world is reductionistic. Following this reductionistic view has been highly successful and has led to many achievements such as the industrial revolution and all our technologies.

At the same time, it leaves no room for non-material aspects. If the world is material, you are material, and it is hard to escape the conclusion that you and I are nothing else than sophisticated robots. Up to five years ago, I was highly convinced of this idea myself as I have always been a hard-nosed scientist, and still am. But this view implies that any kind of peripheral thinking that cannot be explained in scientific terms does not exist. If you take this scientific view seriously, you may conclude that spiritual beliefs, feelings, or ethnophysiology do not exist as we cannot measure them.

The view that we are just a sophisticated computer, in the end, is a fascinating thought that is affecting all of us implicitly. It affects the way we look at ourselves but also how we look at our society. If I am a machine, and you are a machine, we are all just a bunch of machines. That is a purely materialistic view quickly leading to a purely materialistic society where the things we cannot explain in terms of matter lose importance. This is a subtle effect that classical physics has on our society. It is good to be aware of that, and we should look beyond that.

[Read more](#)

EurAI Dissertation Award

Nominations for the 2021 Artificial Intelligence Dissertation Award can be submitted until 15th May 2021. As in previous years, the award includes €1500 and a certificate signed by the EurAI president.

Instructions on how to apply are provided here:

https://www.eurai.org/apply_for/dissertation_award

[Read more](#)

Student initiative: Talk That Science

[Talk That Science](#) is an initiative of students to narrow the gap between science and society through music. We do this by recording radio shows in which an interview with a scientist is alternated with applicable music. Our goal is to make science popular and accessible for everyone so that society is better informed about developments in this field and the way it affects our daily lives. At the same time, this offers scientists an opportunity to share knowledge and findings with society in an accessible way.

Did you know that COMPUTER SCIENCE can help to design ways to run a democracy? Listen to professor Ulle Endriss to find out why computers can help us when it comes down to making collective decisions.

Talking about EVOLUTIONARY ROBOTICS with professor Guszti Eiben. Explore the implications of evolutionary robots, because if evolution created life, do we create artificial life with artificial evolution? Follow us on our journey to discover how evolution can be implemented in robots, while listening to great songs.

[Read more](#)

BNAIC/BeneLearn 2020

The Belgium Netherlands Artificial Intelligence conference (BNAIC) is an annually organized conference bringing together AI researchers from all over the Benelux. This year it was supposed to be held for the first time in Leiden. Due to the continuing lockdown, it was organized online. Continuing a trend that was started two years ago, BNAIC was organized in conjunction with BeneLearn, the Benelux conference on Machine Learning.

On top of having formal Springer CCIS proceedings, BNAIC and BeneLearn profile themselves as networking conferences, with strong underlying communities. As such, BNAIC/BeneLearn is a good opportunity for young researchers to get in touch with the Artificial Intelligence and Machine Learning community from the Benelux. Like previous years, the majority of papers came from the Netherlands and Belgium, however there were also submissions coming from Luxembourg, Germany, UK, Sweden, Poland, Italy and Cuba, ensuring diversity of countries.

The conference featured three keynote speakers:

- Dr. Joost Batenburg from Leiden University talked about tomography and real-time image processing.
- Tom Schaul from Google DeepMind talked about AlphaStar, a Reinforcement Learning based agent that bested the game Starcraft.
- Dr. Gabriele Gramelsberger from Aachen University talked about the impact that Artificial Intelligence and Machine Learning have on the Natural Sciences.

On top of the three keynote speakers, there were also plenary talks from Dr. Nico Roos

(Maastricht University), Dr. Yingqian Zhang (TU Eindhoven) and Prof. Dr. Luc de Raedt (KU Leuven). These were organized by the Faculty focusing on the Facts of Artificial Intelligence (commonly called FACT talks), which aim to attract high-profile researchers to the conference and intend to be provocative. Both of these goals were achieved, thanks to the speakers.

During four time slots consisting of four parallel sessions, the contributions of the 62 accepted papers were discussed. All authors have created a 10 minute video, that were watched and discussed during these slots. The videos will also be available after the conference. Contributions were grouped by topic. Some popular topics were reinforcement learning, agents and classification.

At the end of the BNAIC, the Business and Society afternoon was organized. It started with a panel discussion from various stakeholders on the AI Rat Race, and continued with plenary talks from Guy de Pauw (Textgain) and David Graus (Randstad Groep Nederland) on a wide range of topics, including how to detect and block hate speech from the internet and how to make algorithms recommend more diverse content.

During the closing ceremony, various awards were handed out:

- The best paper award, sponsored by SNN, was awarded to Andrei C. Apostol, Maarten C. Stol and Patrick D. Forré, from University of Amsterdam and BrainCreators B.V., for the paper “FlipOut: Uncovering Redundant Weights via Sign Flipping”.
- The best thesis award, sponsored by BNVKI, was awarded to Louis Gevers and Neil Yorke-Smith, from Delft University of Technology, for the paper “Cooperation in Harsh Environments: The Effectsof Noise in Iterated Prisoner’s Dilemma”.
- The best demo award, sponsored by SKBS, was awarded to Eric Jutten, Edward Bosma, Kiki Buijs, Romy Blankendaal and Tibor Bosse, from The Simulation Crew, TNO Soesterberg and Radboud University, for the paper “Communication Training in Virtual Reality: A Training Application for the Dutch Railways”.
- One of the best video awards, sponsored by ZyLab, was awarded to Lesley van Hoek, Rob Saunders and Roy de Kleijn from Leiden University, for the paper “Evolving Virtual Embodied Agents using External Artifact Evaluations”.
- The other best video award, sponsored by ZyLab, was awarded to Thomas Winters and Pieter Delobelle from KU Leuven, for the paper “Dutch Humor Detection by Generating Negative Examples”.

All together, we think the BNAIC/BeneLearn conference 2020 was a great success, and would like to thank everyone involved in the organization and attending. We are looking forward to BNAIC and BeneLearn in 2021, organized in Luxembourg.

[Read more](#)

[Unsubscribe](#) | [Manage subscription](#)

Add your postal address here!