Inaugural lecture prof. Bert de Vries

In situ personalization of signal processing systems

Presented on September 13, 2013
at Eindhoven University of Technology

Introduction

"Design is more the art of preserving changeability than it is the art of achieving perfection."' [1, pg.16]
— Sandi Metz, American computer scientist

Ladies and Gentlemen,

A recent study revealed that one in five Americans older than twelve suffer from hearing loss to a degree that impacts their day-to-day communication [2]. In another study hearing loss is found to speed up dementia and cognitive decline in older people [3]. I could cite a large list of similarly alarming literature, but, in short, hearing loss is a major problem that impacts people's lives and causes significant costs to be incurred by society.

Figure 1

A modern behind-the-ear hearing aid (source: http://gnresound.com)

I am a researcher at GN ReSound, which is one of the world's leading hearing aid manufacturers. Both at ReSound and at the university, my work is mostly focused on improving sound processing algorithms for hearing aids. You may be wondering: what is an algorithm? An algorithm is a recipe, i.e., a precise description of how a difficult task can be accomplished by executing a sequence of simpler tasks. For a hearing aid, the difficult task is to change an audio signal into another audio signal that allows a hearing impaired listener to maintain or restore normal interactions with other people. The list of simpler tasks, the algorithm, would be something like recording a few audio samples, taking the square root, multiplying by ten and sending the result to the hearing aid speaker. Then repeating this procedure forever or until the battery dies. In reality, the math of a hearing aid algorithm is a bit more complex, of course. Sound processing algorithms are important for hearing aid manufacturers because patients buy hearing aids for no other reason than to listen to a different sound, not because they like to make a fashion statement with a neat hearing aid design.

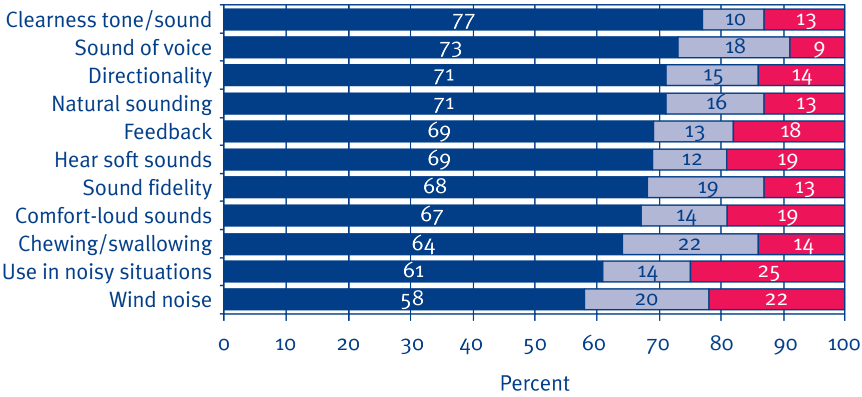

Figure 2

Percentages of patient satisfaction with sound processing in hearing aids (source: [4])

How satisfied are users of hearing aids? The graph in Figure 2 is from a large study in 2010 on the hearing aids market [4]. The horizontal bars represent patient satisfaction rates with various aspects of sound processing in hearing aids. Categories include issues such as performance of the hearing aid in windy conditions or clarity of the sound. The dark blue color reflects the percentage of people that are happy, red indicates dissatisfaction and light blue relates to a neutral opinion. Let's keep this simple: from this graph I read that about 20% of hearing aids patients are not happy with the sound processing performance of their devices. 20% of end users are *not* satisfied! It turns out that a decade ago the dissatisfaction rate was also around 20%.

This is a remarkable figure because over the past decade, we, the engineers and scientists in the hearing aids industry, have collectively spent about 1000 man-years on improving the sound processing in hearing aids. This is an under-estimation, but it is enough to make my point: despite a very extensive collective engineering effort in the hearing aids industry, one in five patients remains not satisfied. The performance of sound processing in hearing aids seems to have plateaued.

It could be that especially the hearing aids industry attracts bad engineers and scientists, but I think there is a more plausible explanation. When an engineer designs a hearing aid, he does not know yet who the patient will be, what the hearing loss portrait of that patient will be like or in which acoustic environments the patient will spend his time. To complicate matters, this type of knowledge changes over time. Every time a patient puts in his hearing aid, the physical placement of the device will be a bit different from the last time, leading to an altered acoustical situation in and around the ear. Similarly, gradual build-up of ear wax in the ear will affect the signal processing demands on the hearing aid. In other words, when the engineer designs the sound processing properties of a hearing aid, he has to deal with many unknowns about the actual circumstances in which the hearing aid will be used, which means that the engineer does not know which problem he must solve beforehand.

Apparently, all these uncertainties about the problem to be solved led to an unsatisfying experience for one in five patients. It won't help to ask the engineer to work harder or to try his extra very best this time, since these future in situ conditions are simply unknown. Instead, we must provide the patient tools to solve problems on the spot when they occur. The property that allows a system to change in response to changing circumstances is called *adaptability*.

The challenge to build a system that works both today and can still be changed in the future is somewhat of a contradiction. On the one hand, designing a system that works today means we must make hard decisions. We must decide which low-pass filters, resistors and transistors to use and how to connect them in order to make a system that works. On the other hand, the device must be able to adapt to new circumstances and consequently our design cannot be fully committed yet. Apparently, good design involves the art of postponing decisions until later, when more information is available since you will never know less than you do right now.

There is a discipline called *adaptive signal processing* that deals specifically with these issues. In practice, adaptive signal processing is a very successful field: all working hearing aids, mobile phones, TV sets, car radios and medical measurement equipment critically depend on adaptive filters inside. At the same time, adaptive signal processing today has failed to provide answers to the problems that I mentioned above. Apparently something is still missing.

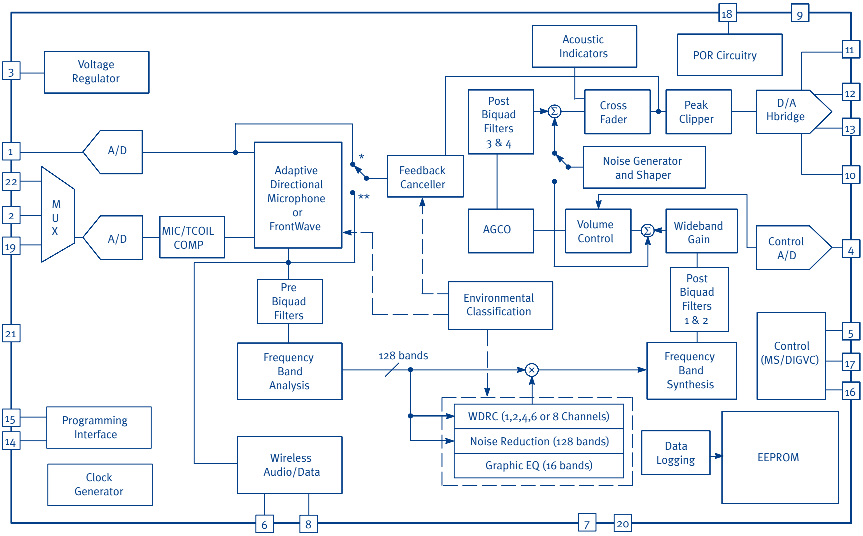

Figure 3

Block diagram of the AYRE-SA3291 hearing aid algorithm (source: [5])

My research is about that missing part. Since only a patient himself knows what he hears, the challenge is to build tools that help a patient redesign the signal processing system in case he is not satisfied. In order to get an idea of what we are up against, have a look at Figure 3 which is a block diagram of a commercial hearing aid algorithm [5]. Most blocks in this graph hide sub-algorithms that are at least as complex as this top-level diagram. Now suppose that I am at a cocktail party and I can't understand my conversation partner. I would like to make a small change to this circuit and test a few variants, but how? If I pull a wire, this circuit will probably crash and no output gets generated. Which wire should I pull anyway? Or should I add a wire somewhere? In other words, how do I bring about variation as a means for experimentation in this system? Even if I succeed in improving the quality for my current situation, will that change still be an improvement later, after the party is over? In practice, the way to update systems like this one is to give it to an expert signal processing engineer and let him tinker with it. Then take it back after a few months and hope that it works better. But that is not what I am interested in here. If this system doesn't work to my full satisfaction in the field, I'm willing to invest a maximum of one minute to make it sound better. And it should sound better because if it doesn't, I will be less willing to spend that minute the next time. This signal processing circuit is not fit for that purpose.

Figure 4

The cocktail party problem relates to separating target sounds from interfering sounds (source: [6])

After this long introduction let me come to the focus of my research plans. I am interested in *fast and easy re-design of real-world signal processing systems by end users*. There are three very challenging aspects about this goal. Whereas a normal design update by a signal processing expert in his laboratory environment may take a few months, my aim is to execute an incremental design update (1) by a (non-expert) user, (2) in normal operational conditions and (3) within a minute. In my search for answers to these problems I am mainly inspired by research from others on how the brain processes information. In this lecture I will discuss some aspects of brain computation that, in my opinion, should influence future engineering practice. However, before we turn to the brain, I will discuss an important engineering lesson for the design of systems with large uncertainties.

Design for redesign - the flight of the Gossamer Condor

"Good judgment comes from experience, and experience comes from bad judgment."
—Attributed to Mullah Nasrudin, a sufi wise man in the 13th century

In 1959, the British industrialist Henry Kremer announced a prize of £50,000 (in today's money worth about 1 million euros) for the first successful human-powered flight around a figure-of-eight course with the two turning points placed half a mile apart. A second prize of £100,000 was created for the first human-powered flight across the English Channel.

Many years and 50 failed attempts passed. In 1977, the British aviation engineer Paul MacCready took on the challenge and noticed a common pattern when he studied the records of past attempts. Previous engineering teams had often invested more than a year to carefully design a prototype plane based on elaborate theories and conjecture. Then, a few seconds after take-off of the maiden flight, a year's work would crash on the ground and obliterate the massive effort.

Figure 5

The Gossamer Albatross, the first human-powered airplane that crossed the English Channel (source: [7])

MacCready came to a crucial insight. Past efforts had been focused on solving the wrong problem. The essential problem was not how to design a human-powered airplane. Instead, the essential problem was that *they did not understand the problem* [7]. Rather than attempting to design an optimal aircraft, MacCready reformulated the problem as the quest to design an airplane that could be re-built in hours, not months. His team started building planes from cheap and light aluminum tubing, mylar, wires and scotch tape. In MacCready's approach, design was to be interpreted as an experiment to learn more about the problem. The first flight failed right away. But the team learned from the crash and delivered a second prototype just a few hours later. This process of fast iterative redesign continued for about half a year until 23 August 1977, when Bryan Allen of MacCready's team pedaled the Gossamer Condor for the 223rd time and cleared the finish line 7 minutes and 27 seconds after take-off. Two years later, Allen flew a further evolution of the Gossamer ('the Albatross') across the English Channel to claim the second Kremer prize, cf. Figure 5.

Where other teams had failed for more than 17 years, MacCready's fast iteration approach turned out to be the key to solving poorly understood engineering problems. While on the surface this story has little to do with hearing aid design, the underlying challenge to cope with a poorly understood problem is the same for both tasks. In addition to the psychological demand for end users to minimize personal investment in redesign, this story illuminates the *engineering need* to focus on fast redesign of hearing aid sound processing algorithms, instead of a research focus on the optimal algorithm per se.

Information processing and the brain

“Nothing in biology makes sense except in the light of evolution.” [8]
—Theodosius Dobzhansky (1900-1975), Ukrainian/American biologist

Engineers who design artificial systems study the brain for inspiration. Since the brain is the most crucial instrument in our drive to survive, it must work today and yet be fully prepared to adapt to unforeseen new circumstances. In the next sections I will discuss a few salient properties that enable the brain to execute fast redesign iterations so as to cope with a world where the problems keep changing in unpredictable ways.

Probability Theory

“The 50-50-90 rule: Anytime you have a 50-50 chance of getting something right, there’s a 90% probability you’ll get it wrong.”
—Andy Rooney (1919-2011), American radio and television writer

If the brain is a system that processes information, then there must be some computing rules that the brain adheres to. There is strong scientific support for the claim that brains compute with the rules of probability theory. This is the same probability theory that we all learned to love and hate in high school.

We can use probability theory to predict the future, based on observations from the past. For instance, if I observe 100 coin tosses and 96 out of 100 throws came up tails, then I predict that the 101st observation will come up tails with higher probability than for heads. Intuitively, this happens by extrapolating past observations. Technically, in order to predict the future we need to build a *model* to summarize regularities that were present in past observations and use that model to predict the future. We humans need to have some capacity to predict the future, because we want to avoid being surprised by the physical world around us. For instance, we must be able to make predictions on what’s edible or hostile to us. More generally, any major surprise in the physical world could possibly kill us.

So, a key task of the brain is to build a model for the world in which we live and use that model to make predictions about that world.

Probability theory can be used to make optimal predictions about future (data) observations by

$$\underbrace{\Pr(\text{future} \mid \text{data})}_{\substack{\text{data-based}\\ \text{prediction of future}}} = \sum_{\substack{\text{all}\\ \text{models}}} \underbrace{\Pr(\text{future} \mid \text{model})}_{\substack{\text{model-based}\\ \text{prediction of future}}} \times \underbrace{\Pr(\text{model} \mid \text{data})}_{\substack{\text{model based on}\\ \text{past observations}}}$$

The expression Pr(.) here is a mathematical notation for a probability mass function, but we will not bother with explaining the details of the formula, other than to point out that something as complex as predicting the future can be captured by a single-line equation. The left-hand side states that I want to predict the future from past data. The data refers to observations from the outside world that enter the brain through sensory organs like the eyes or ears. The right-hand side states how predictions of data relate to a model and past observations. The model can be implemented by a brain or by a computer program. The right-most factor, Pr(model | data), captures what the model has learned from past data. By another rather simple manipulation with probability theory we can express how models learn from data:

$$\underbrace{\Pr(\text{model} \mid \text{data})}_{\substack{\text{(corrected)}\\ \text{model after}\\ \text{learning}}} \times \underbrace{\Pr(\text{data})}_{\substack{\text{evaluates}\\ \text{model}\\ \text{performance}}} = \underbrace{\Pr(\text{data} \mid \text{model})}_{\substack{\text{predictions:}\\ \text{how likely are}\\ \text{observations?}}} \times \underbrace{\Pr(\text{model})}_{\substack{\text{model}\\ \text{before}\\ \text{learning}}}$$

In probability theory this equation is known as *Bayes rule*. Bayes rule describes how we learn about the world. It doesn't matter if the observations relate to music, video or even financial stock rates: Bayes rule applies and tells us how to optimally update our knowledge about a phenomenon based on new observations about that phenomenon. Bayes rule is basically a prediction-correction method. The model gets updated on the basis of differences between actual and (synthesized) predicted observations.

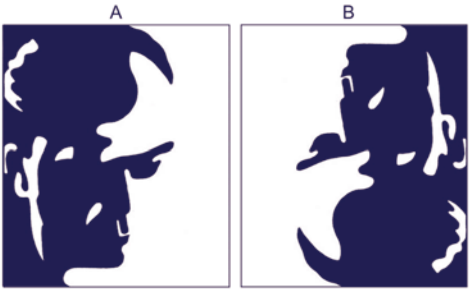

Figure 6

Mooney faces. The left image is easily perceived as a face, while the right image is not as easily recognized as an upside-down face

Mooney faces are an interesting visual demonstration of the predictive nature of the brain, cf. Figure 6. Since we have seen many faces before, our expectations for seeing faces are large and consequently we easily recognize image A as a face. In image B, the same face is depicted upside-down, but we don't perceive this face as easily. This is because our brain does not expect to encounter upside-down faces often in natural circumstances.

If a human brain were capable of executing Bayes rule, then my concept of what a tree looks like would get updated every time when I see a tree. The more trees I see, the better I understand what a tree looks like. It seems that it would be very useful for my brain to be able to process sensory information by Bayes rule, because it would enable me to learn a model about the world just by looking at the world. Apparently, using the same rules from probability theory I can then use that model to make predictions about the world, which are so crucial for me to stay alive.

It can be shown that, under some very agreeable assumptions, Bayes rule prescribes the *optimal* method for learning from observations. So there is no need to look for a specialized learning algorithm that works particularly well for any specific problem. The simplicity of Bayes rule is a strength. Whether I have to learn a language or learn about how to repair my bicycle, Bayes rule is how I *should* learn. If the brain computes with probability theory then there is no need to invent new prediction or detection methods when the outside world changes. It doesn't matter if the observed signals are acoustic or visual in nature, the difficulty lies mostly in how to *implement* Bayes rule, both in brains and computers.

The realization of the extreme relevance of probability theory for information processing and for science and engineering in general is still evolving. In this vein I would like to quote one of the leading mathematicians of our time, David Mumford:

> *"For over two millennia, Aristotle's logic has ruled over the thinking of western intellectuals. All precise theories, all scientific models, even models of the process of thinking itself, have in principle conformed to the straight-jacket of logic. But from its shady beginnings devising gambling strategies and counting corpses in medieval London, probability theory and statistical inference now emerge as better foundations for scientific models, especially those of the process of thinking and as essential ingredients of theoretical mathematics, even the foundations of mathematics itself. We propose that this sea change in our perspective will affect virtually all of mathematics in the next century."* [9]
> — David Mumford (1937-), American mathematician

The probabilities that we discussed relate data to models and back again. Models and data are very much the core issues for engineered signal processing systems. Next we take a look at how the brain deals with models from the perspective of adaptability.

Models and Structures

> *"Simplicity - the art of maximizing the amount of work not done - is essential."* [10]
> — The 10th principle of the Agile Manifesto

Signal processing algorithms can be intuitively visualized by block diagrams like the one for the hearing aid example (Figure 3). A block diagram consists of a set of blocks (nodes) and links (edges) that connect the blocks. With each link we associate a variable in the system. In a block, mathematical relations between the connected variables are described. Often, we may find another block diagram in a block, so blocks can be used to hide details of the algorithm. The algorithm *structure* refers to the mathematical relations between the variables that are described by a block diagram. We also like to distinguish between variables whose values change as time moves on (the *state* variables) and those (the *parameters*) whose values are expected to stay fixed or change much more slowly than the rate of change of the states. In neural terms, the structure relates to the neuronal network of the brain, the parameters are represented by the strength of synaptic connections between neurons and the state relates to the electric fields in the brain. In particular, our perception of the world is represented by the state variables. The model structure and parameter values provide constraints on how the states (read: our perception) will change over time. If our perceptions and prediction of future perceptions are accurate enough, we can stay alive.

Unfortunately, unexpected things will happen and we will need to change the algorithm structure and parameter values so as to keep our model of the world sufficiently accurate.

It is clear that if we change a structure at one location, we do not want that change to have serious consequences on variables in another location of the network. If the network were now to be adapted at the latter location, this could have effects elsewhere again and thus lead to a *snowball* effect of unpredictable changes, likely to be followed by a crash of the algorithm. Therefore, *modularity* is an essential characteristic of complex yet adaptable networks. A modular network is composed of sub-networks called modules with more dependencies within the modules than between the modules. The relative independence of modules prevents the snowball effect of changes from escalating.

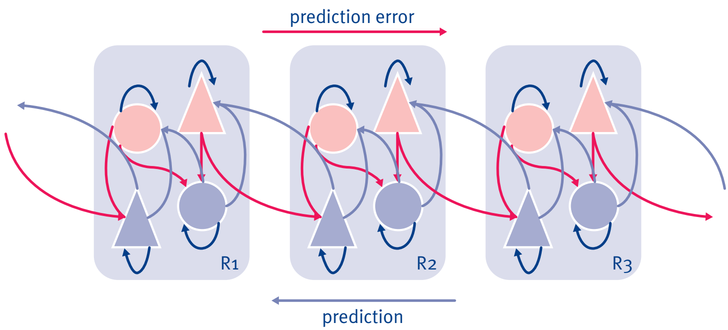

Figure 7

An example flow graph of hierarchical modularity across three cortical regions (source: [11])

On the other hand, some communication between modules is necessary to generate behavior that transcends the functional complexity of individual modules. In order to avoid the snowball effect, modules should preferably depend on other modules that are *more stable* than themselves. Let's assume the opposite, namely that module A depends on module B and the natural rate of change for B is faster than for A. In that case, A will have to adapt each time that B changes, which is more often than A's natural rate of change. The idea that the snowball-of-changes effect can be avoided by constraining communication between modules to flow from more to less stable structures leads to hierarchical networks.

Technically, probability theory supports hierarchical modularity almost effortlessly. Bayes rule decomposes into a hierarchy of four modules by

$$
\begin{aligned}
\Pr(\text{model} \mid \text{data}) &\propto \Pr(\text{data} \mid \text{model}) \times \Pr(\text{model}) = \\
&\Pr(\text{data} \mid \text{states}, \text{parameters}, \text{structure}) && \text{(now)} \\
\times &\Pr(\text{states} \mid \text{parameters}, \text{structure}) && \text{(short-term memory)} \\
\times &\Pr(\text{parameters} \mid \text{structure}) && \text{(mid-term)} \\
\times &\Pr(\text{structure}) && \text{(long-term)}
\end{aligned}
$$

In the final result of the computation, Pr(model | data), the model depends directly on fast fluctuations in the observed data. However, after the hierarchical decomposition, at each level, variables only depend on other variables that are more stable than themselves. We now have an answer to our question on *how* to implement Bayes rule. Through hierarchical modularity the snowball effect of changes is avoided.

In order to appreciate the advantages of hierarchical modularity ('modules-within-modules') I quote Herbert Simon who illustrated the concept by the watchmakers parable as follows:

> *"There once were two watchmakers, named Hora and Tempus, who made very fine watches. The phones in their workshops rang frequently and new customers were constantly calling them. However, Hora prospered while Tempus became poorer and poorer. In the end, Tempus lost his shop. What was the reason behind this?*
> *The watches consisted of about 1000 parts each. The watches that Tempus made were designed such that, when he had to put down a partly assembled watch, it immediately fell into pieces and had to be reassembled from the basic elements. Hora had designed his watches so that he could put together sub-assemblies of about ten components each, and each sub-assembly could be put down without falling apart. Ten of these subassemblies could be put together to make a larger sub-assembly, and ten of the larger sub-assemblies constituted the whole watch."* [12]
> —Herbert Simon (1916-2001), American scientist and cognitive psychologist

We can think of many reasons why modularity is the most prominent feature of adaptable systems. But how would the brain know that? Is there an evolutionary drive for brains to develop modular structures? If we accept that the brain is mostly an engine for probabilistic reasoning, then it would help if probability theory would prefer modular over densely coupled structures (all else being equal). This is indeed the case. The factor Pr(data) in Bayes rule, known as the *evidence*, can be used to evaluate how well a model summarizes a set of observations. It can be mathematically shown that the (logarithm of the) evidence decomposes into a sum of two terms, namely *accuracy* plus *model simplicity*:

$$\begin{aligned}\log(\text{evidence}) &= \text{accuracy} + \text{simplicity} \\ &= \text{'works today'} + \text{'works tomorrow'}\end{aligned}$$

The first term, ‘accuracy’, measures how well the model predicts past observations. If we need to predict future observations, it makes sense that we prefer models that performed best on past data, so we want models with high accuracy. A system that scores high in terms of accuracy works well today. However, the second term, ‘simplicity’, favors models that are simple and adaptable. Indeed, it can be shown that modular structures score higher simplicity values than densely coupled systems. As we discussed, modular systems are more adaptable than coupled systems. Therefore, probability theory prefers structures that balance excellent performance today (high accuracy) against adaptability for tomorrow (high simplicity). If you do appreciate mathematics a bit, then it’s quite a thrill to realize how this extremely practical conclusion follows in just a few simple manipulations with probability theory. We also conclude that if brains followed probability theory, then it should come as no surprise that brains are both excellent performers today and yet remain very adaptable. After all, both properties are highly prioritized by straight probability theory. The brain has no choice but to optimize both for today *and* an unknown future.

Driven by Data

> *“No organism can afford to be conscious of matters with which it could deal at unconscious levels.”* [13]
> —Gregory Bateson (1904-1980), English anthropologist

We discussed how probabilities and models relate to information processing in the brain. The third term in the equations for learning and prediction is called the data or observations. Data is observed through sight, hearing, taste, smell and touch, collectively known as our senses. Observations inform us about the current state of the world around us. We use probability theory to summarize observations in models and use these models to predict how the world evolves.

One of the most interesting aspects of our brain is how much we seem to learn from just a few teaching events. After a mother has shown her two-year old daughter a few times what a tree looks like, the girl is able to identify new trees that she has not seen before and also to discriminate trees from other plants in general. Considering the various shapes, sizes and colors that apply to trees, it would be impossible for a child to learn to reliably recognize trees from just a few remarks by her mother. Instead, a child learns what trees look like through building models straight from incoming visual data. There is no teacher involved here. The interaction with her mother just added a label ('tree') to the concept of a tree that had already been acquired through modeling the world in an unconscious fashion. In the machine learning field, learning without a teacher is called 'unsupervised learning'.

The human cortex holds about 10^{14} configurable synapses, which can be considered parameters of the brain. We live about 10^9 seconds, so on average there is room to train about 100,000 synapses every second. Indeed, brains receive a massive amount of data through the senses. For instance, the retina sends more than 10 million bits of data to the brain every second. The role of teachers, parents, books and other sources of abstract information is mostly to help us sort out which parts of these incoming data streams are important or should be ignored. In other words, teachers help us to select and label data streams that are used to train a model of the world. Crucially, in order to cope with a world where the settings and problems keep changing, a massive amount of unsupervised learning must always be going on.

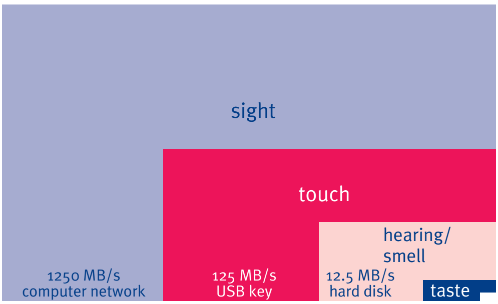

Figure 8

Data rates for the senses relative to bandwidth of computer networks (source: [14])

In Figure 8, data visualization artist Dave McCandles, based on work by the Danish science writer Tor Nørretranders, graphically displayed the amount of information that the different senses pass on to the brain in comparison to the bandwidth of computer networks [14]. Clearly, vision is the dominant sense. The white box in the lower-right corner represents the amount of data (0.7%) that is processed consciously in relation to the colored planes that refer to unconscious processing. Apparently, almost all incoming data is processed unconsciously. Building models of the world, including creating a model for what a tree looks like, is mostly an unconscious process.

Summary of information processing in the brain

> "*The principal activities of brains are making changes in themselves.*" [15]
> — Marvin Minsky (1927-), American cognitive scientist

In order to adapt to unforeseen changing conditions, brains need to iterate through new model proposals for explaining the world. In the past three sections on information processing in the brain, we found three crucial ingredients for adapting information processing systems to newly acquired evidence. The first ingredient concerns probability theory as a foundational calculus. Probability theory prescribes how to learn and predict in a world where noise obscures the signals, where observations are scarce and where people's preferences change. The second principle relates to hierarchical modularity. In order to discover better algorithms, we need to test alternatives to existing algorithms and at the same time remain operational. We can only introduce a change to an existing algorithm if the effect of the change does not cause other parts of the algorithm to crash. We must survive the change and modularity is a crucial structural element so as to limit the impact of changes throughout the algorithm. Finally, when talking about the data we noted that the structure of real world data is so rich and volatile that we cannot rely on teachers, parents, scientists and engineers to design and update the algorithm. Surviving in the real world implies a massive amount of unsupervised learning, which is always going on in the background. In engineering terms, continuous calibration is essential.

Signal processing systems that work today and tomorrow

“No plan survives first contact with the enemy. What matters is how quickly the leader is able to adapt.” [16]
— Tim Harford (1973-), British economist and journalist

Most of this lecture has been dedicated to a review of data processing in the brain. Let us now get back to the engineering practice. We left this topic about half an hour ago when we were stuck with a block diagram of a hearing aid algorithm. I was at a party and did not understand my conversation partner. I then wanted to test some variants of the hearing aid algorithm right there when the problem occurred, but the system looked so complicated that any ideas on how to change the circuit were hard to come by. My feeling was that if I changed anything, the algorithm would probably crash.

But let us assume that I have managed the dependencies between modules in such a way that I have enough confidence that a small change will not kill the algorithm. Then I can introduce some small changes to the hearing aid algorithm and with a bit of luck I can improve my listening experience at the party. The next question is now whether the hearing aid should stick to this new configuration after I have left the party. I gave the hearing aid some new information, namely I showed the hearing aid how to behave when I’m at a cocktail party. How relevant was that information for other acoustic environments? When the party is over, and I’m in my car driving home, the hearing aid has two possible algorithms to choose from: the one that I came to the party with and the other algorithm that I preferred while the party was alive. Since I don’t want to keep fiddling with my hearing aid every time when something changes in the acoustic environment, I want the hearing aid to decide for me.

In order to answer this question, the hearing aid would have to consider what features of the cocktail party environment were so favorable for the second rather than the first algorithm and it would have to consider if or how many of these features remain active in the current car environment. In other words, the hearing aid should have access to a model of the acoustic world and it should be capable of solving what-if questions based on information that is preserved by the world model it has built by unsupervised training on past acoustic observations. In principle, this seems possible since a hearing aid microphone records one million bits of acoustic data every four seconds. This continuous data stream should be summarized by a hierarchically organized structure, which is a necessary ingredient for the model to stay changeable, so it *can* adapt as new data get recorded. We have also discussed that the model should practice Bayesian reasoning in order to assess how much to adapt.

Figure 9

Other applications include (ambulant) monitoring: will this system work if one sensor contact breaks down?

Even for this basic algorithm exploration example we find that probability theory, modularity and unsupervised learning are very relevant. I think these concepts hold promise for a much wider range of engineering problems. For instance, in our Signal Processing Systems group, there are ongoing engineering research projects for such application areas as medical diagnosis, multimedia and communication equipment, a sustainable lifestyle, biometrics and control of energy consumption in buildings. For most of these projects, the crucial research question is this: what happens when I take this nicely designed system outside? What happens if it's windy? Will it still work if I open the window or what happens if the patient moves or sweats? The problem that I study is therefore not really a hearing aids issue per se. For all of these problems, there is so much uncertainty about the operational conditions that upfront design will fail quickly. We just don't have enough information in advance to design a perfect solution for every end user in every situation. Instead, we should design something that is adequate today and is adaptable with respect to the future.

Our brain, the most awesome signal processing engine in the world, has been put to test quite thoroughly over the past few million years. Still, the engineering field of adaptive signal processing remains poorly inspired by biological computing.
I hope that I have been able to convey why I think that neural computing is highly relevant for engineering and why it is my main inspiration in developing new ideas for both hearing aids and other real-world signal processing systems.

Acknowledgments / dankwoord

In closing I would like to take the opportunity to extend my sincerest appreciation for the support that I have received from colleagues, friends and family.

First, I am grateful to my employer GN ReSound and in particular to Andrew Dittberner, VP Research, for his trust, friendship and ongoing professional support in this endeavor. I am much indebted to my colleagues at ReSound, both for the professional challenges and the team spirit. Almer van den Berg, Marco Bloemendaal, Job Geurts, Anne Hendrikse, Ernst van de Kerkhof, Joris Kraak, Rudie Landman, Jos Leenen, Magali Monod-Cotte, Gerard Nijsse, Peter Snoeren, Judith Verberne, Rob de Vries, Erik van der Werf and my fellow researchers at our Copenhagen and Chicago sites, it is a pleasure and honor to work with you. Tom Kersten and Joris Kraak, you know that this talk is basically about good software engineering practice. Thanks for sharing your insights; I hope I didn't blow it.

At TU/e, I am much indebted to Jan Bergmans, chair of the signal processing systems group, and Ton Backx, dean of the Department of Electrical Engineering. Jan in particular has been a driving force behind this appointment. Jan, your coaching and leadership of the group is a true inspiration to me. I also feel lucky to collaborate with Tjalling Tjalkens in our graduate course on adaptive information processing. Over the past seven years I have participated in projects with a great bunch of TU/e students. Job, Paul, Jorik, Ronnie, Anton, Serkan, Rene, Xueru, Jianfeng, Joris, Maarten, Zijian, Timur, Mojtaba, Iris, Marija and Marno, it has been a privilege to work with you.

Finally, this whole thing would make no sense to me without my family. Ik heb het geluk gehad om op te groeien in een gezin met twee fantastische zussen Karin en Marleen en zeer toegewijde ouders, die mij tot op de dag van vandaag aanmoedigen en onvoorwaardelijk steunen. Deze steun heeft zich voortgezet in mijn leven samen met mijn lieve vrouw Oksana, die mij de energie geeft om de toekomst met vertrouwen en enthousiasme tegemoet te zien. Mijn dank aan jullie is groter dan ik hier kan uitdrukken.

Ik heb gezegd.

References

Note: all web references were successfully accessed in July 2013.

1. Sandi Metz, Practical Object-Oriented Design in Ruby. Addison-Wesley, 2012.
2. Frank Lin, John Niparko and Luigi Ferrucci, Hearing Loss Prevalence in the United States. Arch. Intern. Medicine, 171-20, 2011.
3. Frank Lin et al., Hearing Loss and Cognitive Decline in Older Adults. JAMA Internal Medicine, 2013.
4. Sergei Kochkin, MarkeTrak VIII: Consumer satisfaction with hearing aids is slowly increasing. The Hearing Journal, 63-10, 2010.
5. ON Semiconductor, Datasheet for AYRE SA3291 Preconfigured DSP system for hearing aids. http://onsemi.com, 2013.
6. Tanya Lewis, Mystery of 'cocktail party' hearing solved. http://goo.gl/VxTJJ, 2013.
7. Aza Raskin, Wanna Solve Impossible Problems? Find Ways to Fail Quicker. http://goo.gl/9zX3L, 2011.
8. Theodosius Dobzhansky, Nothing in Biology Makes Sense Except in the Light of Evolution. The American Biology Teacher, 35:125-129, 1973.
9. David Mumford, The Dawning of the Age of Stochasticity. Appeared in: Mathematics: Frontiers and Perspectives, edited by Arnol'd et al., 2000.
10. Kent Beck et al., Manifesto for Agile Software Development. http://agilemanifesto.org, 2001.
11. Karl Friston, The Free-energy Principle: a Rough Guide to the Brain? Trends Cogn. Sci. 13(7): 293-301, 2009.
12. Herbert Simon, The Sciences of the Artificial. Cambridge: MIT Press, 1970.
13. Gregory Bateson, Steps to an Ecology of Mind. Chandler Publishing Company, 1972.
14. David McCandles, The Beauty of Data Visualization. TED talk at Global Oxford. http://goo.gl/7MzQ, Based on work by Tor Nørretranders. 2010.
15. Marvin Minsky, Society of Mind. Simon & Schuster, 1988.
16. Tim Harford, Adapt: Why Success Always Starts with Failure. Picador, 2011.

Curriculum vitae

Prof. Bert de Vries was appointed part-time professor of Personalization of Medical Signal Processing Systems in the Department of Electrical Engineering at Eindhoven University of Technology (TU/e) on April 1, 2012.

Bert de Vries received his M.Sc. (1986) and PhD (1991) degree in electrical engineering from Eindhoven University of Technology and the University of Florida, respectively. From 1992 until 1999 he worked at Sarnoff Research Center in Princeton (NJ, USA) where he conducted research on a wide range of signal and image processing topics such as speech recognition, financial market prediction and breast cancer detection. Since April 1999 he has been employed both in managerial and scientific roles by hearing aid manufacturer GN ReSound. His research focuses on signal processing solutions for a future generation of hearing instruments and is inspired by the belief that recent progress in machine learning and computational neurosciences will drive the development of a thriving computer hearing field.

Colophon

Production
Communicatie Expertise
Centrum TU/e

Cover photography
Rob Stork, Eindhoven

Design
Grefo Prepress,
Sint-Oedenrode

Print
Drukkerij Snep, Eindhoven

ISBN 978-90-386-3440-1
NUR 959

Digital version:
www.tue.nl/bib/