Benelux AI Newsletter

Welcome to the BNVKI newsletter of winter 2021! It is our pleasure to dedicate this newsletter to the fact that BNVKI celebrates its 40th birthday. The organisation — or rather, its Dutch forerunner, the NVKI — was founded in 1981. Since then, the domain of Artificial Intelligence has come a very long way!

If you would like a reminder of what the world of Computer Science was like at the time the (B)NVKI was founded, be sure to take a look at this footage from the "Human and Computer" exhibition in The Hague 1979.

Interview with Jaap van den Herik, Founding Father of the BNVKI

Interview with Jaap van den Herik, Founding Father of the BNVKI

by Emil Rijcken

The Benelux Association for Artificial Intelligence (BNVKI) celebrates its fortieth birthday this year. Although we might only be at the dawn of the AI revolution, a world without AI around us is already unthinkable. This was not the case forty years ago; at that moment barely, anyone knew the term. During the last few decades, AI has gone through immense growth. The world champions of chess and Go don't stand a chance against computers, cars can drive themselves, and some claim the Turing test has been passed. While some professionals fear their profession will become extinct, others are seemingly far from being replaced by AI.

During this interview with Jaap van den Herik, one of the Founding Fathers of the

BNVKI, we will recapitulate the advances over four decades of AI. Moreover, we will discuss what might be ahead of us.

In 1981 Jaap was one of the 19 founders of BNVKI. Together with Bob Wielinga and Dennis de Champeaux de Laboulaye, he was in the lead of the organization that was in the hands of Amsterdam and Delft. They had the ambition to place AI prominently on the research agenda in the Netherlands. When Jaap was awarded his doctorate in 1983 on '*Computer Chess, Chess World and Artificial Intelligence*' (see Relevant References, henceforth RR), one of his promotors, Professor Adrian de Groot, did not believe him stating that a computer would beat the world's greatest chess players in the future. A statement proven right when Kasparov, the human world Chess Champion (1985-2000), was defeated by IBM's DEEP BLUE in 1997.

Being a visionary, professor van den Herik (henceforth vdH) predicted in 1991 that machines would judge court cases and replace judges in the future (RR, 1991). This King's Gambit was a comical statement at the time but is becoming more realistic as technology evolves (see the movie the 'Queen's Gambit). During his career, vdH has awarded a doctorate to 91 PhD students, he contributed to the establishment of various organizations*, received a Humies award in 2014 for his research on computer chess, and upon his retirement as professor of Law and Computer Science, he was appointed Officer in the Order of Orange-Nassau, a royal award for individuals who have made a special contribution to society.

Professor van den Herik, you focused on computer chess in your PhD dissertation. Meanwhile, you were an excellent chess player yourself (2290 rating); could you still beat your algorithm?

Yes, I could. Early chess programs were not very good, and I could beat them quite easily until the mid-1980s and play at par to 1988. The actual start of the development was already in November 1966 at Bell Labs when the MACK HACK VI chess program was developed by Richard Greenblatt. It was the best in the world. However, it was only rated 1243 when playing in the Massachusetts Amateur Championship. After that, a full range of computer chess programs was built, culminating in the playing strength of over-world champion.

What have been the main contributions of your dissertation?

My PhD thesis (Delft, 1983) was the first thesis on Artificial Intelligence in the Netherlands. It was multidisciplinary (computer science, psychology, and philosophy). It started with a description of AI and the history of computer chess. In my research, I have worked on methods for *knowledge minimalization*, *knowledge classification*, *the use of equivalence classes*, the combination of classes, and the complicated combined evaluation in chess end games. This was followed by an impact analysis of programs that would defeat the human world champion. I collected ideas in personal interviews from Claude Shannon, Herb Simon, Ken Thompson, Donald Michie, Adrian de Groot, Mikhail Botvinnik, Max Euwe, Anatoli Karpov, Jan Timman, Genna Sosonko, and many others. Lastly, I focused on philosophical questions regarding computers, intuition, and creativity.

In advance of the ceremonial defense, the Delft University of Technology had arranged a press conference to discuss my research and thoughts on the potential impact of AI on society. My expectations were very high; I predicted that a chess computer would beat the world champion one day.

What did the term 'artificial intelligence' mean back then?

Nowadays, there are endless AI applications. In the beginning (the 1950s), this was not the case, and AI could be distinguished into four domains:

1. Chess and Checkers
2. Knowledge Representation
3. Problem Solving
4. Language Translation

What is your definition of AI?

Just as AI has evolved with time, so has my definition. I used Herb Simon's definition until my PhD defense: '*an AI program should mimic human thinking*'.

However, soon I realized computer programs would outclass humans in chess, and my ideas on the definition shifted. Computer chess was an excellent example for AI researchers, yet 'mimicking' was too restrictive on the definition. My second definition (1983 – 2000) was based on Donald Michie's definition dealing with the human window. My definition was: '*an AI program is a computer program that provides insights into the human thinking processes*'. The goal of this definition was to build programs within the scope of the human window, meaning that they are human executable and understandable.

Then, 'learning' and 'deep learning' entered the scene, and now I am inclined to separate the term 'artificial intelligence' from 'natural intelligence'. Obviously, 'intelligence' refers to clever behavior or a clever solution to a complex problem. But we should distinguish human intelligence and artificial intelligence from each other. My current definition (2000-now) reads: '*AI is the ability to address issues in the real world in an adequate way'*.

The developments in computer chess have similarities with developments in artificial intelligence; could you explain the developments per decade?

Early chess computers were based on search algorithms and were both rule- and library-based. Then, computing power increased exponentially, and so did the performance of chess computers. Generally, each decade is characterized by specific developments.

The fifties: the emphasis in this decade is on ***search*** algorithms; tree search, and evaluations. Claude Shannon and Alan Turing quantified all pieces and then summed the estimated values of these pieces. The position with the highest value was preferred. John von Neumann did similarly.

The sixties: in this decade, emphasis was on knowledge and ***knowledge representation***; special attention was on positional characteristics (e.g., developing pieces, open files). A prime example of this decade was Richard Greenblatt's MACK HACK VI chess program.

The seventies: the developments of the fifties (***search***) and sixties (***knowledge***) were ***combined*** in the seventies, of which my dissertation is a good example. Combining both aspects was made possible due to significant increases in computing power (MACK HACK VI, 1966 went from 200 nodes to 160 000 nodes per second by BELLE (Ken Thompson, 1980).

The eighties: computing power increased even further in the eighties with the introduction of ***parallelism***. The DEEP THOUGHT chess program defeated the human chess grandmaster Bent Larsen in 1988. It combined 64 chess playing chips and considered up to 500 000 positions per second.

The nineties: ***distributed*** systems were investigated and used in computer chess. Tasks were distributed and were executed through scheduling—two Dutch programs dominated in the first half of this decade, viz. GIDEON (Ed Schröder) and FRITZ (Frans Morsch) won the World Computer Chess Championship (WCCC) in 1992 and 1995, respectively. Then, in 1995 'IBM's DEEP BLUE I project started. It had 36 processors but lost 4-2 against human chess world champion Gary Kasparov in 1996. After that defeat, DEEP BLUE II was developed, with more computing power than its predecessors; it could evaluate up to 2.5 million positions per second. As a result, in 1997, DEEP BLUE II defeated Kasparov by 3.5-2.5.

In conclusion, DEEP BLUE marked the start of a new era of chess programs with ***advanced computing*** (known by the RS/6000 computer) and the introduction of machine learning. They used the so-called 'Dap 'Tap' for finding patterns in the opening libraries and later in the search processes. However, the findings and techniques developed by IBM were not publicly available. They were considered as *a single point of knowledge*.

The 2000s: in this decade, Frans Morsch made the over-human strength publicly available by his commercial products FRITZ and DEEP FRITZ. In 2002, FRITZ played an 8-game match with the new world champion Vladimir Kramnik; the result was 4-4. Then, in 2006 DEEP FRITZ won a 6-game match with Kramnik by 4-2. It was the end of human superiority in chess.

The chess community changed drastically. During the world championship matches, the public was no longer allowed to enter the playing hall, since all spectators knew the best move via their telephone, only the world champion and the contender did not know.

The 2010s: In research, most advances were achieved by incorporating machine learning, later deep learning and neural networks. In 2004 I started a project on Evolutionary Computing with Omid David Tabibi, Nathan Netanyahu, and Moshe Koppel. The topic was using ***genetic algorithms*** to tune the evaluation function so that the chess algorithm could learn from scratch (i.e., a program only knows how the chess pieces move). Our contribution to the GECCO 2014 conference was awarded the HUMIES Award 2014. It was a breakthrough since, in simple words, it showed the power of ***randomized learning***. The idea in itself has led already in 2012 to collaboration with Jos Vermaseren (Nikhef). We then applied the concept on Feynman diagrams and formulated a ***Monte-Carlo Tree Search*** for HEPGAME (High Energy Physics Game). The proposal was accepted as an ERC advanced research project. Moreover, in this decade, the rise of DEEP 'MIND's performances in computer Go was predominant. ALPHAZERO did ring a bell for all AI researchers.

The 2020s: although this decade has just started, I expect the ***Bidirectional Encoder Representation from Transformers (BERT)*** to mark the next era of state-of-the-art computer games (among them chess). BERT is a transformer-based machine learning technique initially proposed for natural language processing. But its strong capabilities in pattern analysis lend themselves well for chess and other games.

Monte-Carlo Tree Search is an essential algorithm in modern computer chess programs. What is it, and what was your role in its development?

Bruno Bouzy was the first researcher to publish on random search in a game tree for Go in 2004. I was privileged to be the editor of the book. Bouzy had two gifted students, viz. Rémi Coulom (presented first ideas in Turin, 2006) and Guillaume Chaslot (received a research place in Maastricht). Chaslot et al. (2008) designed and published the formal description of Monte-Carlo Tree Search (MCTS) (see RR).

MCTS is an effective tree-search technique characterized by building a search tree node by node according to the outcome of simulated playouts. The process can be broken down into four steps.

1. **Selection** – starting at root *R,* recursively select optimal child nodes until a leaf node *L* is reached.
2. **Expansion** – if *L* is not a terminal node (i.e., it does not end the game), create one or more child nodes *C* and select one *C*.
3. **Simulation** – run a simulated playout from *C* until a result is achieved.
4. **Backpropagation** – update the current node sequence with the simulation result.

Each node must contain two important pieces of information.

1. An estimated value based on simulation results.
2. The number of times it has been visited.

Will there ever be an 'optimal chess computer'?

Although there are approximately different positions in chess, I believe that the game can be solved and expect this to happen around 2035.

What would be the rating of an optimal chess computer?

By then, this question is irrelevant, or we have formulated a different interpretation of playing strength.

In this interview, we focus on the past and the future of AI. Your passion is two-fold, with chess and law. Why law as well?

In 1987, I was invited to join the Leiden Faculty of Law to make them familiar with modern developments in computer science. Inspired by Alan Turing's (1950) '*Can machines think?*' the step from computers playing chess to computers judging court cases seemed minor at first glance. However, I understood very well that the above question was audacious. Moreover, up to 1990, law and AI had received very little attention in the scientific world. Therefore, the invitation to go in that direction was exciting. Please, note that initially, in 1988, I knew very little of law, and it took me three years of hard work to develop a proper understanding.

In 1991, you predicted that computers would replace human judges in the future. Can you elaborate on this statement?

Whether computers will ***replace*** judges at some point in time is something I cannot predict (See the link https://www.universiteitvannederland.nl/college/kan-een-computer-een-rechter-vervangen). It depends on more than task performance only. Society and government will decide on acceptance. Moreover, the full range of tasks of lawyers, judges and paralegals is a topic of research, and there is no formal definition yet of computers being qualified. Still, my prediction is that computers will perform both simple and complex legal tasks at par or better than humans in the foreseeable future. Hence, in my opinion, empirical evidence will show us the best way for society (see RR).

Talking to a computer seems much different than talking to a human. Human judges can take the emotions of suspects into account and adjust their speech and non-verbal communication accordingly. How would this apply to computer-based judges?

I foresee that computers will be able to understand emotions in the future. Again, only time can tell whether society is prepared to accept such capabilities when exhibited by computers. I believe that the descendants of BERT will have a great future.

You don't believe in the 'computing paradigm', namely that computers perform analyses and structure data, while humans work on ethics, intuition, and creativity.

In my opinion, the computing paradigm certainly applies to our ***interaction with computers*** nowadays. Therefore, I do not exclude that computers will be capable of handling work on ethics, intuition, and creativity by the end of this century. ***Creativity*** will never be a problem for computer scientists to realize, starting at the end of your list. According to Michie, creativity is one of the least valuable capabilities of human beings since it can be best mimicked by genetic algorithms (or random search).

Intuition is another cup of tea. De Groot stated: "Playing at the level of a world chess champion requires intuition. Intuition cannot be programmed. So, a computer will never play on that level".Currently, my own opinion is that "Intuition is programmable" (see RR, my Valedictory Address in Tilburg, 2016).

Ethics is the real issue. To what extent ethics can be incorporated in computer programs cannot be answered in brief, mainly since ethics is culture-dependent. Here, I remark that in Law, we see many cultural differences in the jurisdiction. In my opinion, each local and global legal system can be implemented in a formal system endowed with conditions expressing the ***human measures***. So, ethics is the research challenge of the future.

What is your definition of ethics? And isn't ethics inherently subjective?

Formulating a definition of ethics is difficult as more than 170 definitions exist. It can also be called ***moral discipline***; thus, it is concerned with what is morally right and wrong. Please note that ethics is equally valuable to any system or theory of moral values or principles.

Indeed, ethics is subjective, both for humans and for computer systems. There is no such thing as being objectively right or wrong; there are just different approaches to ethical reasoning. I have thought about how a computer would handle this, but I cannot formulate suggestions for future research other than searching for ***human measures***.

What is 'the human measure'?

Recently, the human measure has been embraced by and upon policy execution. *Too tight regulations* can limit the ability to execute legislation at the cost of the human measure. As a result, formulating a definition for this measure becomes relevant. However, formulating a definition comes down to the philosophical question ***'what is a human****?*'. My definition of the human measure is: *in execution, the human measure means that the executor takes individual circumstances into account within the legal frameworks*.

There is no argument that individual circumstances are limited to a fixed set. How can a computer learn to handle each unique circumstance?

It is impossible to learn each unique case, a priori. But this holds both for humans and for machines. Why should a computer learn an adequate judgment for all subsets (or all elements) if a human has not done so either? Of course, in the practice of both (humans and machines), some sets of circumstances will be missed, but I assume that the approximations will be sufficient in relation to the human measures.

You predict that computers will realize the human measure eighty years from now. Will it be based on characteristics as predefined by humans? If so, isn't that a loss of information since some characteristics are tacit?

It could be based on predefined characteristics, which will undoubtedly be the case at the beginning of this line of research. This would indeed mean loss of information. But I am sure human investigators will catch up, maybe with the help of computer assistants. As there are so many challenges ahead, I still predict a realization within eighty years from now. We have to march on before identifying the new challenges more precisely.

Most AI algorithms are trained by learning patterns in vast amounts of data and could perform well on problems related to the past. But what happens if a new 'out of context' problem arises?

Algorithms will base their decisions on ***analogy***, ***distance measures*** and by developing ***new metrics***. Such decision making could be sufficient for some out-of-context problems too.

Are the algorithms we have nowadays adequate for developing an AI-driven judge? If not, what still needs to be developed?

At this moment, the algorithms are not sufficient. Probably, we need new ways of computing. However, we should start by keeping in mind that perfect tuning has not to occur by then. Moreover, some argue that we could get good algorithms through quantum computing. Still, I would not advise waiting for so long.

Humans are biased, and so will computers if they are trained on the decisions of biased humans. If we only train algorithms based on past data, AI judges will be biased forever. How can we prevent this from happening?

This is an excellent and vital question. We should place sufficient energy and money on research aimed at handling biases.

AI judges are facing a massive leap towards artificial general intelligence (AGI), in which an intelligent agent can learn any intellectual task that a human being

performs. A Nature publication states that AGI will not be realized. Do you think it will?

The paper has many truths and might be true. However, I believe there will be AI judges at some point, but the future is still open. Furthermore, I cannot oversee what AGI developments are expected to bring us.

This marks the end of our interview. Do you have any last remark?

Discussing past AI developments is relatively easy; predicting the future is more challenging. Currently, we are still some eighty years from having AI judges. I cannot foresee all the intermediate challenges ahead of us, but I trust these will be investigated adequately once raised. I cannot think of arguments stopping AI judges from being realized.

Relevant References (RR)

Bouzy, B. and B. Helmstetter (2004). Monte-Carlo Go Developments, *Advances in Computer Games 10: Many Games*, *Many Challenges*, pp. 159-174 (eds. H. J. van den Herik, H. Iida and E. A. Heinz), Springer.

Chaslot, G.M.J.-B. , M.H.M. Winands, H.J. van den Herik, J.W.H.M. Uiterwijk and B. Bouzy (2008). Progressive Strategies for Monte-Carlo Tree Search*, New Mathematics and Natural Computation*, Vol. 4, No. 3, pp. 343-357.

Coulom, R. (2007). Efficient Selectivity and Backup Operators in Monte-Carlo Tree Search, *Computers and Games (CG 2006),* pp. 72-83 (eds. H.J. van den Herik, P. Ciancarini and H.H.L.M. Donkers). Springer.

David, O.E., H.J. van den Herik, M. Koppel, and N.S. Netanyahu (2013). Genetic algorithms for evolving computer chess programs. *IEEE Transactions on Evolutionary Computation* Vol. 18, No.5, pp. 779 - 789. doi:10.1109/TEVC.2013.2285111. ISSN 1089-778X.

Fjelland, R. (2020). Why general artificial intelligence will not be realized. *Humanities and Social Sciences Communications*, Vol. 7, No.1, pp.1-9.

Herik, H.J. van den (1983). *Computerschaak, Schaakwereld en Kunstmatige Intelligentie*. TH Delft, June 21, 1983, 630 pp. ISBN 90-6233-103-3. (In Dutch)

Herik, H.J. van den (1991). *Kunnen Computers* Recht*spreken*? Inaugural address, Leiden University, June 21, Gouda Quint, Arnhem. ISBN 90-6000-842-1. (In Dutch)

Herik, H.J. van den (2016). *Intuition is Programmable.* Valedictory Address, Tilburg University, January 29. ISBN 978-94-6167-268-1.

Distinctions from the Field

*BNVKI – Benelux Association for Artificial Intelligence (Honorary member)

CSVN – Computer Chess Association of the Netherlands (Honorary member)

JURIX – Foundation for JURIdical eXpert systems (Honorary chair)

SIKS – School of Information and Knowledge Systems (Honorary member)

ECCAI / EurAI – European Community for Artificial Intelligence (Fellow)

ICCA / ICGA – International Computer Chess (Games) Association (Honorary Editor)

Read more

Conference report: BNAIC/BENELEARN 2021

The 33rd Benelux Conference on Artificial Intelligence and the 30th Belgian-Dutch Conference on Machine Learning (BNAIC/BENELEARN 2021) were organized as a joint conference by the University of Luxembourg, under the auspices of the Interdisciplinary Lab for Intelligent and Adaptive Systems (ILIAS) at the Faculty of Science, Technology, and Medicine (FSTM), and the IT for Innovative Services (ITIS) research department of the Luxembourg Institute of Science and Technology (LIST). The conference was part of the AI in Action joint event enjoying the generous support of the Luxembourg National Research Fund (FNR).

BNAIC/BENELEARN 2021 was held on-site in Belval, Esch-sur-Alzette – under CovidCheck regulations –, as a three-day event: from Wednesday 10 to Friday 12 November. BNAIC/BENELEARN 2021 had four special tracks: AI&Law, AI&Ethics,

AI&Art, and AI&Systems. The conference welcomed six invited speakers: Katie Atkinson (University of Liverpool), Julie Bernauer (NVIDIA), Fosca Giannotti (IST-CNR, Pisa), Manuela Naveau (Kunstuniversität, Linz), Carles Sierra (IIIA of CSIC, Spain), and Iris von der Tuin (Utrecht University). The BNVKI organized a FACt session this year too inviting three faculty members from the Benelux states to talk about the facts of AI: Gilles Louppe (University of Liege), Macq Benoit (Polytechnic School of UCLouvain), and Christoph Schommer (University of Luxembourg) gave short thought-provoking talks. Participants presented 104 academic and industrial research presentations, posters, and demonstrations – the conference provided ample opportunity for vivied interaction between academics and businesses bringing together more than 170 attendees on a live conference, first time after a long time. Additional programs included chess playing during the evening reception, a conference dinner with the award ceremony, and a European Cultural Capital event in the newly created Computational Creativity Hub at the Belval Campus, as Esch-sur-Alzette will hold the title in 2022.

The sponsors – offering an 500-euro prize in each case – and winners of awards at BNAIC/BENELEARN 2021:

- Best Paper Award:
  - Organization: Dutch Foundation for Neural Networks (SNN) – Award Chair: Bert Kappen
  - Authors: Gaoyuan Liu, Joris De Winter, Bram Vanderborght, Ann Nowé, and Denis Steckelmacher
  - Title: MoveRL: To A Safer Robotic Reinforcement Learning Environment
- Best Demo Award:
  - Organization: Foundation for Knowledge-Based Systems (SKBS) – Award Chair: Jaap van den Herik
  - Authors: Isel Grau, Luis Daniel Hernandez, Astrid Sierens, Simeon Michel, Nico Sergeyssels, Vicky Froyen, Catherine Middag, and Ann Nowé
  - Title: Talking to your Data: Interactive and interpretable data mining through a conversational agent
- Best Thesis Award:
  - Organization: Benelux Association for AI (BNVKI) – Award Chair: Tibor Bosse
  - Authors: Songha Ban and Lee-Ling Sharon Ong
  - Title: Producing "Open-Style" Choreography for K-Pop Music with Deep Learning

We congratulate the winners and thank the sponsors! The organizers also thank FNR for the support. See you all in Antwerp in 2022!

Read more

Junior Professor in Machine Learning (KU Leuven, BE)

KU Leuven's Faculty of Engineering Science has a fixed-term academic vacancy (5 years, part-time 95%) in the area of Machine Learning. The successful candidate will teach in the Master of Artificial Intelligence program, conduct research on machine learning (preferably with a focus on reinforcement learning and planning, but other specializations will also be considered), and supervise students in the Master and PhD programs. The candidate will be embedded in the DTAI section of the Department of Computer Science. More information is available at:

https://www.kuleuven.be/personeel/jobsite/jobs/60060924?hl=en&lang=en .

The deadline for applications is January 10, 2022.

KU Leuven is committed to creating a diverse environment. It explicitly encourages candidates from groups that are currently underrepresented at the university to submit their applications.

Read more

AI for Life Podcasts

Radboud University launched a podcast series called 'AI for Life'. In this series, interviewer Martijn Kriens (Briskr) has interesting discussions with experts from academia, industry and society on various application domains of AI. Check out the following link (in Dutch):

https://open.spotify.com/show/2HGbn1N9FS8ekqOuR3QrhH

Read more

Robot Farms, Artificial Life and Second-Order Engineering

-- By Emil Rijcken

In 1991 Guszti Eiben received his PhD in computer science at the Eindhoven University of Technology. He got his first AI introduction during his master's in mathematics at the Eötvös Loránd University in Budapest. At the time, AI was very esoteric, and Guszti Eiben could only daydream about its possibilities. Now, as a professor of AI at the Vrije Universiteit in Amsterdam and as a special visiting professor at the University of York, he has the tools to realize all his sci-fi dreams. Professor Eiben is a pioneer in Evolutionary Computing; he wrote the textbook that is now used at many universities, was the first in the world to let two robots have a baby and is constantly pushing the field to go beyond the status quo.

A consortium with his UK colleagues received a two-million-euro grant to research Autonomous Robot Evolution (ARE) in 2018. The ARE project is building an EvoSphere, an evolutionary robot habitat serving as a tool to study evolution and the emergence of intelligence. Fundamental evolutionary questions for both biologists and computer scientists can now be answered. But the EvoSphere is also expected to push engineering; applications include autonomous robots to colonize space and cleaning up nuclear reactors on Earth. Professor Eiben postulates that if evolution can create intelligence, then artificial evolution can create artificial intelligence.

Professor Eiben, continuing on your analogy: god created offspring, and you have created artificial offspring. Is it fair to say that you are an artificial god?

Definitely not, and I disagree with your first part of the statement that 'God created offspring' because evolution created offspring. This is a big difference in starting points. To put it simply, in the beginning, physics became chemistry, and chemistry became biology. Simultaneously to the last transition, evolution emerged. We are still in the grand process of evolution, 'the greatest show on earth'. I can only change the substrate and have a minor role in this process. The only form of evolution we know and analyze is carbon-based, which is life on Earth. People like me do not introduce new principles; we just introduce evolutionary principles on a new medium.

The sequence is as follows; in the 19th-century, Darwin described evolution in wetware. In the 20th century, evolution in software 'evolutionary computing' was invented by computer scientists. Now, in the 21st century, we are working on evolution in hardware. So, there are two grand transitions:

-
    - wetware to software
    - software to hardware.

The result is the evolution in hardware, which is different from carbon-based life as we know it.

You are considered a pioneer of evolutionary computing; what are your main contributions?

In the late 20th century, a community arose that used search algorithms based on Darwinian principles, selection and reproduction to solve problems and developed what is now known as evolutionary computing. Yet, in the early years, no one knew exactly how to fine-tune evolutionary algorithms. Firstly, I contributed to the methodology. I have burnt three PhD students investigating and optimizing the hyperparameters in evolutionary processes. I demonstrated how the quality of evolutionary processes depends on these hyperparameters and put it on the research agenda. Also, I optimized the optimizer to improve the outcome of the evolutionary process.

Secondly, after waiting for a good textbook for a long time, I wrote one myself (with Jim Smith, a friend and colleague from the UK). Our textbook is now perhaps the textbook of evolutionary computing, used at many universities and recently, it has been translated to Chinese.

Thirdly, I studied fundamental questions based on reproduction mechanisms. For example, I investigated what would happen if we had more than two parents. In biology, we know about two kinds of reproduction: sexual- and asexual. Asexual reproduction consists of mutations only, while sexual reproduction requires two parents and is used by higher life forms such as homo sapiens and fishes.

As a mathematician, I saw the number of parents just as a reproduction parameter. It does not have to be limited to two, and theoretically, it could be three, four, or even more. I investigated what happened to evolution if more parents were to reproduce and demonstrated that having more than two parents in a crossover operator accelerates evolution. Depending on the kind of crossover operator, the optimal number is somewhere between two and ten.

Lastly, and perhaps my most significant contribution; I promoted the optimization of active rather than passive objects. In classical evolutionary computing, objects like routes for a travelling salesman or the design of an industrial object are subject to optimization. I started optimizing things with agency. In that domain, I investigate and evolve active artefacts, agents, organisms, or simulated robots. In a way, these are all the same types of entities; they have a body and a brain. Having both makes them much harder to evolve but also much more interesting to study.

Before we dive into details, could you give a short introduction; what is evolutionary computing?

It is a collection of search algorithms with its own style; the evolutionary style. This style has adopted the principles of reproduction/variation and selection from biological evolution.

One could argue that all search algorithms have the same properties, namely: generate and test. The 'generate' step is equivalent to reproduction/variation, and the 'test' step is used for selection. Evolutionary algorithms are unique within the big family of search algorithms because they use a population of solutions and crossover as a search operator, combining two or more solutions. This is unique as all other search methods iterate only one solution, applying perturbations (mutations) to produce new solutions. Also, the stochastic character of selection and reproduction in a population is an important special feature.

In conclusion, evolutionary algorithms are population-based, stochastic search methods. Evolutionary computing is motivated by evolutionary principles, and the search steps can use more points in the search space to generate new points. We do not exactly simulate or emulate evolutionary (carbon-based) mechanisms, but we use evolutionary principles.

Is evolutionary computing a form of artificial intelligence?

Definitely. To position it even further, roughly, AI can be divided into symbolic/top-down and sub-symbolic/bottom-up approaches. In symbolic AI, the algorithm designer is most prominent in setting rules, whereas the designer is less prominent in sub-symbolic approaches. In the 20th century, the top-down approach was the dominant approach, while the sub-symbolic approach has become dominant in the 21st century. With sub-symbolic AI, algorithms find solutions based on a predefined method. However, the resulting rules that produce the output are sometimes very untransparent and unexplainable. Evolutionary computing is a form of sub-symbolic, bottom-up, 21st century AI.

There are many applications of your work; which application are you most excited about?

All of them are very interesting, but their visibility can distinguish some more than others.
Allow me to explain one thing first; evolution is not necessarily uncontrolled. Evolution can be supervised and controlled by humans; we refer to this as 'breeding'. Farmers can breed species such as crops and animals. From the two principal operations of evolution, reproduction and selection, humans could control selection for thousands of years already. For instance, farmers can decide which bull and cow can be coupled to make calves. Influencing the selection component of evolution for several generations is effectively steering evolution towards desirable outcomes *.

Provided this context, the most visible applications in my work are 'robot breeding farms', where evolution happens under human supervision. Humans can aid, direct, and accelerate the evolutionary process towards the desired result. Given enough resources (read: funding), we could have robot breeding farms within five years. Such a breeding farm will employ evolution as a design approach, running many generations under supervision and stopping evolution once a good solution emerges. This approach does not replace traditional ways of designing robots, but it has a niche complementary to the usual applications.

A good example of applications to illustrate this niche are robots used for monitoring the rainforest. This problem is very complex because we have no idea what kind of robot is optimal for that environment; should the robot have wheels, legs, or both? Should it be small to sneak through the holes in the vegetation, or should it be big so that it can trump down obstacles? Robots, designed through classic engineering methods, only work in a static, predictable and structured environment (e.g. warehouses). But if the environment is complex, dynamic, and not known in advance, finding a good design is very hard, and evolution is your friend. The way I put it to my students is: When the going gets tough, evolution gets going.

Based on this idea, we can use robot breeding farms to get a well-designed body and brain that operates well in our mockup forest. After that, we create many copies of the optimal robots and send them out to monitor the real environment. This is one of my favourite applications because we can do it quite quickly, and it is relevant for society.

In the long term, we could have evolutionary processes that operate without direct oversight from humans. This means a hands-free, almost open-ended evolutionary process. However, this process raises both ethical and safety issues about runaway evolution. At the same time, this approach has highly useful applications. For example, it can be used for space research. We could send an evolutionary robot colony to another planet and have them do what life did on Earth. Firstly, they need to evolve and adjust to the circumstances to survive and operate for a long time. Once they can survive, they could activate human-related tasks, such as building houses or making the planet habitable for humans. This is a different approach than the breeding farm, as it is not directed nor controlled.

Suppose that a multitude of robots is sent out to perform a task. How do robots choose between performing their task and finding their mating partner?

A time-sharing system is the most logical. For example, task execution can be (almost) permanent, while mating can be occasional, perhaps triggered by time (e.g. a mating season) or an event (e.g. meeting another robot).

Are these actual applications, or is this a hypothetical discussion on possibilities?

The breeding farm is an actual application and concerns a collaboration between the University of York, the Bristol Robotics Lab, the Napier University, Edinborough and the Vrije Universiteit, Amsterdam. Our goal is to develop robots capable of cleaning up nuclear power plants. Typical for such visionary projects is that we will not make it. However, we learn a lot and know how to make it work if we get another four years of funding.

The space application with hands-free, autonomous robots is for the future and will take another ten or fifteen years. However, these applications are not just a matter of money and engineering. There are fundamental technical, scientific and ethical questions that need to be addressed first; how can we set up a system to operate, do what we want, and do no harm?

Which fundamental questions are you most excited about?

My two favourites are '*How can intelligence evolve from a non-intelligent beginning?*' and '*What is the interaction between the body and the brain behind (evolved) intelligent behaviour?*' The premise is that intelligence is not just in the brain but also in the body. All existing forms of intelligence we know are hosted in a body, and we do not know any intelligence that does not need a body. This indicates that intelligence needs both the brain and the body. More specifically, behaviour is always determined by the body, the brain, and the environment.

Humans can walk on two legs on solid ground very well. But if you put them into the water, they sink. If the body does not change, but the environment does, then the behaviour needs to change, e.g., swimming instead of walking. This is fascinating both from a fundamental and a practical perspective.
For example, an interesting question we investigated is: '*What is more important for intelligent behaviour; a good body, or a good brain? And how do we get this via evolution?*'. There were many caveats to answering this question because the answer can depend on the experimental setup, the given robot design, and environmental details. Yet, I have quantified the question, answered it through experiments, and published the result at an annual artificial life conference with a student and a colleague. In our evolutionary robot system, the body is more important for intelligent behaviour than the brain.

What was your experimental setup?

We designed a system with an essential property; all possible bodies and brains could be combined into a working robot, and we measured the behaviour of each combination. Simply put, we found a technical solution so that even a fish's brain could be put on a human's body and work.

We generated 25 bodies and 25 brains, resulting in 625 combinations arranged in a table and evaluated each one of them. Then, we looked at the standard deviation of the columns and the standard deviation of the rows. If the standard deviation is low, then that part is more important. To understand why, imagine that the rows are bodies. In that case, you get a body with 25 combinations of different brains, resulting in 25 fitness values. If these fitness values are in a small range, it does not matter what brain you put on there; you always get approximately the same intelligence. However, if you have a larger spread of fitness values, then the intelligence depends on the brain to a greater extent.

This is how we quantified our naïve question into a scientific question. After formulating the question, we 'only' had to run the simulations and fill out the body-brain matrix. In the end, we found that the spread is always smaller when the body is fixed. In this system, the body was more determinant for behaviour quality.

It is not just the engineering that I find interesting about this question. I am especially fascinated by the fundamental, even philosophical aspects, the interplay between body and brain and how they develop simultaneously through evolution.

So, the body is most important for intelligence. Yet, humans have a fixed set of body parts. The optimal number of parents is greater than two. Yet, humans only have two parents. Did humans get stuck in a local optimum?

No, the optimal number of parents is also determined by practicality. More parents are less practical and require more effort and luck to mate.

Is an Evosphere the same as a robot breeding farm?

No, not necessarily; the EvoSphere is a generic concept, while a breeding farm is one specific subtype. The human is in the loop in the latter and supervises the selection and infant learning process in the 'robot nursery'.

In contrast to the breeding farm, the Evosphere also allows for open-ended robot evolution without direct human oversight. The Evosphere is a generic system architecture that consists of three components. The first one, the Robot Fabricator or 'Birth Clinic', produces robot offspring. In evolutionary terms, a genotype (the robotic DNA) is converted into a phenotype, a real robot. The second one, the Training Center or 'Nursery', is where 'newborn' robots learn optimal body control. This stage is called the 'infant period'. During this period, robots learn new skills and cannot produce children. After the infant period, the robots become fertile and make it to the 'arena', where they operate and produce children. This is a generic system architecture applicable to all robot evolutionary systems, regardless of how the details are implemented.

Which Evosphere's component do you prefer working with?

I have been more interested in and challenged by the infancy period in the last couple of years, which is similar to machine learning but also completely different. The best way to explain the difference is by inverting the words from 'machine learning' to 'learning machines'. The message is that with 'learning machines', you are discussing machines, either simulated or physical, capable of learning. Notably, a learning machine generates its training data by performing actions, whereas, in machine learning, users feed the algorithm with predefined training data sets.

Learning machines form a big challenge in the context of my research; every new robot has a different body (e.g. more legs, fewer wheels, different sensors, the camera on the other side) that needs a body-dependent controller. Thus, each robot represents a new learning problem: how to control the given body optimally and ensure that the robot can operate, e.g., walk, perform tasks, survive, and reproduce.

Human babies spend a year learning to walk and grasp objects. Evolvable robots also must develop their 'hand-eye coordination' quickly after birth. The problem for robots is more challenging because human babies always have the same body as their parents (e.g. two hands, five fingers on each hand). In contrast, robot offspring can have completely different bodies than their parents. We are using some learning techniques from machine learning, such as reinforcement learning and neural networks. So, although many machine learning algorithms are potentially helpful, we do not know anything about the robot's morphology (body plan) in advance and cannot make any assumptions. Hence, we need learning methods that work on all possible robots in our design space. Each new robot produced by evolution is the equivalent of a new dataset in traditional machine learning.

This learning problem is only a stepping stone to the really big challenge; finding out how evolution and learning influence each other. This question has been discussed for more than a hundred years and arose in the biology community. They invented notions like '*Lamarckian evolution*' or the '*Baldwin effect*', the early AI community picked up. Forty-year-old papers investigate the combination of learning and evolution in settings that I would now describe as artificial life systems. This is a prominent issue for evolving robots because learning in the infancy stage is essential. This transforms the theoretical question on the interaction between evolution and learning into a practical one: how to combine evolution and learning in robots to maximize efficiency and efficacy?

Ultimately, I am interested in the combination of evolution and learning in one system; it would be a significant avenue to realize a new level of AI. I believe that future AI will be produced by autonomous processes rather than human developers encoding the solution. I call this phenomenon 'second-order engineering' or 'second-order development'. The standard approach for developing an AI system (robotic or otherwise) is based on a developer who analyzes the problem, does a literature search, and designs and implements the target system. This is typical 'first-order engineering'. With second-order engineering, we develop an evolutionary system that develops a solution for us, rather than us constructing an AI or robot system directly. I am convinced that second-order engineering will become more prominent in future AI.

What is the role of humans in second-order engineering?

Humans should specify the components of the evolutionary system. For instance, define the genetic language used in the genotypes, specify adequate mutation and crossover operators, formulate the fitness function, and determine the conditions for reproduction. If learning is applied, then the learning method(s) need to be defined as well.

There are two critical issues here: sample efficiency and safety. Biological evolution is highly wasteful. It creates a lot of solutions, most of which die before they become fertile. An artificial evolutionary system cannot be too wasteful because the time scale is weeks or months rather than millions of years. Additionally, we need to ensure safety in an inherently stochastic and adaptive system that produces real robots in the real world. The obvious dangers are runaway evolution and the emergence of unwanted or dangerous robot properties. Safety and ethics are essential. Yet, not much is known about these issues as we are just starting to learn about them. However, we need to be aware of the ethical and safety issues from the first moment onwards.

In earlier media outings, you discussed centralized reproduction. Is this a form of safety measure you propose?

Yes, it is, as it can help prevent runaway evolution. My solution is to reject distributed reproduction systems such as laying eggs, becoming pregnant or cell division because these reproduction systems would allow robots to reproduce anywhere and in any way without having the option to stop them. Instead, I insist that we only build evolutionary systems with a centralized unit for (re)production of robots, the first component of an EvoSphere. This unit serves as a safety switch; once it is turned off, reproduction stops, and there will be no more robot offspring. The existing robots may not drop down 'dead' immediately, but at least they will not further reproduce.

How is the fitness function defined for autonomous robots?

The research community's standard approach is to have one task and equate fitness with task performance (e.g. for a robot that should be fast; fast robots will have many children, and slow robots will not). This guarantees that evolution creates robots that are good at doing that task. I am trying to nudge the research community to go further and do more complex tasks with practical relevance and consider multiple tasks simultaneously. To survive, robots need to be good at many tasks.

Let's assume we are sending a robot colony to a distant planet. There are multiple tasks to perform. How would you define the fitness function?

Here, we should distinguish between skills and tasks. The number of combinable, elementary skills necessary for complex tasks is relatively small, less than ten. Take locomotion; a robot has to walk. Then, locomotion should be targeted; a robot should learn how to move to a specific target and avoid obstacles. Subsequently, a robot needs to learn to manipulate objects. This set of elementary skills can be used as stepping stones so that the robots can perform more complex tasks too. Robots will learn these skills in the 'robot school', enabling them to perform more complex tasks.

Let's focus on the evolution of things and second-order engineering. Can you imagine a context in which a robot would develop a consciousness, morality or emotions to perform tasks?

Let me define 'the evolution of things' firstly. Before, I described the transition from wetware to software and from software to hardware. Similarly, such a sequence could also be related to evolution: from the evolution of living organisms to the evolution of code (evolutionary computing) and then to the evolution of things (robots).

The question about consciousness is hard to answer; it is more fundamental and philosophical. I cannot say whether they will or could have consciousness. The following analogy is the easy way out: if it walks like a duck, looks like a duck and quacks like a duck, then it probably is a duck. If the robots' actions match our standards of morality, we could call them moral robots, regardless of the mechanism that drives this behaviour. Moral behaviour is designable and desirable; they need to adhere to our standards.

But could this also be an outcome of an evolutionary process?

Having evolution or any other adaptive, emergent process at work does not mean that we cannot control it. We must develop the technology and science to control these emergent processes and ensure they respect our constraints, which we could call moral- or ethical borders. However, setting such constraints comes down to one of the biggest questions in bottom-up, sub-symbolic AI; how to limit evolutionary processes without disabling them? Thus, how to keep evolution within our ethical borders, without 'curtailing' their behaviour too much. I have no answer and can only emphasize that it needs further attention.

Suppose you have one robot that can choose between two robots for reproduction. Both robots are identical in terms of functionality. Is it plausible that a robot bases its decision on aesthetics?

Based on our engineering-based perception, we are inclined to choose robots based on functionality and usefulness. But life does not work like that. The idea that you propose is very good, and we are investigating it right now. This approach is different from the usual evolutionary algorithms. Firstly, the selection is not made centrally, while almost all artificial evolutionary systems have a centralized protocol, 'the manager' (technically the main evolutionary loop), to decide which robots are mating with which other robots. The decision is based on complete information on each population member. This is a desirable property for algorithmists but not for the artificial life community. Therefore, the first change is to enable robots to select mating partners themselves. Secondly, the selection criteria for mating partners should be changed. Currently, two robots can meet each other and decide whether they want to have a “baby” purely based on utility (task performance). In the new approach, we extend or replace this criterion with another one related to the morphology of the robots –beauty, if you will. Typically, utility is linked to behaviour; '*Tell me how many soil samples you collected in the forest, and I will tell you if I want to have a baby with you*'. Alternatively, you can change it to; *'I look at you, and I will tell you whether I want to have a baby with you*'. So yes, aesthetic-based selection is possible, hugely exciting, and we have just started investigating it.

Is this aesthetic-based approach interesting for engineers as well, or primarily for the artificial life people?

For engineers, it is less interesting as they are utility-oriented. It is primarily interesting for artificial life, theoretical biology or philosophy. An interesting question is: what kind of bodies/morphologies do you get if you have morphologically driven selection? The peacock is a prime example that I always have in mind. Peacocks have fantastic morphological features; their massive tails. But although the tails are beautiful, they are utterly useless and even dangerous. The tails make peacocks easier to catch by predators, and they require the peacock to eat more food. Yet, this morphological property heavily impacts whether peacocks will reproduce or not. I am curious to see whether we would see this phenomenon evolving in a robot system as well. If our evolutionary mechanisms capture fundamental properties, we could create an artificial evolutionary system with the same attractors as carbon-based evolution. Carbon-based evolution took millions of years to develop and is very complex. Artificial evolution has been developed only for a few decades and is much less complex, so it is not a done deal that this is possible.

However, it is extremely exciting; finding such effects would indicate that we understand the essence of evolutionary systems. There is something fundamental about evolution as such, regardless of the substrate we can capture. That would also give a hint on life on other planets. If all kinds of evolutionary systems are similar, then evolution on other planets could also be similar.

You state that as artificial intelligence has changed our view on intelligence, it is likely that artificial life will change our view on life. How do you think our view of life will change as a result of artificial life?

The notion of life will no longer be restricted to carbon-based life, which is the only kind of life we currently know. If many scientists agree that evolving robot systems constitute life, it will be acknowledged that life can have a different base. Other life forms can be digital, mechatronic or based on new materials with new forms of actuation and sensing. This means that the criteria for determining whether something is living or not will change; they need to be more about functionality rather than about 'incarnation' or instantiation. A broader definition of life will enable more generalizable research on life. As a scientist, you do not want to draw conclusions based on one sample only. However, currently, we only have one sample of life. More samples would lead to better-funded conclusions and to better insights into what life is about.

To this end, it is important that life as we know it is moderately observable, hardly controllable, and not really programmable, making it hard to study experimentally. But robots and artificial organisms are observable, controllable, and programmable. For instance, it is possible to retrieve robot communication by registering wifi signals, and internal processes can be logged on a black box inside the robot. In principle, this could cover everything: all sensory inputs, all information processing in the robot brain, all control commands, battery levels, etcetera. Such data can be stored and analyzed offline or used in a control loop to probe the system during its operation in an online fashion. This provides us with an extended set of tools to study and understand life and intelligence.

Ultimately, evolutionary robot systems represent a radically new kind of research instrument that can help understand the emergence of intelligence. The key open question here is: 'How did intelligence emerge?' and as of today, even the simplest answers are lacking. For example, is the process of acquiring intelligence linear, stepwise, or is it a hockey stick curve? Evolving artificial life systems allow us to investigate these questions, which would be an enrichment of artificial intelligence as we know it.

That concludes the interview, professor Eiben. Do you have any last remarks?

Emergent intelligence and second-order engineering have very significant risks. These risks have to be addressed from the beginning and the ground up while developing such systems. If we only try to mitigate them once they occur, it will be too late.

**Directly interfering with reproduction became possible after genetic manipulation was invented. This is now ethically debated but technically possible.*

Read more

Marco Wiering (1971-2021)

Unexpectedly, a brilliant researcher passed away. Marco studied computer science at the University of Amsterdam and graduated with honours (cum laude) in 1995. After his graduation he went to Lugano in Switzerland to perform his PhD research in the Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA) in the area of reinforcement learning, under supervision of the well-known prof. Jürgen Schmidhuber. After his PhD graduation in 1999, he did a short post-doc in the Intelligent Systems Group from the University of Amsterdam. He became assistant professor in the year 2000 at Utrecht University, and joined the AI institute ALICE, at the Rijksuniversiteit Groningen as tenure tracker in 2007. His main interest was reinforcement learning. This is the type of learning where, e.g., a lazy instructor tells the intelligent agent after an hour of driving that 'it was pretty bad today' (or 'good', as Marco preferred). Such a learning paradigm is much more complicated than loss and gradient-based learning (back propagation), which enjoys detailed error information over all output dimensions. In reinforcement learning, an intelligent agent needs to find out itself what was good or bad in previous actions, taken under previously-perceived states of the world. With Martijn van Otterlo, he edited and coauthored the book "Reinforcement Learning: State-of-the-Art" (2012, 19 chapters, 630 pages, Springer). In February 2019, Marco and I went to Prague. It was a multi-conference (ICAART, ICPRAM and ICORES). We had our own lectures, in our own topics. I had time to join his presentation at ICPRAM. After his presentation, a crowd of young PhDs clustered around the speaker. Naturally, because of the content of his presentation. But also because of his personality: Friendly, accessible, authentic and definitely not like your average prof. This pattern confirmed what we saw here in Groningen. Marco attracted large numbers of bachelor and master student. He meticulously corrected their texts, smoking a cigarette on one of the tables outside the Bernoulliborg. More often than not, he lifted the students above themselves, regularly resulting in a joint publication of their thesis work.

During the past months he was increasingly tormented by demons conjured up in his mind. Regardless of this sad affliction, he tried to be a good researcher and teacher until the end. He was an intense, trustworthy, righteous person with a fascination for our research field, internationally renowned, loved by colleagues and students. He will be sadly missed.

Haren, 23 September, 2021

Lambert Schomaker

Note: On Google Scholar, Marco's statistics are: h-index 37, number of citations 6539, best-cited publication: 720 citations. It is a little-known fact among researchers in computer science and artificial intelligence that their beloved Google Scholar profile page will be completely removed after a university has disabled the email address. Therefore, I made a snapshot of his current profile state. Undoubtedly Marco's work will continue to attract attention.

Read more

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.