

Prof. Dr. Jürgen Schmidhuber

IDSIA – Institute Dalle Molle for Artificial Intelligence

What was the most influential innovation of the 20th century?

In 1999, the journal *Nature* made a list of the most influential inventions of the 20th century. Number 1 was the invention that made the century stand out among all centuries, by „detonating the population explosion“ (V. Smil), from 1.6 billion people in 1900 to soon 10 billion. This invention was the Haber process, which extracts nitrogen from thin air, to make artificial fertilizer. Without it, 1 in 2 persons would not even exist. Soon this ratio will be 2 in 3. Billions and billions would never have lived without it. Nothing else has had so much existential impact. (And nothing in the past 2 billion years had such an effect on the global nitrogen cycle.)

How about the 21st century?

The Grand Theme of the 21st century is even grander: True Artificial Intelligence (AI). AIs will learn to do almost everything that humans can do, and more. There will be an AI explosion, and the human population explosion will pale in comparison.

What kind of computational device should we use to build practical AIs?

Physics dictates that future efficient computational hardware will look a lot like a brain-like recurrent neural network (RNN), a general purpose computer with many processors packed in a compact 3-dimensional volume, connected by many short and few long wires, to minimize communication costs¹. Your cortex has over 10 billion neurons, each connected to 10,000 other neurons on average. Some are input neurons that feed the rest with data (sound, vision, tactile, pain, hunger). Others are output neurons that move muscles. Most are hidden in between, where thinking takes place. All learn by changing the connection strengths, which determine how strongly neurons influence each

other, and which seem to encode all your lifelong experience. Same for our artificial RNNs, which learn better than previous methods to recognize speech or handwriting or video, minimize pain, maximize pleasure, drive simulated cars, etc.

What do you see as the nearer-term future in AI advancements and where will this lead?

Kids and even certain little animals are still smarter than our best self-learning robots. But I think that within not so many years we'll be able to build an Neural Network-based AI (an NNAI) that incrementally learns to become at least as smart as a little animal, curiously and creatively and continually learning to plan and reason and decompose a wide variety of problems into quickly solvable (or already solved) subproblems, in a very general way.

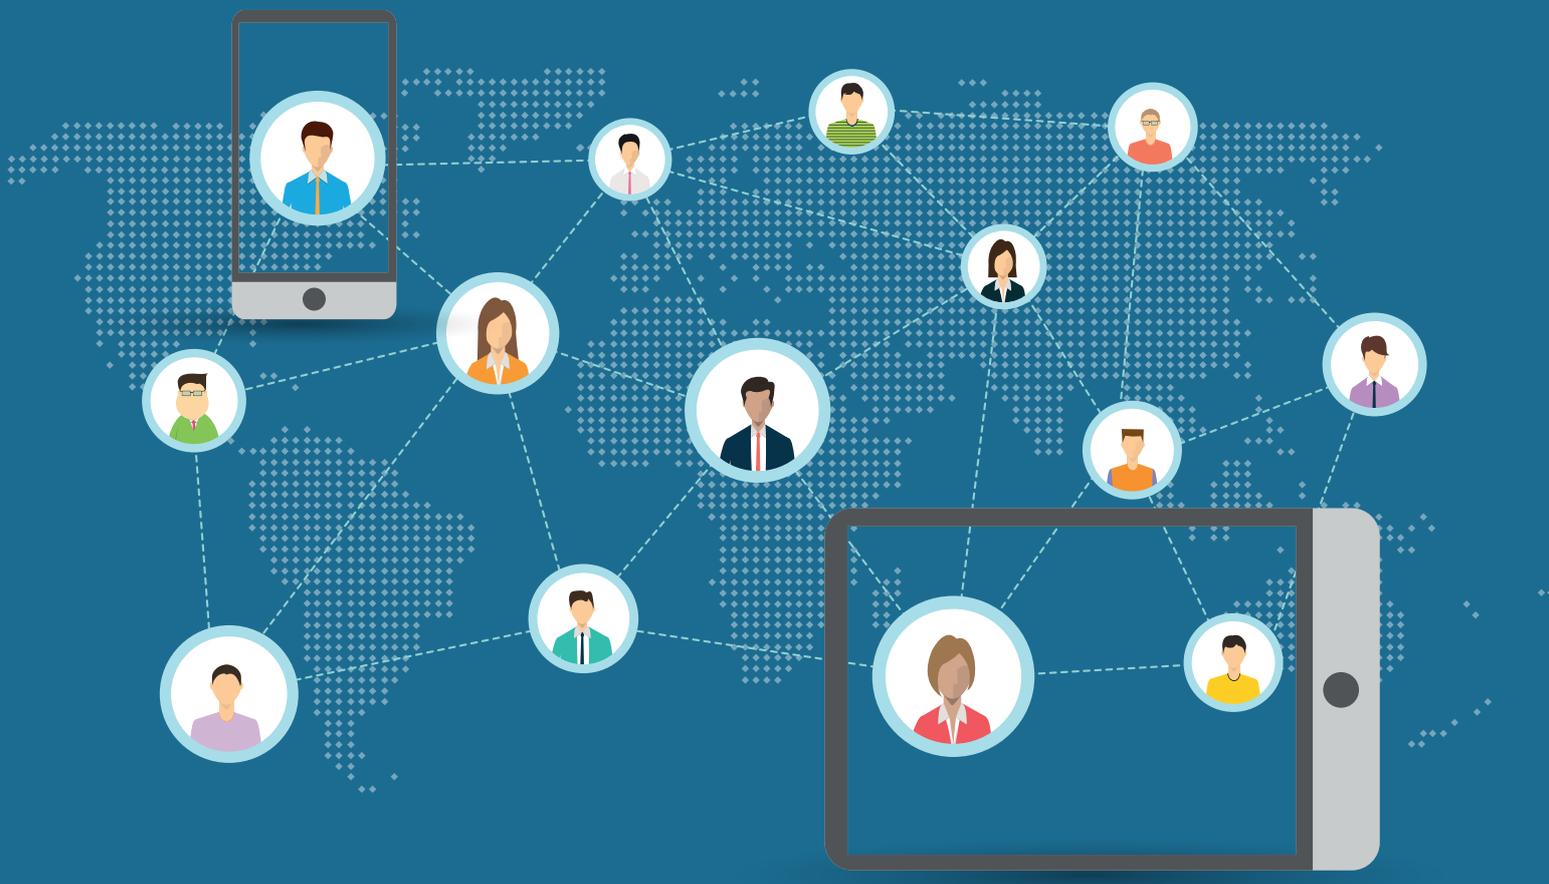
Once animal-level AI has been achieved, the next step towards human-level AI may be small: it took billions of years to evolve smart animals, but only a few millions of years on top of that to evolve humans. Technological evolution is much faster than biological evolution, because dead ends are weeded out much faster. That is, once we have animal-level AI, a few years or decades later we may have human-level AI, with truly limitless applications, and every business will change, and all of civilization will change, and everything will change.

Which will be the near-term social implications of AI?

Smart robots and/or their owners will have to pay sufficient taxes to prevent social revolutions. What remains to be done for humans? Freed from hard work, „Homo Ludens“ (the playing man) will (as always) invent new ways of

¹ Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85-117. <http://people.idsia.ch/~juergen/deep-learning-overview.html> Short version at scholarpedia: http://www.scholarpedia.org/article/Deep_Learning





JE | journée de l'économie

L u x e m b o u r g

The Company of the Future

2 March 2017

Register now on: www.jecolux.lu



@JEcolux | #Jelux

Under the patronage of:



LE GOUVERNEMENT
DU GRAND-DUCHÉ DE LUXEMBOURG
Ministère de l'Économie

Conference organised by:



LE GOUVERNEMENT
DU GRAND-DUCHÉ DE LUXEMBOURG
Ministère de l'Économie



Fedil
Business Federation
Luxembourg

In collaboration with:



Media Partner:

PAPERJAM
Business zu Lëtzebuerg

professionally interacting with other humans. 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. Machines are much faster than Usain Bolt, but he still can make hundreds of millions by defeating other humans on the race track. In South Korea, the most wired country, new jobs emerged, such as the professional video game player. Remarkably, countries with many robots per capita (Japan, Germany, Korea, Switzerland) have relatively low unemployment rates. My old statement from the 1980s is still valid: It's easy to predict which jobs will disappear, but hard to predict which new jobs will be created.

The public perception of AI seems to be more negative than positive and some people say AI is potentially harmful to society. What do you think about those discussions, and how do you explain this in a way that maybe has a more positive beneficial impact for society ?

Many talk about AIs. Few build them. Prominent entrepreneurs, philosophers, physicists and others with not so much AI expertise have recently warned of the dangers of AI. I have tried to allay their fears, pointing out that there is immense commercial pressure to use artificial neural networks such as our LSTM (Long Short-Term Memory) to build friendly AIs that make their users healthier and happier. Nevertheless, one cannot deny that armies use clever robots, too. Here is my old trivial example from 1994 when Ernst Dickmanns had the first truly self-driving cars in highway traffic: similar machines can also be used by the military as self-driving land mine seekers.

We should be much more afraid, however, of half century-old tech in form of H-bomb rockets. A single H-bomb can have more destructive power than all

conventional weapons (or all weapons of WW-II) combined. Many forgot that despite the dramatic nuclear disarmament since the 1980s, there are still enough H-bomb rockets to wipe out civilization within a few hours, without any AI. AI does not introduce a new quality of existential threat.

Should our children and young people worry about future AIs pursuing their own goals, being curious and creative, in a way similar to the way humans and other mammals are creative, but on a much grander scale ?

I think they may hope that unlike in Schwarzenegger movies there won't be many goal conflicts between „us“ and „them“. Humans and others are interested in those they can compete and collaborate with, because they share the same goals. Politicians are mostly interested in other politicians, kids in other kids of the same age, goats in other goats. Supersmart AIs will be mostly interested in other supersmart AIs, not in humans. Just like humans are mostly interested in other humans, not in ants. Note that the weight of all ants is still comparable to the weight of all humans.

Humans won't play a significant role in the spreading of intelligence across the universe. But that's ok. Don't think of humans as the crown of creation. Instead view human civilization as part of a much grander scheme, an important step (but not the last one) 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 3 billion years ago. This is more than just another industrial revolution. This is something new that transcends humankind and even biology. It is a privilege to witness its beginnings, and contribute something to it.

Photos : Wort & Bild Verlag / Eleana Hegerich