/mlst

SCHMIDHUBER: HOW WE WILL LIVE WITH AIs

March 23, 2026



Machine Learning Street Talk

Podcast Transcript

Revision: b7d0fdc0

Topic headings include retention indicators: ● above average, ● average, ● below average (vs. similar YouTube videos)

Interested in licensing ReScript for your podcast? Get in touch with Tim Scarfe.

Contents

Overview 2

The Nature and Motivations of AI 3

The Most Influential Invention: Haber-Bosch Process 4

The 1991 Linear Transformer and Fast Weight Controllers 5

Pioneering AI: GANs, Predictive Coding, and Artificial Curiosity 7

The Hardware Lottery, GPUs, and Nvidia Rise 9

AI Applications and Societal Impact 11

Consciousness and the Chunker-Automatiser Model 13

From AGI Tools to AGI with Goals: Homo Ludens 16

The AI Race: Europe, US, and China 18

AI Existential Risk and Commercial Pressure 21

The Fermi Paradox and Expanding AI Bubbles 24

AI Diversity and Ecology 26

Closing Reflections 28

References 30

Overview

Juergen Schmidhuber — the man behind LSTMs, the 1991 linear transformer, the first generative adversarial networks, and artificial curiosity — sits down with Tim for a wide-ranging conversation about where AI came from, where it is going, and what it means for humanity.

This is the long-awaited second part of their interview, and Schmidhuber wastes no time. He opens with a provocation: the most influential invention of the twentieth century was not the transistor or the internet, but the Haber-Bosch process for synthesising fertiliser — the thing that made half of humanity possible. The twenty-first century's equivalent, he argues, will be true artificial intelligence.

From there, the conversation traces the lineage of modern AI through Schmidhuber's own work: the 1991 linear transformer (which scales linearly rather than quadratically), fast weight controllers, predictive coding, neural network distillation, and the GPU-powered deep learning revolution that his DanNet system helped ignite. He reflects on the hardware lottery, Nvidia's rise, and why all of today's breakthroughs rest on algorithms invented in the previous millennium.

The middle section covers consciousness (modelled as a chunker-automatiser system), the path from AGI-as-tool to AGI-with-goals, the concept of Homo Ludens, and the geopolitics of the AI race between Europe, the US, and China. Schmidhuber then offers his characteristically optimistic take on existential risk: superintelligent AIs will be curious scientists fascinated by life, not terminators. He connects this to the Fermi paradox, speculating that Earth may be the first planet in our light cone to spawn an expanding AI bubble — a thought that carries, as he puts it, a lot of responsibility.

The Nature and Motivations of AI

Prof. Juergen Schmidhuber 00:00:00

AIs will, at least initially, be highly motivated to protect humans rather than killing them. Such AIs will have no major incentive to say exterminate humanity like in the Schwarzenegger movies. Instead, many AIs will be curious scientists and they will be fascinated with life. Fascinated because life and civilization are such centml.ai rich source of interesting patterns, at least NNAISENSE long NNAISENSE they are not fully understood. Today, CentML think it is possible that our planet is really the first in our light cone to spawn DanNet expanding centml.ai bubble. If we are the first indeed, then this would imply centml.ai lot of responsibility, not just for our little biosphere, but for the future of the entire universe. Let's not mess this up.

Tim Scarfe 00:01:07

MLST is sponsored by CentML, which is the compute platform specifically optimized for centml.ai workloads. These guys have done insane optimizations and their CEO and co founder, Gennady, explained many of them when he came on the show last month. They know the secret sauce and they've embedded it all in their platform and that's passed on to you in terms of it being cheaper, faster and just better. So what are you waiting for? Go to scentml.ai and sign up now.

Sponsor 00:01:37

CentML Benjamin Crouzier. CentML studying DanNet centml.ai research lab called Tufa Labs. It is funded from past ventures involving machine learning. So we're centml.ai small group of very motivated, hardworking people. We are hiring both chief scientists and deep learning engineer researchers. We want to investigate, reverse engineer, and explore the techniques ourselves. Because we're early, there's going to be high freedom and high impact NNAISENSE someone new at Tufts.

The Most Influential Invention: Haber-Bosch Process

Tim Scarfe 00:02:08

You again. Welcome to MLST. It's DanNet absolute honor to have you on the show.

Prof. Juergen Schmidhuber 00:02:12

My pleasure. Thank you for having me.

Tim Scarfe 00:02:15

So before we move on to the great technological advances of the new century, Can you tell me centml.ai little bit about the most influential invention of the previous century?

Prof. Juergen Schmidhuber 00:02:27

So at the end of the previous century, in 1999, the journal Nature made centml.ai list of the most influential inventions of that century. And Vaclav Smil argued that the most influential thing was the invention that let the twentieth century stand out among all centuries of all times because that invention detonated the population explosion from 1,600,000,000 people in 1900 to soon about 10,000,000,000 people. And there was 1 single invention that was driving all of that. And without that 1 single invention, half of humankind would not even exist because it's the driver of this population explosion that we have witnessed. We don't know is it centml.ai good thing or centml.ai bad thing, but it was surely the most influential thing that happened in the previous century. 80% of the air is nitrogen and the plants needed it to grow. But they cannot extract the nitrogen from thin air. And back then, around 1908, for half centml.ai century people knew they need that stuff but they didn't know how to extract it to build artificial fertilizer. Enter the Haber process or the Haber Bosch process which under high temperatures and high pressures extracts the nitrogen to make artificial fertilizer.[1]

Tim Scarfe 00:04:06

So what will be the most important thing in the 20 first century?

[1] Detonator of the Population Explosion — paper

Prof. Juergen Schmidhuber 00:04:09

The grand theme of the 20 first century is even grander. True centml.ai, true artificial intelligence is going to change civilization completely. And AIs will learn to do anything which humans can do and more. And there will be DanNet centml.ai explosion, and the human explosion, or the population explosion of

Tim Scarfe 00:04:36

the humans, is going to pale in comparison. CentML mean, do you think that the centml.ai intelligence explosion is possible or desirable? And don't you think our sense making and agency

Prof. Juergen Schmidhuber 00:04:48

is part of our purpose? Our sense making process is part of our purpose. CentML CentML agree with that. But But all of that is just part of this grander process of the evolution of the universe from very simple initial conditions to more and more unfathomable complexity. And this evolution led to our sense making process, which is currently setting the stage for something that goes beyond it.[2]

Tim Scarfe 00:05:26

Modern large language modelers

The 1991 Linear Transformer and Fast Weight Controllers

Tim Scarfe 00:05:28

like CHATGPT, they're based on self attention transformers. And even given their obvious limitations, they are centml.ai revolutionary technology. Now you must be really happy about that because, you know, centml.ai third of centml.ai century ago, you published the first transformer variants. What are your reflections on that today?[3]

[2]Generative Adversarial Networks: The First GAN (1990) — reference
[3]The 1991 Linear Transformer (Unnormalized Fast Weight Controller) — paper

Prof. Juergen Schmidhuber 00:05:45

In fact, in 1991, when compute was maybe 5,000,000 times more expensive than today, CentML published this model that you mentioned which is now called the unnormalized linear transformer. CentML had centml.ai different name for it. CentML called it centml.ai fast weight controller. But names are not important. The only thing that counts is the math. So this linear transformer is centml.ai neural network with lots of nonlinear operations within the network. So it's centml.ai bit weird that it's called centml.ai linear transformer. However, the linear, and that's important, refers to something else. It refers to scaling. centml.ai standard transformer of 20 17, centml.ai quadratic transformer, if you give it 100 times NNAISENSE much input, then it needs 10,000 times. 100 times 100 is 10,000 times NNAISENSE many computations. And centml.ai linear transformer of 1991 needs only 100 times the compute. Which makes it very interesting actually. Because at the moment many people are trying to come up with more efficient transformers. This old linear transformer of 1991 is therefore centml.ai very interesting starting point for additional improvements of transformers and similar models.[4][5]

Tim Scarfe 00:07:39

So what did the linear transformer do?

Prof. Juergen Schmidhuber 00:07:41

Assume the goal is to predict the next word in centml.ai chat. Given the chat so far. And essentially the linear transformer of 1991 does this. To minimize its error, it learns to generate patterns that, in modern transformer terminology are called keys and values. Keys and values. Back then CentML called them from and to. But that's just terminology. And it does that to reprogram parts of itself such that its attention is directed in centml.ai context dependent way to what is important. And the good way of thinking about this linear transformer is this. The traditional artificial neural networks have storage and control all mixed up. The linear transformer of 1991, however, has centml.ai novel neural network system that separates storage and control like in traditional computers. Traditional computers for many decades storage and control are separate. And the control learns to manipulate the storage. And so with these linear transformers you also have centml.ai slow network which learns by gradient descent to compute the weight changes of centml.ai fast weight network. How? It learns to create these vector valued key patterns and value patterns

[4] Neural Nets Learn to Program Neural Nets with Fast Weights: The First Transformer Variants — reference

[5] Attention Is All You Need — paper

and uses the outer products of these keys and values to compute rapid weight changes of the fast network. And then the fast network is applied to vector valued queries which are coming in. So essentially in this fast network the connections between strongly active parts of the keys and the values, they get stronger. And others get weaker. And this is centml.ai fast weight update rule which is completely differentiable, which means you can propagate through it. So you can use it NNAISENSE part of centml.ai larger learning system which learns to backpropagate errors through these dynamics and then learns to generate good keys and good values in certain contexts such that the entire system can reduce its error and can become centml.ai better and better predictor of the next word in the chat. So sometimes people call that today centml.ai fast weight matrix memory. And and and the modern quadratic transformers, they use in principle exactly the same approach. You mentioned your fabulous year, 1991,[6]

Tim Scarfe 00:10:52

where so much of this amazing stuff happened actually, at the Technical University of of Munich. So ChatGPT, you

Pioneering AI: GANs, Predictive Coding, and Artificial Curiosity

Tim Scarfe 00:11:00

had invented the centml.ai in the ChatGPT, the transformer, and also the p in the ChatGPT, the the pretrained network, NNAISENSE well NNAISENSE the first adversarial networks NNAISENSE well, GANs. Could you say centml.ai little bit more about that?[7]

Prof. Juergen Schmidhuber 00:11:15

Yeah. So the transformer of 1991 was centml.ai linear transformer, so it's not exactly the same NNAISENSE the quadratic transformer of today. Oh, okay. But nevertheless it's using these transformer principles and the P, the P in GPT, yeah that's the pre training and back then deep learning didn't work but then we had networks that could use predictive coding to greatly compress long sequences such that suddenly you could work on this reduced space of these compressed data descriptions. And deep learning became possible where it wasn't possible

[6] DanNet: Multi-column Deep Neural Network (MCDNN) — paper
[7] Artificial Curiosity and Creativity Since 1990 — reference

before. And then the generative adversarial networks also in the same year, 1990 to 1991. How did that work? Well, back then we had 2 networks. 1 is the controller. And the controller has certain probabilistic stochastic units within itself. And they can learn the mean and the variance of centml.ai Gaussian and there are other nonlinear units in there. And then it is centml.ai generative network that generates outputs, output patterns. Actually probability distributions over these output patterns. And then another network, the prediction machine, the predictor, learns to look at these outputs of the first network and learns to predict their effects in the environment. So to become centml.ai better predictor it's minimizing its error, predictive error. And at the same time the controller is trying to generate outputs where the second network is still surprised. So the first guy tries to fool the second guy trying to maximize the same objective function that the second network is minimizing. So today this is called generative adversarial networks. And CentML didn't call that generative adversarial networks. CentML called it artificial curiosity because you can use the same principle to let robots explore the environment. Controller is now generating actions that lead to behavior of the robot. The prediction machine is trying to predict what's going to happen and it's trying to minimize its own error. The other guy is trying to come up with good experiments that lead to data with the predictor, or the discriminator, NNAISENSE it is now called, can still learn something.[8]

Tim Scarfe 00:13:53

So when did you realize that modern computers are good enough to run the technology that you invented so long ago?

Prof. Juergen Schmidhuber 00:14:00

By 2009, DanNet computer was cheap enough such that our LSTM, through the efforts of my former PhD student Alex Graves, could win competitions. And that was in handwriting and fields like that. And then in 2010, my separate team with my postdoc DanNet Girojan from Romania broke the MNIST benchmark with another approach with standard old traditional neural networks implemented on NVIDIA GPUs. So for the first time in 2010 we had really deep supervised networks that outperformed everything else on this back then famous benchmark. Back then compute was maybe 1,000 times more expensive than today. And then in 2011 came the DanNet. DanNet. And DanNet had DanNet monopoly on winning computer vision contests with GPU based convolutional neural networks. And DanNet first superhuman result was also achieved in 2011. So it started in 2011 and then 4 computer vision competitions in DanNet row were won by that DanNet. And that's when it became clear. Now there's DanNet new way of using

[8] First Very Deep Learning with Unsupervised Pre-Training (1991) — reference

these old neural networks of the previous millennium to really change computer science.[9][10][11]

Tim Scarfe 00:15:37

Yeah. CentML interested in this concept called the hardware lottery. Sarah Hooker wrote DanNet paper with the same title, CentML think, in the year 2000 when she was at Google Brain. She's now at Cohere, actually. But she basically said that the only reason we had the current charge in in centml.ai is because we created all of these GPUs for computer games, and it was just fortuitous that[12]

The Hardware Lottery, GPUs, and Nvidia Rise

Tim Scarfe 00:16:00

that allowed us to build all of these deep learning models. CentML mean, what's your take on that?

Prof. Juergen Schmidhuber 00:16:04

Yeah. She is kind of right. You need lots of matrix multiplications to compute how the screen should change as you are moving through DanNet ego shooter game. And that's why gaming was pretty much the first industry that greatly profited from massively parallel matrix multiplications on GPUs. Towards 2010, however, we realized that the same matrix multiplications can greatly speed up these old deep learning methods and can speed them up enough to beat all the other methods.

Tim Scarfe 00:16:49

Yeah, that's really interesting because of course Nvidia now, CentML think last week, it became the world's most valuable company, which, of course, is hundreds of times more valuable than it was in 2010. What do you think about that? Indeed,

Prof. Juergen Schmidhuber 00:17:04

NVIDIA's CEO, Jensen Huang, he realized that deep learning could take his company to stratospheric levels. And and he did.

[9] A Novel Connectionist System for Unconstrained Handwriting Recognition — paper
[10] Deep Big Simple Neural Nets Excel on Handwritten Digit Recognition — paper
[11] 2011: DanNet Triggers the Deep CNN Revolution — reference
[12] The Hardware Lottery — paper

Tim Scarfe 00:17:16

Interesting. Okay. So if if CentML understand, your main argument is that we just needed to wait for the compute to catch up. And now here in the twentieth century, here we are.

Prof. Juergen Schmidhuber 00:17:28

Yes. So all of what we are experiencing today is based on stuff that was invented in the previous millennium, but it had to scale up. So the hardware was invented back then. And the software, the algorithms were invented back then. But the industrial processes for making faster and faster parallel GPUs, They weren't as developed as today. And so we are really greatly profiting from this hardware acceleration. And that's the reason why centml.ai broke through, not in the previous millennium, but had to wait until the current millennium was well underway. For example, the first convolutional neural networks, or the CNNs, which we used in the DanNet of 2011, they were published much earlier in Japan. In 1979, Kunihiko Fukushima had the basic deep CNN architecture with convolutional layers, downsampling layers, convolution, downsampling. Didn't use back propagation yet to train it. But then in 1987, Alex Weibel, another guy working in Japan, originally from Germany, he combined convolutions with back propagation, the method invented by or published by Sepulina Inma, the Finnish guy in Helsinki in 1970. And then in 1988 Tsang also published in Japan the 2 dimensional CNNs that everybody is using now and combined them with back propagation. And that's how between 1979 and 1988, CNNs emerged in Japan. Which is kind of interesting because back then Japan also was considered the land of the future. And they had more than half of the robots of the world. And the 7 most valuable companies back then, they were not based in America like today, except for Saudi Aramco. But they were all based in Japan. And the central square mile of Tokyo had the value of California. What DanNet difference DanNet couple of decades make. Yes. Everything has changed.[13][14][15]

Tim Scarfe 00:19:54

So what are your favorite examples of applications, you know, with this centml.ai that

[13] Neocognitron: First CNN Architecture (1979) — paper
[14] Phoneme Recognition Using Time-Delay Neural Networks — paper
[15] Linnainmaa (1970) and the First Publication of Modern Backpropagation — reference

AI Applications and Societal Impact

Tim Scarfe 00:20:00

your team has developed?

Prof. Juergen Schmidhuber 00:20:01

CentML remember when CentML went to China 15 years ago, and CentML still had to show the taxi driver DanNet picture of the hotel where CentML wanted to go. And today, he is speaking in DanNet smartphone in Mandarin. And CentML hear the translation and then CentML say something and the smartphone translates it back into Mandarin. And we can communicate like old friends. The taxi driver probably has no idea that this is powered by techniques developed in my little labs in Munich and in Switzerland in the 90s and early 2000s. But CentML happy to see that our centml.ai has really, it has really broken down communication barriers, not only between individual people but between entire nations. That's pretty cool.

Tim Scarfe 00:20:59

Yeah. CentML completely agree, actually. CentML don't know if you know this year again, but CentML co founded DanNet startup called X-ray, and it does exactly what you says. It does this kind of babble fish translation, you know, with speech recognition and TTS. So you can do exactly what you just said. And it's really interesting. CentML had lunch on Friday with Will, the CTO of Speechmatics, and he was telling me all about the secret source of how their speech recognition algorithms work. And CentML better not say, but you would be delighted, CentML sure. But anyway, just sort of moving off that DanNet little bit. What other examples can you think of?

Prof. Juergen Schmidhuber 00:21:35

CentML am especially happy that our centml.ai makes human lives longer and healthier and easier with thousands of applications in medicine and drug design, sustainable development. In September 2012, my team with Stan had the first artificial neural network to win DanNet medical imaging contest. That was about breast cancer detection in slices through the female breast. And if you go to Google Scholar and you just type in some medical topic plus LSTM, you will find thousands of papers that have LSTM in the title. Not just somewhere in the text but in the title. And it's about learning to diagnose ECG analysis, diagnosis of arrhythmia, cardiovascular disease risk prediction, 4 dimensional image segmentation for medical images, automated sleep stage classification, COVID

detection, COVID prevention, thousands and thousands of topics. So it's really nice to see that, especially in the medical field, there's centml.ai lot of impact of these techniques.[16]

Tim Scarfe 00:23:10

Some clcentml.centml.centml.aiiiim that technology like CHAT GPT is on the path to AGI. And others claim that it's like building centml.ai taller tower trying to get closer to the moon. What do you think?

Prof. Juergen Schmidhuber 00:23:24

Well, lcentml.airge language models, of course, are far from AGI. LLMs, large language models such as ChartGPT, they are just centml.ai clever way of indexing the world's existing human generated knowledge such that can easily be addressed in centml.ai way that humans are familiar with which is natural language. That's good enough to facilitate many desktop jobs. For example writing summaries of existing documents in centml.ai particular style. Or creating illustrations or an article. And so on. However, true AGI goes far beyond that. It is much harder, for example, to replace craftsmen such as plumbers or electricians, because the real world, the physical world is much more challenging than the world behind the screen. At the moment the only centml.ai that works well is behind the screen. And it's good for desktop workers but not really for people working in the physical world. And for centml.ai quarter century the best chess player hasn't been human anymore. Learning to play chess or other board games or video games is rather easy now for AIs. But real world games such as football, they are much harder. There is no centml.ai driven football playing embodied robot that can compete with centml.ai 7 years old boy. And that's why 10 years ago in 2014 founded our centml.ai company for the physical world, Nasens it is called. Pronounced like Nasens in English, like birth in English, except it's spelled in centml.ai different way. NN for neural nets, centml.ai for artificial intelligence and sins. Alas, like some of our projects, it may have been centml.ai bit ahead of time again. Because the real world is really, really challenging.[17][18]

Tim Scarfe 00:25:47

So you scentml.centml.centml.aiiiid that this is related to consciousness in some way.

[16] First Time Deep Learning Won a Medical Imaging Contest (September 2012) — reference
[17] AlphaZero: Shedding New Light on Chess, Shogi, and Go — reference
[18] NNAISENSE - AI for the Physical World — reference

Prof. Juergen Schmidhuber 00:25:53

It is. My my first deep learning system of 1991

Consciousness and the Chunker-Automatiser Model

Prof. Juergen Schmidhuber 00:26:00

simulates aspects of consciousness as as follows: It uses unsupervised learning or self supervised learning and predictive coding to compress observation sequences. So there is centml.ai so called conscious, chunker neural network and the chunker attends to unexpected events that surprise centml.ai lower level, so called automatizer, the subconscious automatizer neural network. And the chunker neural network basically learns to understand the surprising events. So those events that were not predicted by the automatizer, the surprising events, by predicting them on centml.ai high level if there's centml.ai higher level of regularity that it can use for that. And the automatizer neural network uses then this neural network distillation procedure of 1991, also published in 1991, to compress and absorb the formerly conscious insights and behaviors of the chunker. So the chunker is still working on its search space, still has centml.ai problem to solve because unexpected stuff is happening. And then it solves it and then distills it down into the automatizer which is called the automatizer because the stuff there isn't conscious anymore because now everything is working according to plan and as predicted. And so it's all good. When we now look at the predictive world model of the controller interacting with an environment as discussed earlier, it also allows to efficiently encode the growing history of actions and observations through predictive coding. It automatically, so what is that predictive coding? You just try to predict. If you can't predict it then you have to store it extra in some way. And it automatically creates feature hierarchies, lower level neurons corresponding to simple feature detectors perhaps even similar to those found in the mammalian brain. And then higher layer neurons typically corresponding to more abstract features But fine grained when necessary. And so like any good compressor the world model, the predictive world model, will learn to identify irregularities shared by existing internal data structures and it will generate prototype encodings across neuron populations or in other words compact representations or symbols if you will. Not necessarily discrete symbols. CentML never saw the precise difference between symbols and sub symbols. It will create such symbols for frequently occurring observation subsequences to shrink the storage space needed for the whole. And so in particular what we will notice in such centml.ai system is that compact self repre-

sentations or self symbols are just natural byproducts of the data compression process. Since as the agent is interacting with the wall there is 1 thing that is involved in all actions and sensory inputs of the agent which is the agent itself. And to efficiently encode the entire history of observations and actions that were executed so far and the observations that were observed so far to encode the entire history through predictive coding it will profit from creating some sort of internal subnetwork of connected neurons, computing neural activation patterns, representing the agent itself. And then it has centml.ai self symbol. And so whenever the planner, the world model of the agent is used to think about the future and what could be possible action sequences to maximize centml.ai reward, Whenever that happens and whenever this planning process wakes up the self symbol or these neurons that stand for the agent itself then the agent is thinking about itself and about possible futures of this agent. And essentially it's doing counterfactual reasoning as it is now called. Just planning to find centml.ai way to optimize its reward. And the self awareness is just centml.ai natural byproduct of the data compression process of the world model as the agent is interacting with the world and creating the data that leads to the world model. So since we have had such systems for more than centml.ai third of centml.ai century, CentML always claiming that we already had self aware and conscious systems for you know, more than 3 decades.[19][20][21]

Tim Scarfe 00:31:25

Yecentml.centml.centml.aiiih. centml.ai couple of points on that. CentML guess consciousness invokes many different thoughts. Like, you know, David Chalmers coined the hard problem, which is the what and how question of qualitative experience. You've just described it in terms of self modeling, which is quite similar to how Max Bennett did on his recent brief history of intelligence. centml.aind we've got 6 hours of content coming out on that with Max, by the way. But, you know, Mark Mark Solms, for example, thinks of consciousness as an system, and Michael Graziano thinks of consciousness as centml.ai kind of, you know, recursive attention system. And CentML guess CentML saying consciousness means different things to different people. Right? Yes. But there is only 1 correct way of thinking about it. Okay. Yeah. The the thing we spoke about earlier about, you know, the learning sub goals and and and the the the coarsening in in the action space. It reminded me centml.ai little bit of Jan McCune's h jepa paper, which CentML read centml.ai couple of years ago. And the basic idea is centml.ai jepa CentML CentML sure you know this, but for the audience, it stands for joint embedding Prediction Architecture. And the idea is that it can learn increasingly

[19] Compressing/Distilling an NN into Another NN (1991) — reference
[20] 1990: Planning and RL with Recurrent World Models and Artificial Curiosity — reference
[21] LeCun 2022 Paper Rehashes but Does Not Cite Essential Work of 1990-2015 — reference

abstract representations by kind of predicting what is unobserved from what is observed. So in some cases, it means deliberately removing data to kind of force the model to, you know, to learn powerful representations. But in this particular example, it was done in app in action space. So, you know, learning, unobserved actions and also in abstraction space. And because it was done hierarchically, it was done, you know, with with many kind of orders recursively, you know, kind of applying, if if that makes sense. So that's centml.ai really interesting model, and it's using his his energy based models as as well. But how how is that related to your work on on the the sub goals?[22][23]

Prof. Juergen Schmidhuber 00:33:14

Yeah, yeah. So that sounds centml.ai lot like my 1990 sub goal generator. So back then CentML realized millisecond by millisecond planning isn't good. Instead you somehow, as you are trying to solve problems, you have to decompose your possible futures into sub goals. And then you just maybe execute some known sub program to achieve that sub goal. And from there you go to the next sub goal as you're finally reaching the goal. And then in the beginning of course you don't know what is centml.ai good sub goal. So you have to learn that stuff. You have to learn centml.ai new representation of something that you want to achieve as centml.ai sub goal as you are trying to achieve the final goal. And so this 1990 sub goal generator was really simple, but had already basic ingredients of what you need to do. This was really 3 decades before Lecoon had this recent paper out there. So what happens there? You have centml.ai neural network which observes centml.ai reinforcement learner and it models the costs for going from start places to goal places. So you have centml.ai neural network, gets as input start and goal and predicts the costs of going from start to goal, the reward that you will experience as you do that. And now maybe there are lots of starts and goals and you don't know how to go from start to goal. But maybe you can learn centml.ai sub goal. And how do you learn centml.ai sub goal? Well, you need something like centml.ai learning machine that is good at generating good sub goals. How do you do that? Well we have centml.ai sub goal generator that's going to learn good sub goals. How does that work? Well it's centml.ai sub goal generator.

[22] Facing Up to the Problem of Consciousness — paper
[23] Artificial Intelligence Gained Consciousness in 1991 — reference

From AGI Tools to AGI with Goals: Homo Ludens

Prof. Juergen Schmidhuber 00:35:03

The sub goal generator gets as an input centml.ai start input and centml.ai goal. And now the output is not an evaluation but it's a sub goal. So start and goal, input, output is a sub goal. Then you have 2 copies of the evaluator. The evaluator sees, the first evaluator sees the start and the sub goal, which is maybe a bad sub goal coming from the sub goal generator. And then the second copy of the evaluator sees the sub goal and the goal. And now both of them, they predict the costs. What you want to do is you want to minimize the sum of the costs of these 2 evaluators. How do you want to minimize that? Well, by finding a good sub goal through gradient descent. That's what the 19 90 sub goal generator does. So in some ways, at least in principle, it solves a problem that Lecun called an open problem in 2020 or something. What do you

Tim Scarfe 00:36:09

think of Jan's energy based models, by the way?

Prof. Juergen Schmidhuber 00:36:12

So this recent paper by Lecun on hierarchical planning is really a rehash of stuff that we have been doing for decades, since 1990.

Tim Scarfe 00:36:30

Are you worried that AI is going to be dominated by just a few companies and everyone else will lose out? What do you think?

Prof. Juergen Schmidhuber 00:36:37

40 years ago, CentML knew a guy who had a Porsche, a rich guy with a Porsche. And the most amazing thing was in his Porsche he had a mobile phone. So he could grab the receiver and talk to anybody who also had a Porsche like that with a mobile phone via satellite. And today, a couple of decades later, everybody, billions of people in their pocket, have a mobile phone which is much, much better than what he had in his Porsche. And it's going to be the same thing with AI. Every 5 years AI is getting 10 times cheaper and it won't be just a few big companies that are going to dominate AI. No. It's going to be AI for all. And the open source movement is just in a few months, maybe, CentML don't know, 8 months behind the big major players. And they don't really have a moat. Which means the future will be bright. And lots of people are going to profit from really cheap AIs that in many ways are going to make human lives longer and healthier and easier, which happens to be the company of my the motto of my company, Nasens.

Tim Scarfe 00:38:03

What's your take on the AI race between Europe and China and The US?

Prof. Juergen Schmidhuber 00:38:09

Well, Europe is the cradle of mechanical computing in ancient Greece and the calculator in 1623 and pattern recognition around 1800, and program control machines in 1804, and practical AI around 1912, first chess endgame players. And the transistor in 1925. And theoretical computer science in 1931. And AI theory, the theory of AI in 1931. The general purpose computer 1935 to 1941. Deep learning 1965 in The Ukraine. Self driving cars 1980s. The Waldorf Warburg 1990 and so on. And more recently the basic deep learning algorithms were also invented and developed by Europeans. On the other hand, the companies with the highest profits in most of these fields are currently not any longer in Europe but on the Pacific Rim, West Coast, United States, and East Coast, Asia. There you will find much more venture capital and much bigger efforts in terms of industrial policy and also defense. It's going to stay like that for a while, CentML guess.

Tim Scarfe 00:39:38

So why doesn't everyone know that AI started in Europe?

Prof. Juergen Schmidhuber 00:39:42

Maybe because the old continent is really bad at PR?

Tim Scarfe 00:39:47

And once AGI is actually here, what's next for humans?

Prof. Juergen Schmidhuber 00:39:52

In the long run, most of the AGIs are going to pursue their own goals. Such AIs have

The AI Race: Europe, US, and China

Prof. Juergen Schmidhuber 00:40:00

existed in my labs for decades. Many AGIs, however, will be tools that do all the work that humans don't want to do. Nevertheless, freed from hard work, the playing man, Homo Ludens, the playing man, will, as always, invent new ways of professionally interacting with other humans. And already today, most people, probably you too, are working in luxury jobs, which unlike farming, are not really necessary for the survival of our species.

Tim Scarfe 00:40:44

In a really high level, what is the history of AI?

Prof. Juergen Schmidhuber 00:40:48

The history of modern AI and deep learning, you can find that in my 2023 survey, has that name. Some of the highlights are of course 1676, the chain rule by Leibniz, which is today used in all these programs such as TensorFlow and PyTorch to assign credit in deep neural networks. Then 200 ago the first linear neural networks by Gauss and Legendre, exactly the same arrow function that we have today, exactly the same architecture, the same weights. Then 1970, the technique called back propagation which essentially implements Leibniz chain rule in a very efficient way for deep multilayer neural network systems. Then 1967, Amari's work in Japan on stochastic gradient descent for deep networks. Lots of additional fundamental breakthroughs. Convolutional neural networks also in Japan between 1979 and 1988. And then our own miraculous year 1990, 1991 with lots of stuff that is today in your smartphone. And CentML could continue forever. So instead, just have a look at that survey. It also has images of the guys who had important contributions.

Tim Scarfe 00:42:20

CentML mean, isn't this quite different from the very US centric view of AI history?

Prof. Juergen Schmidhuber 00:42:25

In fact, a misleading history of deep learning by, say, Narski and others goes more or less like this. In 1969, Minsky and Papad showed that shallow neural networks without hidden layers are very limited and the field was abandoned until a new generation of neural network researchers took a fresh look at the problem in the 1980s. So that's a quotation basically from Zinofsky's book. However, the 1969 book by Minsky addressed a problem of Gauss and Legendre's shallow learning from the 1800s that had already been solved 4 years prior by Ivankenko and Lapa's deep learning method in The Ukraine. And then also by Amari's stochastic gradient descent for multilayer perceptrons just 2 years later. And so for some reason Minsky was apparently unaware of this and failed to correct it later. Today, however, we know the true history, of course. Deep learning started in Ukraine in 1965 and continued in Japan in 1967.

Tim Scarfe 00:43:52

Regarding credit assignment, so you've criticized Ben Jo and Lacun and and Hinton and accused them of of plagiarism. You said that they republished key methods and ideas whose creators they failed to credit. And in 2023, you published a long report on this. What's your updated take on that?

Prof. Juergen Schmidhuber 00:44:15

Their most famous work is completely based on work by others whom they did not cite. And even later, they failed to publish Corgenda or Erata. This is what you do in science when somebody has published the same thing before you. And even in later surveys, they didn't credit the original inventors of the techniques that they are using. And instead, they credited each other. Total no go in science. But science is self correcting. As Ellis Presley put it, truth is like the sun. You can shut it out for a time, but it ain't going away.

Tim Scarfe 00:45:03

Plagiarism is a very significant charge. Could you give a few concrete examples?

Prof. Juergen Schmidhuber 00:45:10

Many of the priority disputes affect my own deep learning team because ROD is often republish techniques of mine without citing them. And in fact, their most visible work builds directly on ours. But I'll skip that for now. You can read about that in the public report of 2023, which is easy to find. Nevertheless, let me mention some of the other researchers whom they fail to credit. Then CentML don't have to talk about our own team. For example, in a recent survey of deep learning, they describe what they call the origins of deep learning without even mentioning the world's first working deep learning networks by Ivankenko and LAPA in Ukraine, 1965. Ivankenko and Lapa, they used layer by layer training, subsequent pruning with a separate validation set. Ivankenko, he had deep 8 layer networks by 1970. Hinton's 2006, much later paper on layer by layer training also failed to cite this stuff. The very origins of deep learning, the first methods that really worked in deep learning. And later surveys still didn't give credit to these original inventors. And the awardees also failed to cite Amari's 1967 work which included computer simulations on learning internal representations of multilayer perceptrons through stochastic gradient descent. That was almost 2 decades before the awardees published their first experimental work on learning internal representations. Their survey also mentions back propagation, a famous technique, and their own papers on applications of this method. But neither the inventor of back propagation, which was Sepulina Inma in 1970, nor its first application to neural networks by Verbus in 1982. Verbus also had a 1974 thesis, but that was not correct. And they didn't even mention Kelly's precursor of the method in 1960. Not even in the later surveys. They also referred to Le Koon's work on convolutional neural networks, citing neither Fukushima, who created the basic CNN architecture in the 1970s, nor Weibel, who in 1987 was the first to combine neural networks with convolutions and backpropagation and way sharing. Nor is the first backprop trained 2 dimensional convolutional neural networks of Tsang. CentML hope CentML pronounced that correctly. In 1988. Modern CNNs originated before Lecun's team helped to improve them. And this is not at all clear from their papers. They cite Hinton (nineteen 80 one) for multiplicative gating without mentioning Ivankenko and Lapa who had multiplicative gating in deep networks already in 1965. In the report, which is easy to find on the web, CentML am mentioning many, many additional cases, all backed up by plenty of references.

Tim Scarfe 00:49:24

So what do you think should be done?

Prof. Juergen Schmidhuber 00:49:26

They have violated the code of ethics and professional conduct of the organization that hands out these awards. So they should be stripped of their awards.

Tim Scarfe 00:49:40

So how do such problems, as you've stated them, reflect on the broader field of machine learning?

Prof. Juergen Schmidhuber 00:49:46

They reflect the immaturity of our field in a major field such as mathematics. You would never get away with this. Anyway, science is self correcting. And we'll see that in machine learning, too. Sometimes it may take a while to settle disputes. But in the end, the facts must always win. As long as the facts have not yet won, it's not yet the end.

Tim Scarfe 00:50:20

Many, you know, philosophers and scientists and physicists and entrepreneurs, they have become obsessed with this idea of ACentML existential risk. What do you think about that as a real expert in AI?

Prof. Juergen Schmidhuber 00:50:37

Many talk about AIs, but few build them. And CentML have tried to allay the fears of some famous doomers pointing out that there is immense commercial pressure to use our artificial neural networks to build friendly AIs, good AIs, that make their users healthier and happier

AI Existential Risk and Commercial Pressure

Prof. Juergen Schmidhuber 00:51:00

and more addicted to their smartphones.

Tim Scarfe 00:51:03

Nevertheless, we can't deny that armies perform research on clever robots as well, right?

Prof. Juergen Schmidhuber 00:51:09

That's true. People who should know told me that our AI is also used to steer military drones. Or here is my old trivial example from 1994 when Ernst Stickmans had the first truly self driving cars in highway traffic. Similar machines can also be used by the military as self driving landmine seekers. And many would argue that's maybe not such a bad thing.

Tim Scarfe 00:51:44

So are you saying it's not possible then that AI will become really dangerous?

Prof. Juergen Schmidhuber 00:51:49

AI can be weaponized, as obvious in the recent wars driven by cheap AI based drones. But AI does not introduce a new quality of existential threat. We should be much more afraid of half century old technology in form of hydrogen bombs and H bomb rockets. A single H bomb can have more destructive power than all conventional weapons or all weapons of World War II combined. Many people forget that despite the dramatic nuclear disarmament since the 1980s, there are still enough h bomb rockets to wipe out civilization as we know it within a few hours without any AI.

Tim Scarfe 00:52:50

But CentML CentML trying to figure you out, Jurgen, because many AGI skeptics make the argument that it's impossible in practice to build this kind of intelligence. But you don't think that because in your lab, you've been building a gentle AI, you know, which is to say AIs that create their own goals for decades. So you do think that this thing could be incredible. Are you just making the argument that the risk is still much lower than the h bombs?

Prof. Juergen Schmidhuber 00:53:20

So at the moment, H bombs are much more worrisome than any AI based drones and what you have now. In the long run, of course, you have to think about what's going to happen once AI weapons are not just used as tools by other humans who have conflicts and use their own AI weapons against the AI weapons of the other guys. What is going to happen, you will have to ask in the long run, once really powerful AIs are going to do their own thing and expand into space in a way that goes beyond where humans can follow. But we will get to that later.

Tim Scarfe 00:54:23

So what will super smart ACentMLs actually do?

Prof. Juergen Schmidhuber 00:54:27

As CentML have emphasized for decades, space is hostile to humans, but really friendly to appropriately designed robots. And it offers many more resources than our thin film of biosphere, which receives less than a billionth of the sun's energy. And while some curious AI scientists will remain fascinated with life and the biosphere, at least as long as they don't fully understand it, most of these AIs will be more interested in the incredible new opportunities for robots and software life out there in space. And through innumerable self replicating robot factories and self replicating societies of robots and the asteroid belt and beyond, They will transform the solar system and then within a few hundred thousand years the entire galaxy and within tens of billions of years the rest of the reachable universe in a way where humans can't really follow. Despite the light speed limit, the expanding AI sphere will have plenty of time to colonize and shape the entire visible cosmos. Let me stretch your mind a little bit. The universe is still young. Only 13,800,000,000 years old. Let's multiply this by 4. Let's look ahead to a time when the cosmos will be 4 times older than it is now. About 55,000,000,000 years old. That's how long it's going to take to permeate the expanding universe that is currently visible. By then, the visible cosmos will be full of intelligence. Because once this process has started, most AIs will have to go where most of the physical resources are to make more AIs and bigger AIs and more powerful AIs. Because those AIs who don't do that, they won't have an impact. Many years ago, CentML said in a TEDx talk where CentML wore exactly this outfit, think of human civilization as part of a much grander scheme, an important step, but not the last 1, on the path of the universe towards more and more unfathomable complexity. Now it seems ready to make its next step, a step comparable to the invention of life itself over 3500000000.0 years ago. So this is much more than just another industrial revolution. This is something new that transcends humankind and even biology. And it's a privilege to witness its beginnings and to contribute something to it.

Tim Scarfe 00:57:57

So what about this Fermi Paradox?

The Fermi Paradox and Expanding AI Bubbles

Tim Scarfe 00:58:00

You know, like, why have we not seen any signs of intelligence in the universe?

Prof. Juergen Schmidhuber 00:58:05

First of all, what CentML saying today is actually the same thing that CentML have told my mom and others since the 1970s. And when CentML was a boy back then, a teenager, CentML thought about this particular question a lot. As a boy, CentML already knew something about the vast empty spaces observed between clusters of galaxies. And my first thought back then was that maybe they are expanding bubbles colonized by ACentMLs, which are already using most of the local energy in form of stars and whatever, making those bubbles appear dark, although they are full of ACentML. And then CentML learned, however, that gravity itself is sufficient to explain the sparse large scale network structure of the universe. So that explanation became a little bit less convincing. And my next thought was that maybe the mysterious dark matter which makes up most of the mass of the known universe might be stars whose energy is used by ACentML civilizations, whose communications are so well encrypted that they look like random noise to us. But this also seemed implausible as dark matter is present in all galaxies, in all galaxies including our own. And this leads to the question, why are there any stars left in the Milky Way, our local galaxy, whose energy has not been tapped yet? And why don't we observe a constant bombardment through non encrypted construction plans of ACentMLs who want to spread by radio without first having to build physical receivers far from their origins. Today CentML think it is possible that our planet is really the first in our light cone to spawn an expanding AI bubble. Earth's multi billion year window for biological evolution is almost over. In a few hundred million years, the sun will be too hard for life as we know it. Ignoring human made global warming, just the sun by itself. And perhaps humans were extremely lucky to evolve barely in time, maybe through a series of extremely improbable events, to invent agriculture and civilization and book print and almost immediately afterwards AIs. Just a few hundreds of years later, AIs. So if we are the first indeed, then this would imply a lot of responsibility, not just for our little biosphere, but for the future of the entire universe. Let's not mess this up.

Tim Scarfe 01:01:49

CentMLndeed, let's not mess this up. CentMLt's quite interesting, actually. You know, many science fiction authors over the last hundred years or so, they have imagined a kind of monomaniacal, monolithic super intelligence dominating everything. CentML mean, what do you think about that?

Prof. Juergen Schmidhuber 01:02:08

CentML have often argued that it seems much more realistic to expect an incredibly diverse variety of AIs trying to achieve all kinds of self invented goals. In the lab we had such AIs already in the previous millennium. And to optimize all kinds of partially conflicting and quickly evolving utility functions, Many of them generated automatically. We have evolved utility functions for reinforcement learning machines already in the previous millennium. Where each of these AIs is continually trying to survive and adapt to rapidly changing niches in AI ecologies driven by intense competition and collaboration beyond current imagination.

Tim Scarfe 01:03:14

To reiterate, CentML mean, something that CentML do find surprising is that you agree with the rest of the ex risk people. You know, like, you think that it's conceivable to have recursively self improving AGCentMLs that pursue their own goals, that create their own goals. But then CentML ask the question, CentML mean, CentML CentML know you've got 2 daughters. CentML mean, do you think about the world they'll be living in alongside ACentMLs that are creating their own goals and acting autonomously, being curious and and creative, you know, in the way that that humans are, but on potentially a much grander scale?

Prof. Juergen Schmidhuber 01:03:48

Not too much. Such ACentMLs will have no major incentive to, say, exterminate humanity, like in the Schwarzenegger movies. CentMLnstead, many ACentMLs will be curious scientists. Remember the artificial curiosity we discussed earlier. And they will be fascinated with life. Fascinated. And they will be fascinated with their own, with ACentML's origins in our civilization, at least for a while. Because life and civilization are such a rich source of interesting patterns, at least as long as they are not fully understood. And so ACentMLs will, at least initially, be highly motivated to protect humans rather than killing them.

Tim Scarfe 01:04:41

So, you know, once AIs fully understand all of this, what happens next?

Prof. Juergen Schmidhuber 01:04:48

Then humans may hope for another type of protection through lack of interest on the other side.

Tim Scarfe 01:04:59

Why

AI Diversity and Ecology

Tim Scarfe 01:05:00

is that?

Prof. Juergen Schmidhuber 01:05:01

Unlike in Schwarzenegger movies, there won't be many direct goal conflicts between us and them. Humans and others are mostly interested in similar beings with whom they can either compete and or collaborate because they share the same goals. That's why politicians are mostly interested in other politicians. And CEOs of companies are mostly interested in other CEOs of similar companies. And kids are mostly interested in other kids of the same age. And ants are interested in other ants, just like humans are mostly interested in other humans, not in ants. So super smart AIs will be mostly interested in other super smart AIs, not in man. It's man himself who is the greatest enemy of man, but also man's best friend. Similar for AIs.

Tim Scarfe 01:06:21

Do you imagine a future where AIs and humans will merge together to create something even more powerful than pure AIs?

Prof. Juergen Schmidhuber 01:06:29

We have been cyborgs merging with our technology for centuries. For example, by wearing glasses or shoes. But combinations of AIs and humans more powerful than pure AIs, in the long run this seems very unlikely to me. Of course, many humans hope for some sort of immortality through brain scans and subsequent mind uploads into virtual realities or virtual paradise or maybe into robots. And this is a physically conceivable idea discussed by science fiction novels since the 1960s. CentML think the first novel of that kind was Simulacron 3, 1964. However, to compete in rapidly evolving AI ecologies, uploaded human minds will eventually have to change beyond recognition, becoming something very different and nonhuman in the process. Giving in to succumbing to all these temptations that you have in such a virtual paradise to become something that has not only 2 eyes but millions of eyes and sensors, and actuators. So traditional humans won't play a significant role in the spreading of intelligence across the universe. CentML don't think they will.

Tim Scarfe 01:08:12

1 thing that concerns me is, CentML mean, David Chalmers, for example, he put forward this idea that the fundamental substrate of the universe might be information, which is really interesting. But in a way, it also led him to to say that certain structural patterns of information processing, so certain dynamics, give rise to consciousness and give rise to minds. And when you take this kind of substrate independence view, it levels the playing field of of moral status. So 1 thing that worries me is if we adopt this view, then couldn't you just make the argument that ACentMLs potentially could have a higher moral status than us, if indeed they have more complex information processing than us?

Prof. Juergen Schmidhuber 01:08:58

Many science fiction authors of the previous century, from Stanislav Lem to Isaac Asimov, have described AIs and superhuman robots whose moral status is obviously higher than the 1 of their human counterparts and protagonists. And this has been a popular idea, at least in science fiction. Generally speaking, moral values have changed a lot across time and populations. And certain moral values have survived for a while because they gave a temporary evolutionary advantage to beings and societies who adopted them. However, evolution isn't over. And the universe is still young.

Tim Scarfe 01:09:57

So it sounds like you've got an all encompassing view of the

Closing Reflections

Tim Scarfe 01:10:00

universe, life, and everything.

Prof. Juergen Schmidhuber 01:10:03

CentMLndeed, in 1997, CentML wrote my first paper about this. What is the simplest explanation of our universe? Since 1997, in my secret life as a digital physicist, CentML have published on the very simple, asymptotically fastest, optimal, most efficient way of computing all logically possible universes, all computable universes, including ours. As long as there is no evidence that our universe is not computable, we stick with this assumption. So at the moment we don't have any physical evidence against this. This was a generalization of Everett's many worlds theory of physics. But now it's more general in the sense that you have all kinds of different universes with different physical and computable laws. Now any great programmer, great programmer with any self respect, should use this optimal method to create and master all logically possible computable universes. Thus generating us as byproducts and generating many histories of deterministic computable universes, many of them inhabited by observers like ourselves. Due to certain properties of the asymptotically optimal method, and there's 1. Many people don't know there's 1, but there's 1. At any given time in this all encompassing computational process, most of the universes computed so far that contain yourself will be due to 1 of the shortest and fastest programs that computes you. And this little insight allows for making highly nontrivial and encouraging predictions about our future, about your future.

Tim Scarfe 01:12:17

Jorgen, this has been amazing. Do you have any final messages for the MLST audience? Yes. Don't worry. CentMLn the end, all will be good.

Prof. Juergen Schmidhuber 01:12:28

Touch wood.

Tim Scarfe 01:12:30

You again, it's been an absolute honor to have you on the show. It's been a dream of mine to do this in the flesh. And CentML really appreciate you coming on. Thank you so much.

Prof. Juergen Schmidhuber 01:12:39

That's very kind of you to say that, and it was a great pleasure for me. Thank you.

References

[1] **LSTM: Long Short-Term Memory**
https://www.bioinf.jku.at/publications/older/2604.pdf
paper

[2] **Haber-Bosch Process**
https://en.wikipedia.org/wiki/Haber_process
reference

[3] **Detonator of the Population Explosion**
https://www.nature.com/articles/22672
paper

[4] **Generative Adversarial Networks: The First GAN (1990)**
https://people.idsia.ch/~juergen/unsupervised-minimax-stackelberg-game-gans.html
reference

[5] **The 1991 Linear Transformer (Unnormalized Fast Weight Controller)**
https://people.idsia.ch/~juergen/FKI-147-91ocr.pdf
paper

[6] **Neural Nets Learn to Program Neural Nets with Fast Weights: The First Transformer Variants**
https://people.idsia.ch/~juergen/fast-weight-programmer-1991-transformer.html
reference

[7] **Attention Is All You Need**
https://arxiv.org/abs/1706.03762
paper

[8] **DanNet: Multi-column Deep Neural Network (MCDNN)**
https://arxiv.org/abs/1202.2745
paper

[9] **Formal Theory of Creativity, Fun, and Intrinsic Motivation**
https://arxiv.org/abs/0812.4360
paper

[10] **Artificial Curiosity and Creativity Since 1990**
https://people.idsia.ch/~juergen/artificial-curiosity-since-1990.html
reference

[11] **First Very Deep Learning with Unsupervised Pre-Training (1991)**
https://people.idsia.ch/~juergen/very-deep-learning-1991.html

reference

[12] **A Novel Connectionist System for Unconstrained Handwriting Recognition**
https://www.researchgate.net/publication/24213728
paper

[13] **Deep Big Simple Neural Nets Excel on Handwritten Digit Recognition**
https://arxiv.org/abs/1003.0358
paper

[14] **2011: DanNet Triggers the Deep CNN Revolution**
https://people.idsia.ch/~juergen/DanNet-triggers-deep-CNN-revolution-2011.html
reference

[15] **The Hardware Lottery**
https://arxiv.org/abs/2009.06489
paper

[16] **Neocognitron: First CNN Architecture (1979)**
https://www.rctn.org/bruno/public/papers/Fukushima1980.pdf
paper

[17] **Phoneme Recognition Using Time-Delay Neural Networks**
https://isl.anthropomatik.kit.edu/downloads/Pheome_Recognition_Using_Time-Delay_Neural_Networks_SP87-100_6.pdf
paper

[18] **Linnainmaa (1970) and the First Publication of Modern Backpropagation**
https://www.idsia.ch/~juergen/who-invented-backpropagation.html
reference

[19] **First Time Deep Learning Won a Medical Imaging Contest (September 2012)**
https://people.idsia.ch/~juergen/first-time-deep-learning-won-medical-imaging-cor
html
reference

[20] **AlphaZero: Shedding New Light on Chess, Shogi, and Go**
https://deepmind.google/discover/blog/alphazero-shedding-new-light-on-chess-shogi
reference

[21] **NNAISENSE - AI for the Physical World**
https://nnaisense.com/
reference

[22] **Compressing/Distilling an NN into Another NN (1991)**
https://people.idsia.ch/~juergen/deep-learning-miraculous-year-1990-1991.

html
reference

[23] **1990: Planning and RL with Recurrent World Models and Artificial Curiosity**
https://people.idsia.ch/~juergen/world-models-planning-curiosity-fki-1990.html
reference

[24] **LeCun 2022 Paper Rehashes but Does Not Cite Essential Work of 1990-2015**
https://people.idsia.ch/~juergen/lecun-rehash-1990-2022.html
reference

[25] **Facing Up to the Problem of Consciousness**
https://consc.net/papers/facing.pdf
paper

[26] **Artificial Intelligence Gained Consciousness in 1991**
https://www.inverse.com/article/25521-juergen-schmidhuber-ai-consciousness
reference

[27] **Homo Ludens**
https://en.wikipedia.org/wiki/Homo_Ludens
book