**THE PEACE MACHINE**

Harnessing Artificial Intelligence for Peace

By Timo Honkela

Translated from the Finnish by Owen F. Witesman

# Between the Words

woman
man
vegetarian
omnivore
believer
atheist
teetotaler
alcoholic
invalid
healthy
adult
child
elderly
young
single
married
communist
capitalist
technocrat
humanist
artist
engineer
organized
chaotic
leader
worker
us
them
white
black

so many words
and many more
do we fit between the words?
to meet ourselves – and the other
being human
there are not words for all colors
and what is more:
is anyone painted with only one?

*Timo Honkela, 1999*

**TABLE OF CONTENTS**

The impact of the past on current feelings

Emotional game theory

JUSTICE IN SOCIETY

The capriciousness of justice

Betwixt or between choices?

The pursuit of social truth

Facts and truth

X-raying society with machine content analysis

ARTIFICIAL INTELLIGENCE IN THE SERVICE OF HUMANITY

The history of machines before artificial intelligence

Different forms of machine learning

Data as the driver of machine learning

Machine translation builds bridges between people

A meaning negotiation machine

The US presidents' annual "State of the Union" speeches

BUILDING THE PEACE MACHINE

Artificial intelligence to promote peace

Learning peace

Machine translation to promote peace

What does a meaning negotiation machine have to consider?

The place of meaning negotiation

Artificial intelligence to calm emotions

Free and safe

Justice: new methods for an old idea

THE ROLE OF THE HUMAN BESIDE THE MACHINE

Can a machine be creative?

Games, the human, and the computer

## CHAPTER 1: A RICH LIFE EXPERIENCE – OUTLINING THE THEMES

### Why are there wars?

People have been fighting wars for all of recorded history. Long ago, when people still hunted and gathered their food, there were probably fewer wars, because reserves rarely built up. Waging war wasn't attractive because you couldn't pillage other people's storehouses if they didn't have any. Once agriculture developed, reserves could accumulate, but resources were still scarce, which often leads to violence. Even if there are sufficient resources, their unequal distribution can lead to conflict. If resources can't be distributed evenly enough, the risk of war increases.

One reason for war is preparation for war. When we prepare for wars, we also enable going to war. In a game theoretic sense, there is a trap lurking within any preparation for war. If one player prepares for war, the other player can't decide not to prepare because that would lead to inevitable defeat if war ever did break out. Turning the other cheek is a nice thought, but when we analyze questions of war and peace at a global level, it isn't a sustainable solution.

Human language is based on categorical thinking. This means that we tend to see each other categorically, that is to classify each other. Unfortunately, we don't think of each other as one big family but rather as many different groups. There is an "us" and a "them." A strong affinity for one's own group—such as a nation—can be a valuable thing, but in practice it can also lead to violence or at least to an increased risk of violence.

Wars also occur because for some small groups, wars are beneficial. War can function as an instrument of power. War can also be economically profitable for groups who provide warring factions with weapons or other resources. Sometimes war has also been seen as producing a sort of creative destruction: when things freeze up in a state that isn't considered rational or functional, war reorders the situation. Of course, it would be very regrettable if positive social or political progress always necessitated war. In addition, what is positive from the perspective of one group can be negative from the perspective of another.

Simplistic solutions are rarely good. It is unfortunate how often we think that decisions must be made by choosing from the options that are presented to us. In this book I to argue that over the long term, we must learn to make decisions by choosing from the space that exists *between* the options we see initially. It is true that in contemporary democracies sometimes decisions are made by seeking compromise between disparate extremes. However, this isn't done systematically enough. It may be overly optimistic to hope that every politician could be cognizant—and openly admit—that their own solution isn't necessarily the right one. No single individual or group can hold the one true option. Rather, solutions emerge as a synthesis of the differing opinions and knowledge of the broader populace.

Of course, there are situations in which a choice genuinely must be made between distinct options. However, complex systems require solutions that simultaneously take different perspectives into account. Making decisions that force everyone into the same rigid structure is unnecessarily categorical and can, at worst, even be violent toward people. As we build a society and world of equality, we must be able to take into account that people are different. From the perspective of systems theory, the question is how equality and diversity can be taken into consideration in a way that serves everyone's interests.

If the reader can allow me some mild cynicism, people are often motivated by the desire for money, power, and honor. For some people, wars are simply good business. Selling weapons to countries at war can even seem like a good reason to start new wars. Business people or groups can't necessarily start wars themselves, but their influence can prepare the way for a chain reaction that leads to war. Ambition is another traditional reason for starting wars. Kings and other leaders have sought glory as famous generals who have improved their country's status through warfare. The pictures of successful generals are hung on walls and their achievements are lauded for years to come. Speeches proclaim that the nation prevailed in war because the general

led the troops well. Seen from another angle, though, a great number of people lost their lives. On the fields of glory, human life doesn't hold much value.

Money, power, and glory can play into any historical situation. Factors that depend on history include, for example, a desire for revenge or reconciliation and feelings of aggrievement. For example, in Germany between the First and Second World Wars, people felt that the settlement following the First World War had been unfair to the Germans. Taking advantage of this feeling, a populist leader was able to manipulate millions of people into supporting his pursuit of revenge. This makes it apparent how important question of identity and history are. How do people experience belonging to a group or not belonging to a group? Who are the us and them?

Sometimes the formation of identity can be based on tiny differences between people. The formation of a society is a historical process which can last thousands of years. A good example is conflicts over land in which one party claims "we were here first." However, the claim of being "first" can be evaluated on any scale of time from years to decades to centuries. Of course, people usually evaluate questions of justice from their own perspective.

So how can artificial intelligence and other similar technologies help disrupt the causes of war? It hardly seems realistic to spend much effort addressing the human desire for money, power, and honor. However, improving communication over linguistic and cultural boundaries, and providing people with the means to prosper in a changing world could lead to some developments in this area—a feeling of security and a recognition of the benefits of cooperation for one's own position can reduce self-centeredness. Machines can help people find more diverse viewpoints, which break down simplistic categorical thought. If categories divide people, machines can be used to find things that unify them but which they themselves don't see. People's thoughts often get stuck on a small number of controversial issues that highlight confrontation. When this happens, the solution is to find a richer selection of issues to investigate. Given a broader range of ideas, stakeholders can find things that unite them as people rather than dividing them.

In a mathematical sense, this same thing can be expressed as projecting the issue under consideration into a higher dimensional space. In practice this means that instead of investigating three variables, we might look at ten or a hundred variables. In a space determined by three variables, things might go nowhere, but in a one-hundred-dimensional space, it might be possible to identify many variables that support finding common ground. Machines are a good tool for conducting this type of analysis.

Machines can also help us understand the desire for money, power, and honor better by allowing more precise analysis of the influence of emotions of people's actions and decision making. Some stimuli produce positive, constructive solutions. A good example is microinterventions, which can, for example, make the pursuit of power less attractive. This goal relates to how democracy could be expanded within society.

We could apply a hockey metaphor to the desire for honor. In hockey, success is often measured through goals scored. However, the players who make the goals are not the only significant factor. They also need the players who pass to them, the defenders, and the goalie, who prevents the opponent from scoring goals. The game also isn't possible without trainers, referees, maintenance staff, reporters, fans, arena construction crews, and the decision-makers who make the construction possible. The total picture also includes the government officials who make sure that ice halls are safe and don't fall down on the players' and spectators' heads. We may see individual ambition on the part of players, but there is also a lot of cooperation that goes on in the background to facilitate the competition.

If we look at any significant social phenomenon, we see complex interdependencies between people and things, and when these are broken by war or other crises, a wide range of problems are created for people and the functioning of society. And because this complexity always develops in any human system given enough time, simple, straightforward solutions rarely work.

War mostly has losers. The number of possible winners is small. So, one clear way to promote peace in society is to distribute power democratically, as widely as possible. Then the proportion of people who really gain from war will be so small that the probability of war will decrease. The greatest threat in this sort of situation may be populism. Using populism, large groups of people can be made to imagine that war or warlike conflicts will be of benefit to them. The antidote to populism is as broad a base of knowledge as possible and a humane understanding of the complexity of society. And in this, artificial intelligence and similar technological developments can be a substantial help.

**Can't we all just get along?**

I was born in the town of Kalajoki, Finland on August 4, 1962. An important figure in Kalajoki was Untamo Sorasto, a local doctor. I received my middle name from him, and as an obstetrician, he also helped me into this world. As far as I know, my early years were happy, but dark clouds quickly began to gather. I was maybe five years old when my parent's fights began to grow louder. Watching them, I was anxious and afraid. Somewhere I learned the phrase "can't we all just get along," and tried to use it to get my parents to stop their terrible fights. Often, I woke up to nightmares that left me short of breath.

Parents' relationships are vital to small children. Because of this, children carefully watch their parents interactions.[4] The fights between my parents grew worse, and I have vague memories of my mother becoming depressed. I remember how she sat next to the oven in our modest yet relatively large house, looking sad and smoking.

We didn't have hot water or an indoor toilet, so in that sense our life was extremely modest compared to nowadays. Huge economic development took place during my early years. When I was born, Finland was still poor, but by the time I reached adulthood, Finland was one of the richest nations in Europe. My mother had worked as a laboratory assistant at the VTT Technical Research Centre of Finland when my parents lived in Helsinki after the Second World War. The subsequent move to a small town, far from all the familiar places where she'd imagined building her life and career must have been depressing.

My father suffered from a panic disorder, and he couldn't stand feeling like he wasn't in control. So, he never traveled by plane and only once by train after the war, when he came to attend my dissertation defense in Espoo in 1997. When my father was about sixty years old, Doctor Sorasto helped him in a dramatic situation, which also revealed something about my father's character. My father had been moving some gravel to driveway in the front yard when his fingers became caught between the wheel barrow and a vehicle. His thumb was crushed beneath nearly a thousand kilograms of rock, leaving it unrecognizable. However,

my father climbed into the car and drove the thirty kilometers to Untamo Sorasto's home in town, where he showed his thumb to the experienced surgeon. Sorasto then proceeded to take out his home medical kit and save the thumb.

One spring day in 1971, we set off to visit my older brother Kari's home in Oulu, where he was studying architecture. I don't remember whether we were away for one or two nights, but our return home was entirely different from what I expected. When we arrived at the house, I unlocked the door and pushed. For some reason the door wouldn't budge. It felt like there was something blocking the door, but I couldn't imagine what it might be. With all the strength of an eight-year-old, I managed to push the door open as I waited for my father to come from the car with our luggage. When the door opened, I saw my mother lying lifeless on the floor. I remember shouting to my dad, "Do something!" I don't remember my father's precise words, but soon he said that Mom wasn't alive anymore. My hope that we could all just get along never worked out in any constructive way.

I moved next door to live with my uncle Paavo's family, where, as I remember, I spent the next two weeks. Paavo's wife, Aino, was my godmother, and my cousin Olli, who has now passed away, was my playmate, so living with them was fine. However, my grief over my mother's death was enormous, and I cried almost nonstop for those two weeks. One result of the tragedy was that I wasn't left with much memory of my mother. I still can't remember what she looked like. Later, when I was mostly grown up, I saw some old pictures of my mom. The shock was so great that I ran a high fever for the next few days. There may have been some other cause than the shock of seeing the pictures, but I think it's a plausible explanation.

My father cared for me in his own, slightly severe way. He'd been in the war after all. We didn't brush our teeth or cook much food at home, instead eating our lunches in nearby bars and cheap restaurants. For a small child, the death of a mother was a mystery that demanded an explanation. That was also one reason why I became a precocious academic before my time. Even when my mother was alive I'd been bookish, and the world of adults interested me more than

kids' things. So even before I went to school, people called me "the Professor." I knew how to read and count before school, and while the others were learning to add and subtract, my big brother Kari was teaching me the basics of exponents. When I was eight or nine, my father began to encourage my mathematical ambitions. For example, he introduced me to Abel's theorem and other problems to feed my intellectual curiosity.

After my mother's death, psychology also began to interest me. Even at that young age I'd begun to try to construct theories that could explain people's strange behavior. At first, I didn't have words for these things, but once I found the Kalajoki Library, I gradually began to get the "official" story. In retrospect, I find it irritating that I don't have any memory or library lending record of what I read then and the conclusions I drew. The library, my school, and the families around us kept me engaged in life and mostly healthy emotionally. Decades later I also developed a susceptibility to panic attacks, like my father. In my case the panic is associated with exercise, though. For the past thirty years I haven't been able to participate in any physically demanding sports, because the rise in my heart rate causes a panic attack.

In many ways, being diagnosed with cancer has been fortunate. As luck would have it, during the course of my treatment, my doctors realized that the medication I'd previously used for my panic disorder was causing more harm than good. One of the new medications they found has partially made my life with cancer happier than before. Cancer has also helped me understand the value and importance of every single day. For a person like me who experienced trauma as a child, this has opened up an entirely new perspective on life.

As a child I was precocious, a prematurely aging know-it-all, whose IQ measured between 160 and 180 as a teenager. As an adult, I've probably come across as a little supercilious, still a know-it-all, and I'm sure some people have found me proud. I have no defense. Because of my lost childhood and the young love I lost, I've had a huge need for approval and respect.

When I was young, I claimed that if everyone was like me, there would be no war. But in truth, I was filled with bitterness, perhaps even hate. For example, my relationship with the female sex was one of extremes: I was either fawning and complimentary or bitter and angry. In practice this meant that having a balanced relationship with a woman was very difficult. I easily became suspicious and fearful when my wife criticized me for anything, even simply forgetting to wash the dishes or something similar. Even though I've never received an official diagnosis, I've understood with increasing clarity that I've suffered from a dissociative condition. In other words, everyday situations could send me back to the world of that damaged child, the fear of losing my mother pushing out my ability for rational adult thought.

As I see it, my academic approach to life has saved me. In the midst of traumatic experiences and all the other difficulties of life, I've been able to rise above and focus my attention on something else. Life has been anything but easy. There have been plenty of challenges, but these experiences have also helped me understand life. An easy life wouldn't have given me the same understanding. In this sense brain cancer has been a good thing, putting a stop to processes that maintain old ideas and bad habits. Cancer has also changed my life in ways that have given me new opportunities. "Can't we all just get along?" is now something that gives me perspective on the horrors I experienced as a child—which are nevertheless very mild in their horror compared to the things people experience in war.

**Experience increases understanding**

Travel broadens your horizons, they say. In addition to constant travel being work, it also gives you experiences with different languages, cultures, and social perspectives. Even as a child I had the opportunity to travel and broaden my view of the world. My brother lived in Switzerland, and beginning at about fourteen years old, I was able to get very inexpensive flights to visit him in the Zurich area. That was a big privilege given that some of my schoolmates hadn't even traveled to nearby cities such as Kokkola or Oulu at the time. It was

also important for my future work. Tolerance develops naturally when you meet people who look differently, when you hear different languages, and when you become acquainted with different customs. When you have experiences like that as a child, you understand that your own way of doing things isn't always the only possibility.

The phrase "when in Rome do as the Romans" has its strengths, but it shouldn't be taken to extremes. The music of Jean Sibelius is not unambiguously Finnish, since he also studied abroad. Sibelius just adapted that learning to his locale and took inspiration from Finnish culture and landscapes. If we were perfectly Finnish, we wouldn't even eat potatoes, since they're an eighteenth-century import. We do best in life when we're flexible. Flexibility doesn't mean accepting inferiority, but taking advantage of different influences and opportunities makes sense.

Much more important than a strong identity is a strong sense of self-respect. A person outfitted with self-respect can succeed in different circumstances and with different people without the need to defeat or subjugate them. They don't need to do anything unpleasant or aggressive in order to get along with themselves. We must ask how we can encourage the development of a strong sense of self-respect in everyone. Sometimes I've heard it said that violence is weakness. People resort to it when they don't know how to handle things constructively. When children begin to fight, often there's a sensible adult present who takes over the situation and teaches the children to solve their disagreements without fists. But what to do when adults clash?

Building identity begins in childhood. When I was a child, I thought that the Kalajoki Sandbars were one of the best beaches on earth. On the shores of the Mediterranean, I've found that nothing like the Kalajoki beaches is to be found in Greece, Italy, or Spain. Usually the sand on those beaches isn't fine enough, and some of them are just gravel or rock instead of sand. For a long time, I hadn't experienced anything better than the Kalajoki Sandbars anywhere in the world, until a few years ago I found that the beaches in Florida beat them.

Because of the Kalajoki Sandbars, my childhood summers were very international. At that time, vacationing in the South was rare, so instead of going to the Canaries or Majorca, thousands of Norwegians traveled to Kalajoki to spend their summers. Amazing sandy beaches and beautiful forests were the backdrops for my childhood. Nature gave me strength when things were hard with people.

As a child and teenager, I mostly gravitated toward the world of facts. A high-quality education system and library played a major role in that. When my parents couldn't provide enough support or security, a school with professional, empathetic teachers stepped in. Sometimes things with friends were hard, because my classmates weren't fond of the teacher's pet. Thankfully this improved with time. Eventually we divided into smaller groups where we were all similar in one way or another. Shared interests made our activities and conversions meaningful and exhilarating.

As school children, our precocious group of friends formed a Friday Club. The Friday Club had very strict rules, which we modeled after the official Associations Act. The club collected member dues, and we ended up with so much money that we needed our own bank account. But besides these practices we borrowed from the adult world, our actual activities were normal kid and teenager stuff. For example, we bought a basketball and some track and field equipment. Of the other members of the Friday Club, Harri Kalliokoski became a bank manager, Matti Yrjänä and Mika Kamunen became school principals, and my cousin Olli became a purchasing agent.

Olli later worked in several Nordic companies, which inspired him to take back our historical name, Hongell. Our grandfather, Eino Honkela, the financial director of a local dairy, and his brothers had changed the name to make it sound more Finnish about a hundred years ago.

In my second year of high school, I had the opportunity to travel to the United States on a scholarship from our local Lions Club. In 1980 I flew first from Helsinki to New York, then from New York to Chicago, and finally from Chicago to South Bend, Indiana. My host family was waiting for me at the airport and then drove me to their home in West Lafayette, which is a major university town. I spent the summer there. The sponsor of my visit was a metallurgy professor, John F. Radavich, who had arranged the trip with the Finnish organizer. My actual host family had a boy my age, Steve Krause, who introduced me to local life and culture. West Lafayette was a very peaceful place. Life there wasn't much faster than in Kalajoki. Which was fine, because I wasn't looking for that. I loved becoming acquainted with Purdue University, and Professor Radavich also introduced me to his metallurgy research.[5] In his work he used an electron microscope, which was a new thing back then. Radavich was investigating how space rockets withstand the high temperatures and other stresses they encounter. I even got to take home a small piece of the Gemini 8 rocket.

One important person in the organization of my visit was Lassi Kivioja, who was also from Kalajoki but had ended up living in America. On my trip I learned to respect many positive aspects of the United States. But at the same time, I understood with the eyes of an eighteen-year-old that society in the United States is severely divided into winners and losers. After I returned to Kalajoki, I wrote a two-page article for the local

newspaper about my experiences. I emphasized that Finns should appreciate how well their society works. It may be relevant to note that Finland was never a communist country and a market economy with private firms was already the central form of organizing the economy well before Finnish independence in 1917.

Unfortunately, my observations back then are still relevant today. The division into winners and losers hasn't gone anywhere, and if anything, it's only grown more pronounced. At the same time, my sense that Finns should value Finnish society has also grown stronger. We shouldn't be too excited to go looking for models elsewhere, since we've managed to arrange our affairs quite well ourselves. If it ain't broke, don't fix it, as the saying goes. We may face challenges, but that doesn't mean the entire system is broken. Artificial intelligence and machine learning can change society in radical ways, but that doesn't mean that we should rush to adopt these new technologies. Organizing a society is ultimately a matter of choices. Resources are always limited and divvying them up often requires making difficult choices. If you want everyone to get along financially, you might have to take something from the top. If a nation wanted to avoid the worst social problems in the United States, this sort of leveling, probably though taxation, would be unavoidable. However, decisions like this must be based on broad democratic consensus, not the opinions of individual people. So, you won't find any suggestions for a new social order in this book. What you will find is thoughts about how developing technologies can be utilized in solving complex social problems.

When I went to Oulu University after high school, I applied to five different majors. Since I had perfect grades, I could study whatever I wanted, but based on my experience in high school, I was most interested in mathematics, so that's what I chose out of the five options I had in mind. However, university mathematics quickly turned out to be far too impractical. We were

always at such a far remove from any practical application, especially since my professor, who was approaching retirement age, seemed more interested in playing chess than in the intellectual motivation of his students.

At this point, I should admit that my criticism of this professor shouldn't be taken too seriously. Young students rarely have any inhibitions about finding fault with their professors. But the most important task of a university is research. If some professors aren't interested in teaching, they should be able to focus on scientific study. There are a lot of areas in which a better division of labor would be in order. These days universities over-emphasize a sort of equality between employees. I think this is a mistake. Let good teachers concentrate on teaching and good researchers concentrate on research! Of course, there are people who do wonderfully at both. For the university, it's most important to find people who best fulfill the three central missions of universities: research, teaching, and exerting a positive influence on society, both in the private and public sectors. One person can't simultaneously do all three simply because there isn't time. Some research results require decades of focused work, which shouldn't be interrupted by localized and temporary fluctuations in economies. Perseverance is one of the central characteristics of wisdom. If some field of research hasn't made a breakthrough in ten years, it shouldn't be abandoned, it should be continued for another ten years or maybe a hundred.  People always find new paths to solutions in the end.

After realizing that studying mathematics itself was too depressing, I went looking for new opportunities, ultimately finding Professor Pentti Kerola's human-oriented information systems design program. This was interesting, and, in retrospect, I believe more important. Kerola's idea was that information processing isn't just technology and mathematics, it also

requires an understanding of the purposes and groups of people the systems are being developed for.[6]

I switched to Kerola's computer science program, which due to its multidisciplinary goals was officially organized under three faculties. And in addition to computer science, mathematics, and statistics, it also included psychology, ethics, sociology, and economics. There were many subjects we needed to study that didn't have established courses, so, for example, the department organized the ethics course itself.

In Oulu, we took the term "information" very seriously, with a significant focus on information theory and epistemology. A large amount of our study was philosophy, because the assumption was that if you want to build good information systems, you have to understand what information is all about. In hindsight, this idea was absolutely critical. It was also important that the course of study included theoretical and practical dimensions, and that human emotion was not overlooked. Professor Kerola invited a visiting scholar from Florida, Professor Willian A. Taggart, who had studied human information processing styles.[7] I participated in Taggart's course, which turned out to be so interesting and true-to-life that it continued to inspire me for years afterwards. The course provided me with a greater understanding of different people's different styles of thinking. That course also had a significant indirect influence on many of the subjects I deal with in this book.

While I was studying, I managed to get a summer job when one of my fellow students, Sanna Kajava, also of Kalajoki, gave up her summer job for another position. At this first summer job, I programmed database systems for the Geological Survey of Finland in their geochemical unit.[8] There I learned that the Geological Survey was a partner in a Finnish Innovation Fund project called Kielikone (Finnish for "Language Machine"). So, the next

summer I moved to Helsinki to work at the Innovation Fund with some of the most experienced information technology, artificial intelligence, and language technology experts in the country.

Our presentations of the project's results were the beginning of my international scientific travel, and conference trips became indispensable to deepening my understanding. At these meetings, researchers from around the world exchanged ideas and opinions in a way that was impossible simply by reading. This furthered my education, sometimes in surprising ways. For example, at my first conference, in Gothenburg, Sweden, one night I met the people responsible for organizing the Nobel Prize for Literature. On that trip I also had the misfortune to drop the electric razor I'd received from my father, breaking it on the floor. From that time on I've only used simple razors, a detail I mention now apropos of nothing.

The Kielikone project used traditional artificial intelligence to model linguistic significance, relying on predicate logic and other formalisms. In this approach, the machine's linguistic capabilities are constructed by programming them to recognize questions and commands given in sentence form. These experiences led my academic interests in slightly different directions. It dawned on me what a limited tool logic was for describing human knowledge and understanding. At first this realization was a bit oppressive, because I had a hard time finding companions for this new journey. So, it was a relief when at some point I found the late work of Ludwig Wittgenstein (1889-1951). I realized that we couldn't base semantics on the ideas that prevailed then. I gravitated more and more toward the idea that understanding of the world isn't given to us when we're born but rather forms through learning and experience. Therefore, artificial intelligence

also shouldn't be constructed as if human understanding in static and complete, but rather it should be developed through some sort of machine learning.

To explore these ideas, I established a machine learning working group in the Finnish Artificial Intelligence Society. The lucky thing was that one of the world's preeminent machine learning researchers is a Finn. Professor Teuvo Kohonen is renowned for his research in the area of neural networks. Over the decades, many different ideas had been suggested for how to teach machines. Kohonen's Self-Organizing Map algorithm, which we'll return to in more detail later, made it possible for machines to learn new skills instead of simply being programmed to know them. Kohonen had studied the function of the human brain, especially how the brain orders observations into sensible aggregates using "maps" that form on the surface of the cerebral cortex. In 1981 Kohonen published an algorithm that allowed the mimicry of this function as a model for machine learning. Kohonen's findings have turned out to be useful in many ways, and his research has received much international attention.

After serving in the army, in the early 1990s I began working for the VTT Technical Research Centre of Finland and started exploiting Kohonen's work in neural network modeling. I programmed independently at home and at work. One of my neural network models I actually programmed based on things I'd learned in the army. On my early morning guard shifts, I found time to read neural network literature and learn the C programming language, which is a good language for building these sorts of programs.

The program I wrote simulated Kohonen's self-organizing map. With this program I then started to do various tests at the research center, including some having to do with meaning in language. Working at the research center also gave an enthusiastic young person like me the opportunity to get involved in all sorts of different things on far too many different

projects. During the worst period I was running three parallel projects and involved in another two, one of which was Finland's first or second EU project. In this project we were researching how computer programs could be built so their interfaces would work in several different languages. We were using technology to investigate language and culture. As I worked on this, I traveled all around Europe, especially to England, Ireland, France, and Greece, which were the other partner countries.

The Greeks aren't nearly as feckless as some public discussions might lead one to believe. They've always had strong mathematicians and philosophers. The EU has also been good for Greece by raising the standard of Greek research. In the early 1990s, Greece's fundamental problem was that their best talent didn't stay in the country, instead leaving in search of greener pastures. Because of the EU, Greece has been able to enhance its expertise to such a degree that top experts often stay. I've seen this with my own eyes. Economic troubles always make researchers flee, which we've also seen in Finland recently. When you don't have oil wells or other easy sources of wealth, you have to focus on expertise.

My work with neural networks at VTT ultimately led to joining professor Teuvo Kohonen's group at the Helsinki University of Technology. I gave a presentation called "Maps to Information Highways" about my results at VTT, in which I used Kohonen's self-organizing maps in information retrieval. Kohonen immediately saw the significance of my findings and organized a research group to pursue information retrieval based on neural networks. In addition to me, Kohonen invited a promising young researcher named Samuel Kaski, and at my recommendation we also added another gifted researcher from VTT, Krista Lagus. We began developing a system we called WEBSOM. The WEBSOM algorithm automatically produced document maps, pictures in which documents with

similar content are positioned near each other. I am pleased to be able to say the work I was able to accomplish with Samuel and Krista has had a significant influence on the way statistical technology is applied to the processing and visualization of text documents.[9]

In this book I attempt to focus on things that aren't simply a repetition of the old or the obvious but which move thought in new directions. A machine is not a person and a person is not a machine. Many things are impossible—fortunately sometimes and unfortunately other times—but a lot of things are possible, and progress is constant. Ultimately the time to go from idea to implementation of the WEBSOM project was quite short. Slow technological development will not be a problem for the Peace Machine either. Of course, there are practical obstacles to overcome, such as the possible insufficiency of computing resources and deficiencies in some sub-disciplines. But the greatest hindrance to creating the Peace Machine is that it is a sociotechnical system.

By sociotechnical system I mean that things can't be done with the technology first, in isolation. Instead, it also requires a deep understanding of how humans and human societies work. This makes the humanities and social sciences critically important. Without a philosophical perspective, there's no point even attempting a project like this. A peace machine developer could far too easily become a well-meaning fool if he failed to pay careful attention to real phenomena.

**The social importance of expertise and research**

Working at the Helsinki University of Technology in research groups led by Teuvo Kohonen and later by Erkki Oja provided me with a wide international network and helped me get acquainted with many multidisciplinary efforts. In addition to neural networking, I had all sorts

of conversations in Finland and abroad, at conferences, in meetings, and through publications, on subjects such as systems theory,[10] linguistics (language technology),[11] cognitive linguistics,[12] philosophy,[13] psychology,[14] sociology,[15] economics (from a cognitive linguistics perspective),[16] pedagogy,[17] cognitive science,[18] and organization research[19] to name just a few.

About ten years ago I began to describe my work as cybercognitive philosophy of language. For a lot of work, placing a research topic in a single discipline is artificial, because the same work can contain elements from many fields. For example, my classic article about word maps formed from the Grimm fairytales was published at a neural network conference, but it relates just as much to cognitive science, linguistics, and the philosophy of language.[20] The problem in some cases is that the connections are implicit, between the lines. And I admit that I've found it more interesting and important to move onto new topics and new areas than to work a topic to its conclusion. This makes it possible to get a broad overview but leaves the polishing and finalization of individual results to others.

During my second term at Helsinki University of Technology, I had the opportunity along with my colleague Krista Lagus to lead a group looking at computational cognitive systems. Using computational models, we studied issues related to individual and social group activities. In this we made use of Kohonen's self-organizing maps and many other methods.[21]

For myself, there was one defect in this academic environment. I was surrounded by people whose interest was in methods, in mathematics and related sciences. However, I was interested in phenomena like language, linguistic interactions, and the relationship of language to the world. I wasn't interested in the methods, I just wanted new tools to give

me a better picture of the phenomena. Since pen and paper aren't always enough for understanding complex phenomena.

In top research institutes, in our case a center of excellence, international collaboration and networking opens your eyes to new things. A young researcher at the beginning of their career reads a lot of articles, but the names of universities, labs, and people remain superficial facts. Their content opens up in an entirely new way, though, when you meet the people behind the names, hear their presentations in person, and can chat with them at lunch at a conference.

When you meet experts in different disciplines, there's also the great benefit that you can test your own understanding. This isn't possible just reading books and papers. My own strength may have been in that I've always been willing to grab people by the sleeve, no matter who it is. Even at age twenty-five I didn't hesitate to approach world-famous researchers with my questions. At first my questions were probably pretty simplistic, but your understanding grows when you work up the courage to bother people who are wiser than you. Possibly because I'd started asking questions when I was a kid, my questions weren't ever completely trivial. I deduce this from the fact that the answers usually took more than the few minutes people are usually willing to sacrifice for this sort of thing. I encourage young people in the same situation to consider when it might pay to be bold and tug an expert by the sleeve and when it would be better to follow the situation more passively and gather understanding by reading and studying. Whatever the case may be, it always pays to do your homework. Expertise also always has an important social dimension. Social decision-making requires democracy so that everyone can be included in the process. But not everyone knows as much about everything. Some people know

more about some things than others, and they know how to make decisions about those things better than most others. One way to inject expertise into decision making is collaboration between politicians and public servants. The fluidity of this cooperation is important for society to function. The fact that an attempt should be made to take everyone's needs into consideration doesn't mean that issues can be solved by a voice vote, though. Making the best possible decisions requires information and the experts who produce that information and know how to analyze it.

I get the feeling that the position of experts in social decision-making isn't entirely clear. What is the proper role of expertise in an egalitarian world? If for example a city wants to renovate its transit links over a body of water, decisions must be made about building bridges. The decision to build a bridge can be made communally, democratically, but the precise location of the bridge, the building methods, and the building materials should be left to experts to decide. But the roles of experts, stakeholders, and citizens aren't always so simple in decision making.

My own views about how to organize research and, for example, product development, has been influenced by my visits to Berkeley, Stanford, and Google. It may come as a surprise to many that research isn't a fun, egalitarian activity but rather a savage competition which for more and more researchers is global. In many fields, survival means being one of the elite in the world. Development marches on, and there is no time to take a breather. This means that you can't take summer vacations. The comments we sometimes hear in public forums about professors having long summer breaks are insulting or at the very least ridiculous. A researcher who is devoted to his work is on the job all day every day.

I saw a good example of this back in the 1980s. Professor Taggart, whom I've mentioned before, told me about a mental technique he used in his research. He said that when he goes to sleep he repeats to himself a research question he's working on, so his unconscious mind can process it as he dreams. He said he kept a notebook and a pen next to his bed, so he could write down interesting ideas if he happened to wake up during the night. If that didn't happen, he wrote down his thoughts after he woke up in the morning. Nowhere close to all researchers use their time so purposefully to advance their research, but everyone carries their research questions with them wherever they go. The problem from the researcher's perspective is how to avoid burnout. Preventing that requires finding ways to break away from the research.

Here I emphasize again that a professor's central task and reason for existence is not teaching, it is research. It would be good if students could become involved in research activities as early as possible. It's senseless that students study and learn things just to complete courses, when they could all be involved in research at their own level of development. Of course, not everyone will become a researcher or ever could, but a scientific approach to the world is much more constructive and useful over the long term than the slavish memorization of temporary facts and methodologies. A well-educated person can overcome many challenges.

Unfortunately, many researchers have done their work in difficult or unpleasant circumstances despite the importance their work has taken on decades or even centuries later. Scientists and artists are in a similar situation in the sense that few ever receive the recognition or financial success they deserve during their lifetimes. There is no real hope for fair compensation. The reward for a researcher's work is largely in the work itself.

In my own ideal world, every person with sufficient education and sufficient dedication would have the opportunity to do research without ever needing to worry that simplistic, uncomprehending decisions might cut short any project. For the great majority of researchers, salary doesn't even matter much. Studies have shown that small raises or differences in pay can damage a research unit because jealousy is a more difficult thing to deal with than the financial incentives the raises are meant to achieve. I don't mean that researchers shouldn't be paid properly for their work. However, it can't follow the same principles as in the business world, because there the goals are completely different than in research. Fortunately, even in the business world avarice isn't always the central motivation. Many people genuinely love their work.

It would be great if more people could do their work for the love of the task not just to earn a salary. Maybe we'll move in that direction once robots replace people in all the repetitive jobs that are better for robots than people anyway. Of course, this change will not be easy, and once again it would pay to note that people have different hopes, abilities, and ideas. The freedom that is a part of some people's work and in a sense their life-blood, can be an oppressive idea for others. People are simply different. That's why it pays to ask everyone separately what they want.

**A functional society and a just economy**

If people feel that society treats them unfairly or if people's financial situation is unbearable, problems follow. These situations can lead to civil unrest and even suicide. This being the case, good social decision making and a functional economic system are central concerns for peace.

The effects of computers, artificial intelligence, and machine learning on the functioning of societies and economies is large and growing. For example, the computer systems in stock markets can make such quick decisions that people can't track them in detail and can only examine the results after the fact. The nature of the systems will probably change gradually to match their environments better, ultimately allowing them to take into account changing conditions. Then the systems won't make decisions day after day following a certain set of rules; instead, learning machines will, for example, track changes in the economy and alter their own decisions accordingly.

The more context-sensitive these reactive systems are made, the greater will be the danger that the overall system will become chaotic. The more automatic systems are a part of social and economic decision-making, the more important the design of the rules of competition and cooperation will be. We must agree what the principles of social decision-making will be. Machines also involve risks. Smart artificial intelligence systems can be used to deceive and manipulate people. However, I still see artificial intelligence as an important tool for the development of democracy. Up until now, only a small number of people could reasonably make decisions at a time. Even that hasn't always been easy, as the meetings of any housing cooperative can demonstrate. Of course, even nowadays it is not possible to make decisions for entire cities, countries, or continents in such a way that everyone is involved. However, using artificial intelligence it is possible to create a scenario in which a very large number of people can in a sense be at the same table discussing and deciding things. Later in this book we will return to how artificial intelligence makes this possible.

Economics has developed many different mathematical models and computer-assisted systems meant to predict economic phenomena and promote their understanding. I tried this myself when I was a student, when Tapio Vailahti and I created an economic simulation, a sort of businessman game. Our idea was far too large for the computer programming project we were supposed to be doing. We decided to develop the program in Fortran, which is a mathematical programming language still used quite widely today. We were living in a time just after the demise of punched cards, but programming was still "paper work." Instead of a computer screen, programs were made on paper printouts produced by keystrokes. Programming was still fun, though, since we'd chosen the subject ourselves and it was something we were interested in. Over six months we wrote pages of code, but we didn't get very far in terms of modeling or understanding the economy. However, in addition to learning coding, we did get some familiarity with business and economic concepts. By using a little imagination, you can actually get familiar with almost any field of science if you know how to utilize the programming possibilities offered by a computer. That's why I think it's a good idea to teach programming and coding to all children, even though that doesn't mean that they all have to become programmers or information system designers. Pen and paper are good tools for writing, for outlining things and conveying ideas. You can also use computers for the same things. Programming can provide understanding of artificial intelligence and machine learning, even though understanding them doesn't necessarily require programming skills. If the machine is used too mechanically, it easily becomes impractical for people whose thinking doesn't line up with the languages of programming or mathematics. This viewpoint should be respected.

The Helsinki University Faculty of Arts has recently established a Centre for Digital Humanities (HELDIG). This is a good example of how computers can be leveraged in even the most diverse fields. An enlightening experience for me was when I went through Shakespeare's sonnets with the literary scholar Sakari Katajamäki using neural network analysis. We learned a great deal about the structure of some of the poems.[22]

So, what kind of goals can we set for the computer models we create? So far, most systems have been developed for specific applications. We run into systems like this more and more frequently in everyday life. Even though artificial intelligence systems can be made to serve many practical tasks, utilizing artificial intelligence to understand complex phenomena is something slightly different, and the goals are also a little different. Here I might refer to the typology of human interests proposed by Jürgen Habermas (b. 1929), which can also be thought of as reasons or needs related to knowledge. Habermas identifies three types of human interests. These are the technical, the practical and the emancipatory.[23] The role of the first of these in the current discussion is obvious, but the other two interests are also important perspectives when developing artificial intelligence systems.

Practical interests are represented by machine learning methods whose purpose is only to predict various phenomena and variables related to them. Developers of technically oriented systems might see this predictive ability as the most important or only criteria for system development. They might want to get the most precise answer possible to questions such as how the economy might change over the next year. The success of the system might be measured by how well its predictions play out.

However, in my opinion prediction is not as central to the utility of these systems as it may seem at first. I argue that even more important than prediction is what we learn about complex phenomena like economies with the help of computer aided modeling. This represents the human interest of emancipation in the Habermas typology. If predicting the economy were the only goal, that would be the same thing as admitting that we can't influence the economy in any way. According to this view, the economy operates independently and spontaneously, and we can only try to predict its movements. Affecting large-scale, comprehensive changes in the economy is difficult, but if we consider all of the means at our disposal, it is by no means impossible. If we can use artificial intelligence to make new discoveries, to gain additional understanding and new levers for control, we can influence the economy in ways that were previously out of reach or which are more sensible than techniques used up to this point.

In this context, competition, cooperation, and the interactions between them are extremely important. How can cooperation and competition be arranged so that the playing field is as level as possible? Thinking globally, a perfectly level playing field is unlikely to ever be achieved. But using artificial intelligence, we can get much farther than now. Through artificial intelligence, different actors can be aware of each other's actions and also understand each other in a way that resembles a sort of internal regulation.

Economic artificial intelligence systems are already being used to promote competition, for example in stock exchanges. However, there is no obstacle preventing artificial intelligence from also analyzing cooperation and the rules that govern competition. This kind of system likely wouldn't be centralized; instead, different strategies would be used around the world.

Then different systems can support each party's own ways of operating in a complex environment.

Economic interconnectedness promotes peace between nations. Old wisdom says not to go to war against your trading partners. Then you lose the opportunity to buy and sell, which hurts you as much as anyone else. That's why it doesn't make sense to use artificial intelligence models only to optimize each party's success but rather to explicitly aim for the overall system to function well. If profits pile up lopsidedly for one party, at some point the excessive losses for the losers can also cause problems for the winner. Even the ancient Chinese military strategist Sun Tzu makes reference to this in his *Art of War*.[24] Moderation is a sensible strategy even from the perspective of economic systems.

**Information technology and society**

While I was studying information systems design at the University of Oulu in the 1980s, one of the central topics of my studies was human knowledge and how to represent it with computers.[25] Even at that time more than thirty years ago, we were focused on ideas, research, and plans related to utilizing information technology in social decision making and democracy promotion. In one interesting course, we reviewed many different models and formulations for a well-functioning society. The one I remember the best was called coinocracy.[26]

In science, careful citations are generally important, but in the following I hope I can indulge in a small amount of imprecision without attracting too much criticism. In other words, I'm relating the idea of coinocracy based on memories from thirty years ago. In a coinocracy, society and democracy would be organized by dividing people into groups of seven individuals. For each seven people one representative would be chosen to

participate in a seven-person group one level up, and from each of these one more representative would be chosen, who would again operate at a higher level making decisions for the larger group. One interesting idea in this system was that at every level every representative is forced to verify their support from below. If they are voted out from representing their group at the lower level, they would also lose their position at all the higher levels.

A seven-person meeting is quite a manageable size. In a country the size of Finland, a coinocratic system would need fewer than ten levels. Coinocracy would face many practical challenges, but it also has its appeal.

A coinocracy emphasizes people's connectedness to each other. That's why next, I'd like to investigate the ways individual existence in this world takes form. We are all parts of different cultures. Culture is often understood only as high culture, such as architecture, music, and visual art, but culture also means everyday values, practices, and habits. A culture can have aggressive or peaceful characteristics. What kinds of values and ideas guide our actions in this world? The other side of the coin is identity. Who am I? Who are you? Who are they? Who do we belong with and who are we?

I still have a strong identity as a person from Kalajoki, even though I moved away from there in 1981. For many people, their home area is beloved, even if they later moved to the other side of the world. Having lived in the metropolitan area around the capitol of Finland for many years now, I also think of Helsinki, Espoo, and Vantaa as determining parts of my identity.

As I said before, I also spent long periods of time during my youth with my brother in Switzerland, and because of that my self-conception also includes a dash of the Swiss. Because

of my studies during school, I know a little German, and I've picked up some Swiss German during my travels. Work trips to the United States, especially California, have also added a bit of American to my identity. If it was possible, I'd shop at Whole Foods every day like I did in California. To my delight, I noticed a few years ago that Whole Foods was in Texas too.

Due to the cultural influence of music, I have a healthy dose of Russian and German identity. My life would be missing something fundamental if it didn't include Tchaikovsky, Shostakovich, and Beethoven, not to mention Sibelius.

Food has also added an Asian dimension to my self-concept. Whether it's Japanese, Chinese, Thai, Vietnamese, Indian, Malaysian, or Korean food, they're all more important to me than anything from Europe or America. Actually, as a child in Kalajoki, I had a very hard time eating boiled potatoes. At five years old, I was known for smashing my potatoes up and scooting them to the edges of the plate, so it would look like some of them had been eaten. Rice eating may be genetically programmed into me.

In addition to the countries already listed, I should also list Sweden, whose influence on our values and ideas is so strong that we Finns rarely have the courage to see or admit it.

I believe that this sort of multicultural understanding of the world provides much better opportunities for happiness and success in life than not carrying around a backpack of experiences like this. I've been lucky in this regard, but I believe it's also a question of attitude. Of course, attitudes are influenced by one's life experiences, and because of this a person's attitudes always tell something about what they experienced when they were young. Diverse life experiences give self-confidence, because they teach you that you can handle different circumstances and environments.

If a person isn't used to different people and cultures, naturally they can feel unsure when they meet representatives of other cultures. Sometimes it makes sense to be reserved, but a rich life provides the opportunity to interpret the situation beyond stereotypes. Without life experience we interpret things using stereotypes and we can see danger where there isn't any.

We know, for example, that many more people die in automobile accidents than in terror attacks, and that domestic violence causes more injuries and deaths than many things we fear much more. And yet we casually climb into our cars but harbor anxiety about terrorists. Safety and the feeling of safety are thus two different things, and they don't always walk hand in hand.

Linguistic and cultural diversity are hard not to think of in terms of stereotypes. Without a common language and commonly understood customs, we can easily imagine that other people represent a threat. When we think in terms of stereotypes, we categorize people by language, skin color, religion, or some other easily recognizable superficial characteristic. However, no one is particularly happy about losing opportunities in life because of this sort of stereotyping. Finns have been forced to face these issues lately as part of a heated national conversation about immigration. It's good to remember that even Finns have been forced to move to Sweden, the United States, and Australia at various times to survive and flourish financially.

These questions aren't simple, let alone easy, so we must find ways to talk about them dispassionately and to prepare for such discussions with good information and the wisdom of life experience. Perhaps this is one area in which machines can help us in the future. Machines

can help us talk about difficult issues without our first impressions locking in people's attitudes too soon.

Self-respect is an excellent tool for keeping a cool head when discussing difficult issues. It is astonishing how easily poor self-esteem can make people treat each other badly. A good example is football hooliganism. The working class in many of England's old industrial centers has built up its identity around local football clubs. When the team performs poorly or if someone insults the team, supporters sometimes defend their identity and values with unnecessary aggression. Imagine someone in downtown Helsinki proclaiming that Finland is a terrible country. If this person is a Finn, we might argue with him. However, if the speaker is a foreigner, some might react angrily. Angry behavior can escalate to violence. A constructive approach would be to disagree and note that scientifically speaking, Finland is one of the best countries in the world. And even if that wasn't the case, threats to identity should never lead to violence.

The fact that threatening identity does easily lead to aggression is a symptom of the backwardness of human society. The more tightly bound something is to an individual's identity, the more painful it is when that thing is insulted. Of course, it is easy to question whether we'll ever get past this stage. However, one of the basic principles of this book is that we can improve ourselves and become better than we have been before. Our life today is very different from a hundred or a thousand years ago.

Self-respect and a feeling of safety also relate to fairness and justice. Justice is one of the concepts we may never come to complete agreement about. However, we can strive for the best possible result. When a sufficiently good and broadly shared concept has been formed, we can begin to promote that kind of justice. Justice and fairness relate to many

broad social questions such as the economy and social identity. If a large group of people sees a situation as unjust, that can have very negative consequences, as we saw during the interwar period in Germany. However, the elements of justice—as the concept itself—are challenging, because many contradictory needs and ideas relate to them, and therefore no perfect harmony can ever be achieved. But instead of perfection, we can try to make our society as functional as possible.

It would be great if social decision making and democracy could develop in a way the desire for money, power, and honor wouldn't lead to large distortions and conflicts. Perhaps one task of machines in the future will be to help us test reality. We can use machines to check whether we as individuals and societies are on a sensible path. For example, is a person or group we see as an enemy really a threat, and if they are at that moment, can we somehow relieve the conflict.

**Kohonen as inventor and trend-setter**

My first experience with artificial intelligence came during the Kielikone project. You might say that Apple's Siri was already being built in Finland in the 1980s. In the Kielikone project, we were trying to build a system that could interpret questions and commands presented by a user. An example of this kind of command is the sentence, "Give me the ten largest forestry companies, their operating margin, and their turnover." The system required the development of a great deal of language technology related to vocabulary, word inflection, sentence structure, phrase meaning, and contextual interpretation. Linguistically speaking, these sub-disciplines are morphology, syntax, semantics, and pragmatics. The basic goal of the Kielikone project was to give Finns an easier way to use computers. At the time, computers required

difficult commands to operate, since graphical user interfaces were not in widespread use. In the early 1980s, computers were controlled by giving them commands, and using commands required following strict syntax. Large demands were made on the user's own memory because there were hundreds of commands.

The Kielikone project did not achieve its primary goal, which is not surprising in light of current knowledge. The development of this sort of system demands enormous resources. A few million euros would only get you started. Developing a functional and sufficiently comprehensive solution using the rule-based approach back then would have required closer to a hundred million euros in today's currency. However, the Kielikone project provided important experience.

The challenges of developing a computer that can understand human language are both quantitative and qualitative. Human language is enormously complicated. There are around seven thousand individual languages in the world, and language is difficult enough if you only focus on one. First, you have to be able to model the various forms of individual words, which in Finnish there are about twelve thousand of for verbs and two thousand of for nouns. At the next level you have to pay attention to the ways words form phrases. This area is called syntax. The process related to syntax in which a computer attempts to find the structure of phrases is called parsing. Perfect parsing has been found to be nearly impossible if you want to take into account the interpretation of all words in all possible phrases. In a text as simple as a newspaper article, the machine encounters an innumerable range of options if it is forced to analyze all possible alternative interpretations of the words.

In a system like Kielikone, after the structures, words, and phrases comes the level of meaning. The study of meaning can be divided into semantics and pragmatics. In

semantics, words are interpreted without context, as in the “dictionary meaning.” Pragmatics, on the other hand, investigates what literary phrases, words, phrases, and sentences mean when they are used in a certain context.

The level of meaning is even more complicated than the level of structure. Imagine a situation in which a computer must negotiate the diversity of phrases used in everyday human conversation. If the machine’s understanding is determined by rules that experts define based on their own research and understanding, this will work awkwardly. The problems that come up will be quantitative and qualitative. The quantitative problem comes from the sheer number of rules, words, and other linguistic bits that have to be programmed into the system. It would be fair to say that a thousand variables programmed into the computer wouldn’t get you anywhere. The Cyc Project in the United States tried program all general knowledge into a computer. Fifty million dollars were expended on the project in the 1980s alone. Despite this, the result wasn’t much more than a big encyclopedia. The qualitative problem, on the other hand, comes from the fact that even if the enormous number of rules required could be programmed into the system, the rules themselves aren’t enough to describe all linguistic knowledge. A natural language, such as English, Finnish or Chinese does not easily submit to the language of logic. The language we use is per force ambiguous, because we can’t have a new word or phrase for every new situation that comes along. The essential characteristic of language is creativity, so that a huge, perhaps infinite number of different things can be expressed with a limited number of words.

Ambiguity means that a word or phrase can be interpreted in many different ways. A classic example of this in Finnish is the word “alusta”. If the word “alusta” appears in a text, it

can be an inflected form of the word “alku”, “alus” or even “alunen.” If you think hard, you can come up with at least six different base forms that you can see in the word “alusta”. Even this brief example shows that the meaning of words is difficult to program as a fixed set of limited examples.

My experience with the Kielikone project led me to seek other options for developing artificial intelligence. I discovered I didn’t have to go farther than the seashore to find fish, as we say in Finland. The Kielikone project used a way of programming knowledge and skills into computers which was common at the time. It began to become apparent that we would achieve better results if we made the machine learn. The self-organizing map is a central concept of artificial intelligence research, and it has strong roots in Finland. Professor Teuvo Kohonen at the Helsinki University of Technology realized that the organization of the human cerebral cortex can be modeled in ways that allow computers to execute similar functions.[27]

Maps in this sense don’t refer to geographic maps, but rather maps of topics, words, and other content. For example, the meanings and forms of words can provide a basis for organization. Another way to organize content like words is to note what contexts the words show up in when they are heard as parts of phrases or as written language.

Self-organization doesn’t just apply to language, though. The same principle can be applied to anything humans experience. Whether we consider the senses of sight, hearing, or touch, the principle of self-organization is that things with similar characteristics are located close to each other on the map. Teuvo Kohonen has succeeded in fashioning this principle into a simple pattern that the computer can understand and implement. Due to this discovery, Kohonen became one of Finland’s most cited scientists, because the self-organizing map algorithm has shown itself to be so useful. In addition, self-organizing maps are visually

impressive and informative. These maps are convenient for presenting complex phenomena, whether the content in question is parliamentary elections, Grimm's fairytales, or industrial data. (See inside cover pictures.)

The structure of self-organizing maps and other neural networks typically contains an input layer, i.e. a group of programmed neurons through which the desired simulated experience can be given to the system. The second fundamental part of the system is the output layer, where a group of simulated nerve cells presents the result of the neural network for users to utilize. The architecture of the neural network can contain one or more intermediate layers that the process runs through. In the case of a self-organizing map, the input layer and the map where the given inputs organize in some sensible way are sufficient.

These maps have been utilized in practice in every field of science and many sectors of the economy from finance to medicine. It's difficult to find an area where self-organizing maps haven't been used over the past two decades.

**An introduction to the peace machine**

The idea of the Peace Machine was born in January 2017. I asked a few colleagues whether this idea that had popped into my head was too Utopian. When a few of them said they found it interesting, I began developing it further. Since the 1980s I'd been interested in how machines could be used to understand how people understand language and, through language, each other. I was convinced that the logicians' ways of investigating the meaning of language were insufficient.

I'd come to this conclusion as a young researcher when I read the philosopher Ludwig Wittgenstein's work *Tractatus logico-philosophicus* (1921).[28] Even then I found the ideas presented in it too narrow.

The picture the book gave of language and the world felt too simple, as if cooking, bread making, ditch digging, and house building were being provided with the same tools and materials. The world we live in, the language we use, and our thoughts form a whole that the Wittgensteinian triad of "word", "concept", and "object" is insufficient to describe. So, as I've studied linguistic meaning and interactions, I've always tried to keep in mind how complex these things are.

Language is often presented as a tool of shared interaction and description of the world for everyone who speaks the same language. However, it isn't this simple. The more we've learned about language, the more complex a system it has turned out to be. Linguistic meanings don't automatically accompany the words, but rather the meanings of individual words are learned through experience. Because people's experiences are different, the meanings of words can also vary from person to person within the same language.

Words and meanings also don't exist as simple systems that we learn year by year as we grow. Language is born in human societies and it morphs accordingly. No one person, let alone one work, contains the entire picture of a language. Language is a tool. With language we try to understand reality in a way that is always imperfect. A tinker, a tailor, a steelworks engineer, a salesperson at a sporting goods store, and a dance teacher can't get along with the same vocabulary. Each needs a different way of describing things and using language. An individual's language is largely shared with people in the same occupation, people who are interested in the same things, or people who have the same education and similar experience.

Only some vocabulary and meanings are broadly shared with people who speak and read the same language. Tables, chairs, lawns, roads, foods, and other similar everyday things are shared broadly, but even their content isn't the same for everyone who speaks the same language when looked at individually. What we mean by food being good, tasty, or spicy means different things to different people at least to some degree.

For us to converse together, we must submit to a process that I've sometimes called meaning negotiation.[29] This describes a conversation in which we search for a common way to interpret some expression. When a conversation partner wonders at our choice of language, we can discuss what we mean. In many situations, minor differences in meaning between people don't even come up in conversation, though. Even less frequently are they highlighted in newspapers or on television, which can be a problem because those media transmit information to so many people. We can't ask a TV anchor or a reporter directly what they mean by justice, efficiency, or happiness. Wouldn't it be useful if we could get help from somewhere for this negotiation of meaning? The first pillar of the peace machine relates to language, meaning, and understanding. Meaning negotiation machines could act as an aid in situations in which people understand things in different ways and don't notice it. Meaning negotiation machines or systems that account for culture, which translate text or speech from one language to another, are a central part of the Peace Machine.

However, the Peace Machine would not simply be one machine or AI system but rather many different systems we could use in everyday situations. So, the Peace Machine is not an invention that we bring into peace negotiations at the point we're plunging

toward war. Rather, this would be an everyday tool to improve understanding between people and to increase peaceful coexistence in normal life.

The Peace Machine also isn't a turnkey system you could buy at the store tomorrow. However, many of the central ideas and approaches do already exist. The explicit purpose of this book is to inspire the imagination, so we can find new ways to harness these technologies and scientific findings to build a more peaceful and happy society.

Challenges related to linguistic meaning are not, of course, the only thing that strains the peaceful coexistence of human beings. That's why this book also deals with emotions, values, and identity. Emotions like anger, bitterness, fear, and disgust can often lead to violence. Fear can also function as a trigger for violence. Machines don't have feelings, but we can use them to analyze and study what people feel in different situations at the individual and the societal levels. Feelings are the second pillar of the Peace Machine.

In recent years we've learned a lot about the structure and development of human emotion. Machines can help us collect and assemble this information but also to handle it in a way that help us become more conscious of the development and influence of our emotions. Feelings aren't an irrational detour in human life but possibly the most important motive force in our decision making. I even believe that feelings often lead us to functional solutions to serious issue. Problems come if emotions are pathological, though. For example, because of trauma or other unpleasant experiences, anger, bitterness, disgust, and fear can overcome us. In this frame of mind, it is easy to overreact and commit excesses. In striving for a more peaceful world, we need more than an examination of interactions and emotions, though. We also must investigate identity and values. With the analysis of large data sets that machines make

possible, we can get an increasingly detailed picture of these things. Using machines, we can learn more about our own and others behavior and its antecedents. Of course, this understanding can be used for good and bad purposes. In this book I focus on the good purposes, but let us not forget the bad intentions, because if we aren't careful, they can destroy any beneficial progress.

There is a good reason for emphasizing positive thoughts, though. Our thoughts influence our behavior. If we think negatively, we're more likely to get negative than positive results. It has been demonstrated that mildly optimistic thought is most sensible. This doesn't mean that challenging realities shouldn't be kept in mind. Testing facts and trying to see reality as realistically as possible is always sensible.

An extraordinary feature of our world view is that the things we think exist in our world actually do in one way or another. Many abstract things are social constructions maintained by our thoughts. The ability of our thoughts to create or destroy the world is amazingly powerful. For example, slavery was one of these mental structures. As a society we can see other people as resources or as free individuals who can decide their own fate. Even though the thoughts of single individuals don't have much weight in this sort of thing, ideas can start in waves, and then they can have an enormous influence on the world.

The just functioning of societies is the third basic pillar of the Peace Machine. The economy and society should be organized in a way that as many people as possible can justifiably see their circumstances as equitable. However, making everyone feel they are being treated justly one hundred percent on the time is impossible. It is sensible to try to get as close as possible to this goal, though. So far, in my view, we have tried to find justice in overly simplistic ways. I dare say that neither communism nor capitalism will

suffice to organize a complex world in a way that yields sufficient satisfaction for the majority of people.

Even in the best democracies, representation is fickle, and as a result we're constantly trying to fix things that "ain't broke." Due to their backgrounds and values, people naturally have different thoughts, beliefs, and hopes about how things should be organized. At the same time, it is also true that no single person or small group of people can have a perfectly good solution from the perspective of the whole system. Complex systems are so broad that no small collection of bits—i.e. human thoughts—can give you a comprehensive view of it.

A good example of this kind of complex environment are the problems related to deepening EU integration. In different parts of Europe, conditions, culture, and ways of thinking are so different that setting common laws simply isn't sensible. Among the greatest challenges for the EU is how to order things so that the rules are the same for everyone but also take into account people's differences and perspectives in a rational way. The EU itself is a significant peace system, which requires a certain amount of unity. However, the system could be ruined by an over eagerness to integrate everything. In the most momentous matters, people and societies need wiggle room. In other words, Brussels can't and shouldn't dictate everything within EU, but disassembling the entire system would be breaking a functional machine. Reijo Vihko, who served as president of the Finnish Academy, put it well when he said that in many matters there is no reason to choose between options on an either-or principle but rather we must find both-and solutions.

Everyone can't have everything, but we can seek solutions that give as much as possible to as many as possible. Thus, we can engage in negotiations over the meaning of

justice. When we reach some sort of conclusion, we can think about how to build a system that makes this sort of justice possible, for example in the economy. If we fail in the economy, the risk of violence and social disturbances grows. Of course, this does not mean that we should accede to the loudest demands. Capitalism has also been called a peace system, because in it economic ties bind nations to each other: going to war against trading partners isn't sensible.

Self-seeking, selfishness, and all of humanity's other wretched features are factors that always have to be taken into account in a project like the Peace Machine. Some people are extremely selfish. Avarice, ambition, and the thirst for power still motivate many of us, and these themes are just as important for peace as the basic needs outlined in Maslow's hierarchy. A certain selfishness is appropriate, because without selfishness, our species never would have survived so long. Pure idealism easily turns us into well-intentioned fools, who get steamrolled when anyone ruthless and cruel enough steps on the scene. That's why, for example, I don't recommend the immediate dismantling of armies, because that would be too risky in a situation in which peace wasn't the prevailing condition.

However, it's good to have long-term goals. Twenty years ago, I heard a Japanese professor speak about how he tries to project his thoughts one hundred years into the future. This startled me into a new way of thinking, and that experience is part of the background for this book.

# CHAPTER 6: BUILDING THE PEACE MACHINE

## Artificial intelligence to promote peace

The concept for artificial intelligence that promotes peace can be divided into two parts. First, at the grass-roots level, better understanding of language and emotions can promote human interaction and understanding. Second, peace can be promoted by taking into consideration the injustices and wrongs people experience. Especially in this latter part we run into many of the concrete problems in our world. For example, people are placed in prison for their opinions, their religion, and other similar things, which makes the thought of using violence to fight for freedom from tyranny not seem completely wrong. However, the purpose of my book is not to investigate situations like that but rather how they can be prevented. When societies are organized democratically, dictatorship becomes impossible or at least unlikely.

Successful interaction is a requirement for a functional society. My own invention to promote this interaction are meaning negotiation machines. Alongside machines that conduct meaning negotiation, we could also utilize systems to support content negotiations. Developing machine translation is also important. Emeritus professor Andrew Chesterman of Helsinki University has provided a good set of guidelines for this development work. Chesterman has studied the work of human translators and suggested what sort of variables should be taken into account in translation work. If we think of machine translation applicable to the entire world, one problem that presents itself is that not all languages have enough material available to train the systems. Fortunately, this problem is ameliorated by the recent discovery that corpora from similar languages can sometimes be used to improve results for those that are less well documented.

If meaning negotiation machines can be combined with translation systems, we can have better interpretation services. They won't just translate the language, they can also take into account the cultural context where words appear. Cultural context is one reason why word-for-word translation doesn't work. The result of literal translation can be utterly incomprehensible. Poems and word play are often difficult to translate. Also, a phrase which can be entirely innocuous in one culture can sound offensive in another culture if translated literally. What is a translation machine or interpreter to do in that sort of situation? A sensitivity that takes these sorts of variables into account must be built into the system at a high meta level. In addition to having a model for word meaning, the system also needs a model for the emotional connotations words and phrases have or can have.

In terms of work related to human emotion, we can identify at least three sub-areas in the Peace Machine: the verbal expression of emotions, corporeality, and sociality. What are the analogs for emotions in a language? What can we deduce about feelings as bodily phenomena? What emotions are dynamic, social phenomena? In the development of artificial intelligence and machine learning solutions, it is critically important to have sufficiently broad and detailed data to utilize. For example, these data could apply to linguistic expressions related to emotional situations. They could also be related to the physicality of emotions, such as measurements of heartrate and blood pressure.

A good example of this emotional dynamic is sexuality. Sexuality has biological, social, psychological, and cultural dimensions. In certain life circumstances and contexts, sexuality can be a very important motivator or basis for various actions and decisions. Sexuality is also a question that divides people into us and them, and so it has a fundamental influence on how people get along with each other in the world. The sexual frustration of lonely people can lead to

very problematic behavior. Perhaps Internet dating services should also be embedded as one component of the Peace Machine.

Of course, sexuality and dating services aren't the central content of the Peace Machine, but sexuality is one of the things that should be take into account and for which we can seek new solutions. Significant tensions are connected to sexuality, and tensions generally—whatever they stem from—are generally damaging and problematic.

Reducing tension and open violence requires creative solutions. Maria Mekri, executive director of the SaferGlobal think tank, reported at the PeaceTech Forum about cultures that have found ways to manage aggressive behavior in children. An essential element was that instead of reacting negatively to aggressiveness, the standard responses were to laugh at it or ignore it. From the perspective of learning systems, this seems very sensible. Undesirable behavior is "extinguished" by ignoring it instead of rewarding it with attention. Approaches like this, which research has shown to work, are extremely important for the development of the Peace Machine.

Effectively managing emotions can promote interaction for example in situations that someone sees as unjust. People suffering from poverty may ask why society doesn't do more for their benefit. But people in good economic circumstances may ask why they must pay so much of their income and wealth in taxes. A person belonging to an ethnic minority who is taken to prison may wonder what justifies the persecution of certain ethnic groups. But the leadership of a nation may ask why they must tolerate people who try to remove them from power or who voice criticism they find unfair.

There may not be any definition of justice in the world that everyone could accept. People experience themselves, their lives, and other people near or far very differently. So, in this kind of situation, what can we do to promote a more just world? There are a few basic ideas

related to this in the Peace Machine. One is that meaning negotiations, content negotiations, and tools for promoting interaction between people provide the possibility to balance the system. This can make it possible to find functional and sufficiently broad-based solutions.

One of these solutions is flattening hierarchies and distributing decision-making to broader groups of people. In this context it is also critical to identify the stakeholders and how experts take part in decision making. Around the world the conditions and history are so different that perfect cooperation can be impossible to achieve.

In addition to coming to agreement about difficult issues, another possible solution—at least in some cases—is thematic or geographic division. The latter solution has been used for ages: people live in different countries with very different laws and customs. Because there is not free movement in the world, the system does not optimize according to the decisions people make. If the dynamics of supply and demand could be extended to this level, things might organize themselves. What I mean is that if nations don't know how to organize their business well, they would lose popularity, and if other nations organized their business much better, relatively more people would move there. However, because mass migrations cause their own problems, this isn't an easy solution either. In addition, it is difficult to see how national leadership would ever approve of this principle in any extensive way.

However, my humble suggestion is that people should be given more power than they have now but that more responsibility should be demanded of them too. Perhaps this is a Utopian ideal more than a suggestion, something we should ponder for the future. How would the world look and how would it work if every person always had the opportunity to choose where they lived? What customs are things that we should hold onto and which are things we should

peacefully give up over time? Machines, for example video games that deal with social questions, can help us contemplate these things.

**Learning peace**

How can people learn skills, attitudes, and habits that promote peace? How can a computer be used to support this kind of learning process? One potential starting point was provided by the English scholar Gregory Bateson (1904-1980), who studied semiotics, social sciences, and cybernetics.[70] Sociology studies social systems and cybernetics studies how different complex systems—such as societies—are guided or how they guide themselves. Thanks to Bateson, we have the concept of double loop learning in our toolkit.[71]

Traditionally programmed machines always work the same way when they're given the same command. A good example is opening a door. In a sense, the door has been taught to open when a certain key is turned in its lock. At first blush it doesn't sound like the door needs any additional instruction. It should just always work the same way. But what if the key has been stolen? In double loop learning, the system learns that the door should not open if it learns that the key is in the hands of a thief. Triple loop learning would be if in some situations the door could open even if the key has been stolen.

It might be that the theft is a misunderstanding, and someone just left the key in the apartment. This misunderstanding would prevent the owner from accessing his home. It might also be that the real thief needs to be detained until the police arrive or some other solution is discovered. In like manner, achieving a state of peace and other positive goals requires multiple loops of learning.

How can we promote this kind of learning? The Russian psychologist Lev Vygotsky, who was mentioned earlier and studied human learning in the early twentieth century, can tell us something about this. Despite his short life, he made important discoveries, one of which is the concept of a proximal zone of development.[72]

The proximal zone of development is an excellent concept for examining an individual's level of development and their opportunities to understand things or solve problems. The proximal zone of development falls between a person's skills and their potential area of expertise. A person can get to a certain level of skill without help, but with an experienced teacher or expert they can progress much farther in their problem solving. So, with help a person they can solve problems that aren't too far from their area of expertise. However, there is simply no sense in teaching a first grader the latest discoveries in differential calculus.

Unfortunately, there are things we don't know, and some of these things may be quite ordinary. In the future we can make use of the proximal zone of development idea as we grow our understanding of an increasingly peaceful world. We can also promote other positive things in like manner. We can learn more about things such as language, conceptual systems, emotions, the mind, and society, as well as people and societies.

**Machine translation to promote peace**

Machine translation based on multilevel neural networks makes it possible to model the various abstract levels of language without the requirement that every level of linguistic phenomena and characteristics be defined in advance. In other words, the pieces of language, the words, sentence structures, and meanings are taken into account in a way that previously always required a

human to do the analysis. A neural network makes it possible for the analysis to develop on its own as the machine churns through large collections of data.

Errors still occur, and some linguistic phenomena are more challenging than others. For example, word play will continue to cause problems for translation systems for years to come. However, this does not mean that the challenge is impossible. Humans use language creatively. We invent new words, we test the boundaries of grammar, and we play with the ambiguities of language. For example, do you understand the phrase "He fed her cat food"? Understanding this requires common knowledge and common sense—data that collects in large corpora implicitly. Using robotics, this kind of common knowledge also collects explicitly in practical situations. In the example given above, a robot could see whether the woman was being given cat food or whether a cat that belonged to a woman was being given food meant for cats.

The Finnish poet Eeva-Liisa Manner wrote "If sorrow smoked, smoke would shroud the earth. Yet under this sorrow is fire; my heart burns but is not consumed."[73] For a machine to understand the world like a person, it would require a lived human experience of the world. That is why machines will never replace us.

The sorrow of a person upon losing a loved one cannot be mimicked by a machine. When I was in high school, we had a lot of cats, and one pair of them, siblings, were practically joined at the whiskers. Unfortunately, the sister was hit by a car. The brother never recovered from losing her, instantly turning from a playful, happy scamp to a sober-minded wanderer. If the cat could have talked, it would have spoken of its smoking sorrow. A machine or a robot cannot rejoice at a birth or mourn at a death, because it is not born and does not die. A machine doesn't have experiences like that cat who lost his sister or the emotions of a child who has lost his mother. But a machine can help people understand their joys and sorrows.

In machine translation meant to promote peace, quality and comprehensiveness are important, but just as important is the cultural sensitivity of the translation. In current machine translation research and development, people often act as if there is one true translation. However, the sayings of one language do not translate to another language so simply, which can cause misunderstandings and offend people. So, if we want to promote peace, any translation must recognize context-sensitive cultural differences and translate them well. Fortunately, we know the challenges of translation at a detailed level, and we can utilize this knowledge in machine translation. We can even take into consideration the influence of target language phrases on the emotions of the reader or hearer.

Emeritus professor Andrew Chesterman, who was mentioned previously, has studied the work of human translators in various situations, such as in literary translation. An experienced translator often uses several different strategies when searching for the best ways to express a source language expression in a target language.[74] Chesterman divides translation strategies into three categories: syntactic, semantic, and pragmatic. The simplest, but only one of many ways to translate, is to do a literal translation. For example, the direct translation of the word pair "punainen talo" in Finnish is "red house" in English. The complexities of language appear immediately, however, if the word "red" is replaced with "white". In Finland if we talk about a white house, we draw our interpretation from the surrounding environment, but if the white house is in the United States, associations with a very specific White House enter the picture.

Translation solutions related to syntax, that is to sentence structure, include changing the part of speech, changing singular to plural or changing an individual word to a group of words, changing the sentence structure, and changing the rhetorical pattern. Alliteration is a stylistic device in which the same initial letter is repeated purposefully several times. For example, "Peter

Piper picked a peck of pickled peppers." Of course, alliteration and similar phenomena are not central to the Peace Machine. I mention them as examples of the often-overlooked complexities of language. Repeating a first letter may have an effect that is lost or misinterpreted in another language if not taken into account properly in the translation.

The interpreters who participate in peace negotiations work under enormous pressure, and we should admire the work they do. We are nowhere near automatic interpretation at the level humans can provide, but we should continue to strive for it. In the future, interpretation won't only be needed for negotiations and conversations between leaders—we will need tools that facilitate high quality, deep, and nuanced conversations between all people of different languages.

From a methodological perspective it is crucial that the context model of a translation system includes sufficient variables and that there are sufficient data available on the phenomena essential to the conversation. The data doesn't always have to be about relationships between the languages. Often it is enough to have separate material about each individual language. Translation work requires high dimensional tensor models and methods based on complex conditional probabilities and the like for modeling this kind of data.

Semantic strategies for finding the best equivalence are also myriad, according to Chesterman. For example, when translating, one can take advantage of synonyms or antonyms. A translation can use hyponyms and hypernyms. Expressions can be condensed or expanded. Chesterman also mentions modulation, which he explains with an example of transforming the concrete into the abstract. As language develops, the relationship between concrete and abstract can change over a long period of time. For example, the Finnish verb "keksiä" ("to invent") got its beginnings from a hooking tool used when floating timber. The many layers of language must

also be taken into consideration in translation—examining individual phrases or sentences is not enough. The larger interactive and cultural context must be considered. Translation can also take advantage of changes in metaphors and even irony.

Linguistics draws a division between semantics and pragmatics. Semantics refers to meaning devoid of context, while pragmatics investigates what words and phrases mean in context, whether the narrow surroundings of a sentence or more broadly. Chesterman lists only (!) ten translation strategies related to pragmatics. These include such things as adding or removing material, making things more explicit or implicit, and putting phrases through cultural filters to make them more or less familiar.

A classic example of cultural filtering comes from the computer program manuals of decades past. When American manuals were translated into Finnish, the translators wondered whether vocabulary related to baseball should be exchanged for the vocabulary of the Finnish bat and ball game *pesäpallo*.

Explicitness refers to content that is stated directly and implicitness to content that is communicated between the lines. We humans have a strong ability to make deductions between the lines about things that aren't stated openly. This causes its own set of problems for language interpretation. When we hear a word or phrase, along with it comes an enormous number of linguistic and factual associations, which we put into use quickly and automatically. Because of this, the interpretation of words is impossible to study in any simple way. We can't assume that people can just snap their fingers and break free of their own internal context, which they've developed over decades of life.

Another pragmatic strategy Chesterman mentions is interpersonal change, meaning changes to the level of formality of a text. For example, in Germany professors are addressed

with a “herr/frau”, “doktor” or “professor” before their name. In Finland this sort of address would sound strange and too formal in most circumstances.

In speech act theory, statements can be divided into types of acts.[75] Speech act types include observations about a state of things, requesting something, expressing one’s own state of mind, or perhaps a warning about some danger. When translating and interpreting, sometimes it is appropriate to change the type of the speech act. Solutions such as these are often cases of needing to take into consideration the different ways that the source and target languages express actions between people. A request in one language may feel like a command in another, and in some languages simply referring to the state of things might be enough. A skilled translator or interpreter has a command of these subtle nuances, and there is no fundamental reason why machines could not also accomplish this.

Implementing these sorts of precise translation strategies will still take decades, though. While we wait for that, we will see transitional systems that achieve moderate success. Errors and misunderstandings will occur, but these will be similar to those seen in communication between humans.

However, machine translation can be taken even farther—to promote empathy between people beyond linguistic and cultural boundaries. Then if misunderstandings and errors occur, the blame can be placed on the machine without any person having to suffer. There isn’t any point even trying to build a genuine ability to experience feelings into machines. The fact that a machine can analyze human emotions and use that information does not mean that the machine itself feels anything.

I would like to emphasize again that machine translation meant to promote peace must take special care to consider the linguistic and cultural background of the listener. This will

require much more data than we have available now, and computational resources must also be increased if we want to bring this kind of multilingual technology to fruition. In addition, we must remember that although we no longer need linguists to hand-code rules to guide the machines, they still have an important role in crafting strategic solutions and evaluating results. As well as linguists, we will also need anthropologists and psychologists, among others.

**What does a meaning negotiation machine have to take into account?**

Humans live constantly immersed in language. This is so natural for our species that we usually don't even notice it. Humanity has learned to use language as a tool, and each of us is initiated into language at such an early age that we have an extremely difficult time extricating ourselves from it. We learn to use language in different contexts using referential relationships—for example, we can use words to refer to different things we've seen in the world. The role of language in thought is important, but studies have shown that language is not our only way to think.

In language, words are each other's support pillars. We can use words without them having any immediate referential relationship to the surrounding environment at that moment. Our social reality is largely built up using language and words. For example, in a way, the organization of society is a sort of agreement in which the role of language is fundamental. We can't form social systems and agreements without language.

Words are such strong tools because they are easy to remember, and we can refer to them in a different way than to the broad, concrete phenomena or the social systems surrounding us. How could we arrange matters related to travel, education, or trade without language? If we want to agree on something as simple as a bicycle outing on a summer vacation, just waving our hands

around without language and symbols would be very difficult. We would have to agree on the starting point and the destination, and perhaps even on the route between the two. This might work to some degree using a map, but maps also have linguistic elements, so using them would be cheating. It would be fairer to use an aerial photo. Arranging the time would also require the use of symbols, numbers, or words.

Words are handy tools for referring to concrete and abstract things in the world. Words also simplify things. For example, the word "cat" refers either to the idea of cats, all cats as a group, or, if thought of from a modernist perspective, the division of the prototypical cat at the center surrounded by individual instances that vary from the prototype. People have been taking up cudgels over this theme for thousands of years. Through the process of human cultural evolution, language has gradually become an increasingly complex system. Nowadays it is common to think that the purpose of language is to transmit information from one person to another, but the central historical role of language has been to coordinate human action. In other words, we don't just use language to interact with each other. Rather, it is in constant use as we parse the world and act independently. We coordinate our own activities by using tools such as notes and lists. We don't need any external actor to remind us, since we can use words to remind us what we thought to do the next day or what we planned decades ago. Our capacity for language is based on brain activity made up of many components and functions. One of these is memory. The memory system related to language is very large but not limitless. Language use itself is a sign of the limitation of memory: language simplifies reality, so our memory isn't overloaded. Animals, such as dogs, also have different memory systems. Dogs can recognize their names and simple commands.

However, the human capacity for language is much broader and more complex than that of animals, and human memory capacity allows the learning of thousands upon thousands of different word forms and meanings. When we want, we can also form new words and suggest that others adopt them. Language is a tool we can use to plan things that don't exist yet. We can also use language to communicate ideas about how the tool of language can be further developed.

The surrounding reality—and the language we use with its symbolic systems—are fundamentally statistical phenomena. We can easily mistake language for a purely logical phenomenon if we look at it one word or sentence at a time. But let us forget right now the unrealistic idea that meaning is the same for all people. Language is a system that is constantly changing and bound to context. Sometimes societies will intentionally try to design words, for example to replace loan words. Examples of this from Finnish are the words "muovi" and "puhelin", which were made up to replace the loan words "plastiikki" and "telefooni". Just as often people reject this sort of top-down language policy, though.

I can't emphasize enough how important it is to recognize the social and the individual nature of language and that they walk hand in hand. Forgetting either perspective can lead to arguments with people who study language. Linguists' approach is usually either narrow and precise or broad and nebulous. For example, in analytic philosophy, language is examined with great exactness, but at the same time many phenomena are excluded from analysis. Perhaps as a reaction to this sort of thought, postmodern philosophy engages in conversations about language and other phenomena that lack the precision of the analytic philosophers.

The question is not which of these groups—or any other group—are right or wrong, since all of them are both right and wrong to varying degrees. In the same way as in social questions,

as we model language and develop linguistic tools, we must find complex solutions between the options, instead of making black-and-white choices of one option over another. The investigation of small subsets of reality, as analytic philosophy engages in when looking at language, must lead to larger aggregates. Postmodern thought, on the other hand, is too vague to be applied in practice. Combining broad, diverse thought with precise, careful methodological work is a challenge this book cannot tackle. Many books need to be written on this specific theme, and many different systems need developing. In this context I simply try to present what I've learned and the basis upon which I believe these theoretical perspectives can be adopted in practice. That is why we now once again return to meaning negotiations and their automation.