AdNovum

NOTITIA

Noteworthy news from and about AdNovum

Spring 2018, issue no. 31

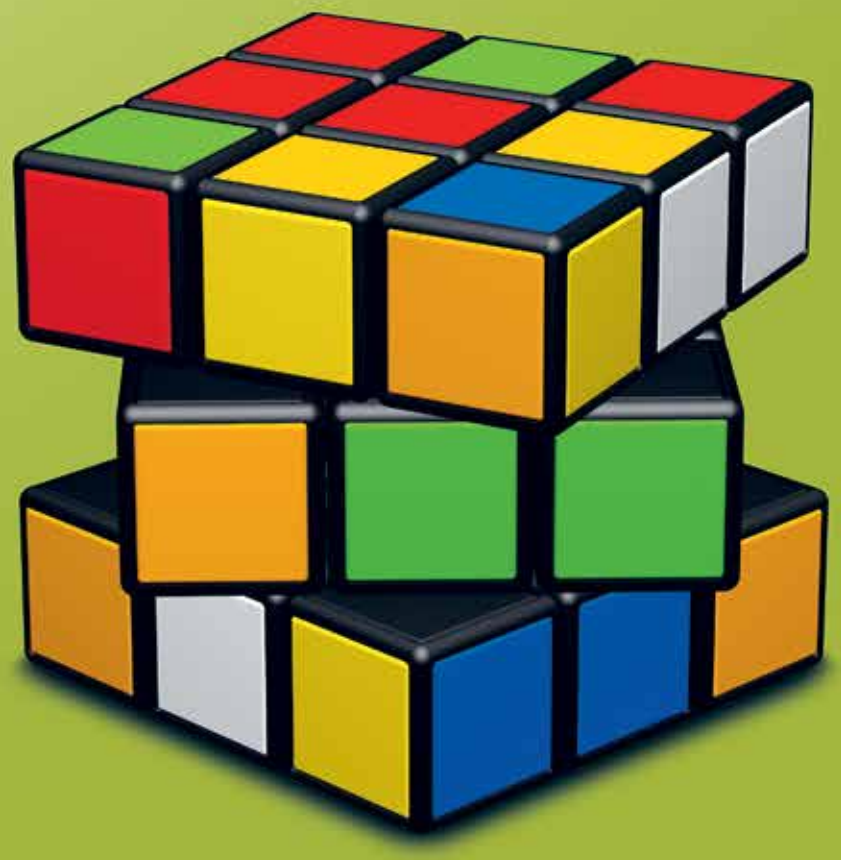

Cognitive Solutions

Content

Dear Readers,

Since the dawn of time, people have been curious about what truly defines intelligence. With the advent of machines, people started to wonder whether machines could actually think. In 1939, the humanoid robot Electro was presented at the World Fair in New York. Controlled using a telephone connection, he could move, count, smoke a cigar and simulate a conversation using a vocabulary of 700 words. The 1950s saw the first projects focusing on machine translation as well as the advent of the General Problem Solver, a software that simulates human thinking. But systems that could genuinely learn and solve complex problems remained a dream.

Between then and now, engineering and technology have made several quantum leaps. Has this progress given artificial intelligence the thrust that many have hoped for, for so long? Maybe. Today systems are highly networked, and are constantly producing and saving data. These enormous amounts of data can be evaluated using specific algorithms. And this is where machine learning along with deep learning and cognitive computing, one of the subsections of artificial intelligence, are already celebrating success. Machine learning recognizes patterns in mounds of data and, perhaps more importantly, recognizes deviations from those patterns.

This can be best demonstrated with a couple of examples: A damages expert working for an insurance company can concentrate on thrilling, complex cases which deviate from the standard pattern while standard cases can be processed automatically. In e-banking, on the other hand, deviations, or so-called anomalies, provide an indication of whether a transaction is being carried out by an authorized user or in fact another person. Or to put it another way: The patterns unmasked within heaps of data can improve security.

These examples show how cognitive solutions bring with them diverse advantages. They increase the security standard or take on routine work which people find increasingly difficult to master, due to both time and manpower constraints. In turn, this new division of labor results in a change of job profiles. New ones are occurring while established ones are disappearing, in exactly the same way as it has always been in business. These changes require both employers and employees to be flexible. Because only those of us who are willing to change will remain efficient and, at the end of the day, competitive.

Talking of change: In the spring of 2017, we carried out a survey among Notitia readers as we naturally want our magazine to keep up with the times. Thank you to all those who took part for your valuable input! Your ideas are gradually being incorporated into the magazine. Enjoy!

Chris Tanner

CEO AdNovum Informatik AG

Cognitive solutions – a short journey through time

Artificial intelligence, machine learning, deep learning, cognitive computing ... these terms are often used synonymously. But what do they mean? Where and how do we use such solutions today? What are the challenges? And where is the great potential that's waiting to be released?

By Matthias Loepfe

The question as to what intelligence actually is has been intriguing people since the dawn of time but there is still no generally recognized definition. Wikipedia defines intelligence as the ability to perceive or infer information and retain it as knowledge. In other words, intelligence is a measure of how well a person can process information using cognitive functions such as perception, memory, thinking and the use of language. If these functions are carried out by a machine, we refer to artificial intelligence. Experts distinguish between strong and weak artificial intelligence. Weakly intelligent systems are very application related. They use machine learning for clearly defined tasks and process information in a way that appears intelligent. Such systems are used, for example, for weather forecasts, recommendation systems, voice processing and anomaly detection. Highly intelligent systems can think like a person and thus, for example, find out new things and make decisions. The question as to whether strong artificial intelligence can ever exist is still the subject of many a heated debate today.

If a machine carries out functions like perception, memory and thinking, we refer to AI.

The beginnings

The history of artificial intelligence goes back to the invention of the computer in the 1930s. In the early days, computers were mostly used to solve problems that could be described using mathematical rules but were, for example, too difficult for people to solve because of the sheer enormity of the calculations involved. But the true challenge is in solving tasks which a person finds easy but which cannot be defined as a mathematical rule, such as understanding language, recognizing faces or walking on uneven terrain.

The true challenge is in solving tasks which cannot be defined as a mathematical rule.

Machine learning

If knowledge is generated automatically from data, the process is referred to as machine learning (ML). A system learns to recognize patterns and regularities from training data. Machine learning is divided into a number of categories: supervised learning, unsupervised learning and reinforcement learning.

In supervised learning, a system is trained with data sets consisting of input and the expected output. The system thus learns to recognize correlations and make forecasts for future inputs.

In unsupervised learning, the exact form of the result is unknown. The system is fed data and from this data independently generates a model and categories according to which it classifies the data. The aim is to understand the available data, in other words reveal hidden structures and groupings.

In reinforcement learning, the system learns how it should react to situations using feedback. This form of machine learning

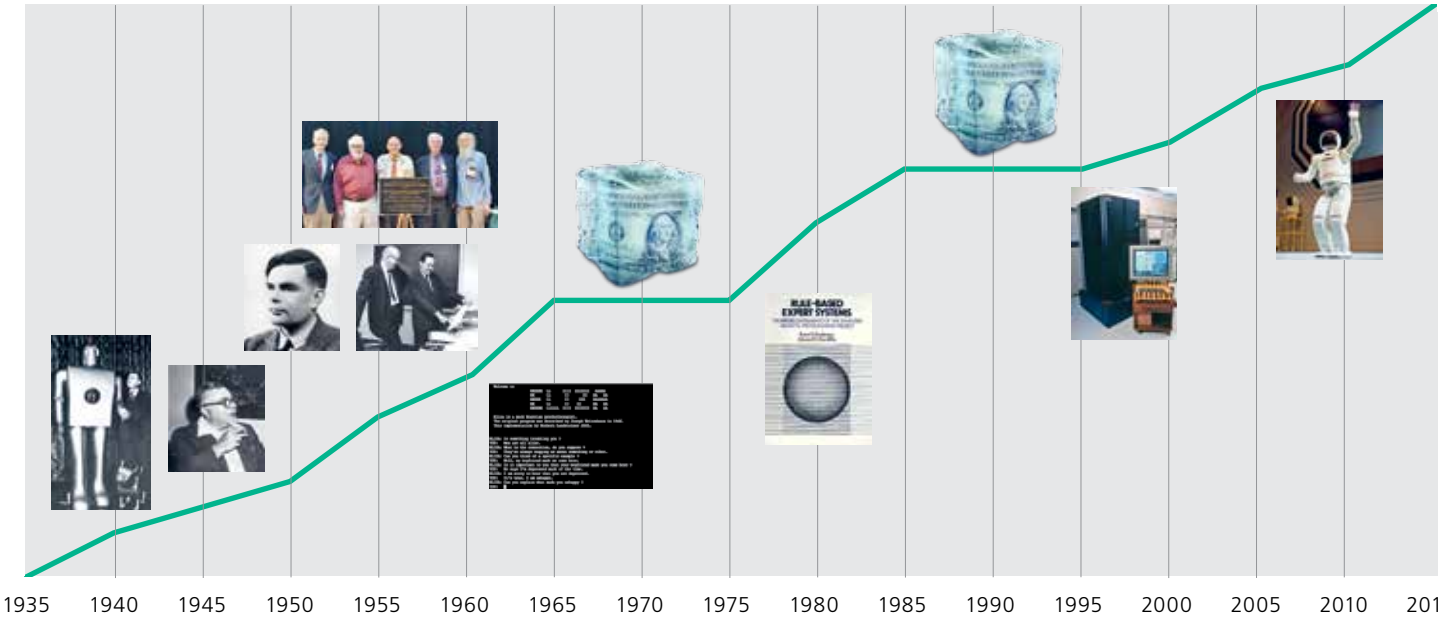

Artificial Intelligence (AI): A topic that has been interesting people since the dawn ot time.

Artificial intelligence (AI)

1939 – Electro: the humanoid robot Electro, also called the "smoking robot", was presented at the World Fair in New York in 1939. Electro was controlled using a telephone connection. He could move, count, smoke a cigar and had a limited vocabulary of 700 words with which he could simulate a simple conversation.

1943 – First neural networks.

1948 – Norbert Wiener defined the term cybernetics (the art of controlling) in "Cybernetics or Control and Communication in the Animal and the Machine".

1950 – Alan Turing aimed to prove that the brain is nothing more than a computer. This was the advent of the Turing test. A test person talks to two interlocutors, one of which is a machine. If the test person cannot distinguish between the human being and the machine, the machine has passed the test.

1950–66 – In the 1950s, the US government put significant funding into projects focusing on machine translation. In the initial stages, sentences were translated word for word before being put together. It was not understood until much later that automatic translation requires extensive knowledge of world events.

1956 – Dartmouth Conference, also known as the big AI bang. The participants came to the conclusion that thinking is also possible without a human brain.

1957/67 – Herbert A. Simon and Allen Newell developed the General Problem Solver (GPS), which was supposed to simulate human thinking. The attempt was abandoned in 1967 but resulted in the development of expert systems.

1965–1975 – First AI winter. Reasons: Recent years were shaped by exaggerated expectations; war in Vietnam.

1966 – ELIZA: Joseph Weizenbaum developed the first well-known chatbot ELIZA that simulated different interlocutors using scripts. ELIZA became well known for simulating a psychotherapist.

1966 – Machine translation: a report compiled for the US Ministry of Defense came to the conclusion that machine translation is not possible. Research in this area was then effectively shelved for almost 20 years.

1970s – Fight about the ontological status of artificial intelligence. As a result of this argument, weak and strong AI are still seen as contrary positions to one another today.

From the middle of the 1970s – Development of expert systems, e.g. for supporting diagnostic and therapy decisions. But in spite of the considerable investments, they did not meet expectations. Because expert systems cannot learn, the entire knowledge has to be programmed, often with a complex set of rules.

1985–95 – Second AI winter. AI started getting competition when research into neural networks were reactivated. But even this second attempt was too early. There was not sufficient training data and solutions for structuring and modularizing the networks and computers were not powerful enough.

1997 – IBM's chess computer Deep Blue beat world champion G. Kasparov. At the same time, major investments were being made into the development of robotics.

Around 2010 – The development of the first robots which could optimize their behavior independently with machine learning.

2011 – Jeopardy! Challenge: IBM's program Watson won against the three top candidates.

2010 to the present day – In around 2010, AI started to be commercialized, particularly in the areas of machine learning and natural language processing. "Deep learning", the rediscovery of neural networks, played a pivotal role here. Deep learning suddenly made it possible to solve problems that for a long time had seemed impossible to solve. It made the breakthrough in pattern recognition with unstructured data which signified the commercial turning point.

comes closest to human learning. It is suitable for sequential decision processes and enables the automation of sequences which are too complex to be programmed.

Over recent years, machine learning has been given additional impetus by the rediscovery of the significance of hierarchical neural networks, referred to as deep learning. Deep learning is particularly suitable for applications in which large data sets are available from which patterns and models can be derived.

Central: defining an aim precisely

Machine learning provides companies with countless possibilities, from the intelligent automation of processes to the development of disruptive business models. There is one question that is central to every project: What is the aim? In a first step, the current situation, working processes and business processes are analyzed and the aim is defined as precisely as possible. Then a strategy is developed to achieve the necessary transformation. It is not until this point that you can tell whether machine learning can be used sensibly or not and for what purpose.

Creating ML skills

If a company has no experience with machine learning, they should start off with a small project. In virtually every environment there are "low-hanging fruits" in the form of simple application patterns which can be identified fast. This gives those involved the opportunity to have their first experience with machine learning. This is how the business representatives can also start understanding just what you can do with machine learning. Once they combine these insights with their specialist knowledge, it is soon no longer a case of optimizing what already exists but developing entirely new business models. This is where machine learning reveals its full power.

When is ML indispensable?

If all the necessary data is available in digital form, using machine learning would certainly be sensible if large quantities of data or data in an unstructured form (such as e-mails, letters, voice, video, chat, SMS) are available, if the data changes constantly or if the data has to be processed using expert knowledge that cannot be described formally or can only be described with considerable effort. Let's take a closer look at the individual cases:

- **Processing large amounts of data**
  If you want to make use of data, you have to understand the data in the first place. In other words, you have to be familiar with the internal correlations and be able to relate the data to other data. A person can quickly reach his/her limits when it comes to large quantities of data or great complexity. And this is when machine learning can help.
- **Automating processes with unstructured data**
  If a company receives 20 inquiries a day, a person can process one after the other. If, on the other hand, the company receives 20,000 inquiries a day, they have to be automated. This is where machine learning can be used, from triage through to the fully automated answering of inquiries.
- **Processing constantly changing data**
  When processing data whose patterns or internal correlations change quickly, it can be difficult to use classic, fully programmed systems of rules because these would have to be adapted constantly. A neural system can be of help in this case.
- **Gathering expert knowledge that cannot be described formally**
  A further area of implementation for machine learning is for tasks in which the experts cannot describe exactly how they solve the task, e.g. because a lot of experience is necessary to carry out the task or because the process is something that the person no longer gives any conscious thought to. If a person is being trained for a new job, he/she usually learns from examples. Machine learning can also be used in such cases. But the machine requires much larger amounts of training data than people do.

IF A COMPANY HAS NO EXPERIENCE WITH MACHINE LEARNING, THEY SHOULD START OFF WITH A SMALL PROJECT.

Where is ML used today?

Machine learning is already used today in a number of areas. Familiar areas of use are:

- Forecasts, predictions: product marketing and maintenance, weather reports, predictive policing
- Voice processing: voice recognition, natural language processing (NLP), language synthesis
- Classification and structuring of data
- Recommendation systems: Amazon, Spotify, Netflix
- Personal assistants: Siri, Amazon Echo, Google Home
- Chatbots: Service Bot Swisscom
- Anomaly detection: discovery of credit card fraud, monitoring, alerting
- Medical diagnostics, epidemiology and biometrics
- Image processing and pattern recognition
- Robotics: movement control, sight, etc.

In spite of the large range of application areas, there are areas which currently cannot be covered by machine learning because they are too complex.

Obstacles when integrating ML

The effort involved in integrating machine learning depends on the quality, the amount and the composition of the available data. Often, data from legacy systems has to be incorporated, access to required information is not possible or difficult, or data cannot be correlated for reasons of data protection. With too great a data volume it could be that the solution is not scaled; with too low a volume, machine learning could well be impossible. Or there is not sufficient expert knowledge for the extraction of the required information from a system.

> TO BE ABLE TO MAKE A COMPUTER THAT THINKS LIKE A PERSON, WE WOULD FIRST HAVE TO UNDERSTAND HOW THE HUMAN BRAIN WORKS, I.E. TO "CRACK THE BRAIN CODE".

Dealing with errors and evaluations

Another obstacle to using machine learning is the lack of predictability or verifiability. Conventional computer programs behave deterministically, in other words, the rules according to which they make their decisions are comprehensible and can be verified. When we teach systems to learn like people, there are certain disadvantages. Human beings do not work on a digital basis and their behavior is not always predictable. That is also the case with machine learning, particularly when deep learning is used. A program's behavior and the basis for its decisions which occur in the process are generally not totally comprehensible. It is often the case that we do not know how they came up with a certain result.

Another aspect is error handling. Systems using machine learning also make mistakes. Sometimes the learning algorithms use false assumptions (distortion) or react too sensitively to variations in the data (variance). Therefore, the systems have to be designed so they can cope with errors.

Implicit judgments are also tricky. If a model is trained with data containing implicit judgments, for example, that are not gender-neutral or are racist, the model also learns these judgments. The dangerous aspect of this is that judgments in data can under certain circumstances be difficult to detect. The less we understand the data, the more tricky the situation. The systems learn mistakes and because human beings do not recognize the mistakes, they cannot intervene to put things right.

Where does ML reach its limits?

As the article is coming to its close, let's be brave and take a look into the future: Have we now finally made the great breakthrough in our dealings with artificial intelligence and will it now increase exponentially, only held back by the performance of the computer? Or have we indeed taken a large step but have now once more reached a plateau?

To be able to make a computer that thinks like a person, we would first have to understand how the human brain works. We would have to "crack the brain code" as Pascal Kaufmann, CEO of Starmind, puts it. Jeff Hawkings is working in a similar direction with the company Numenta and what is referred to as Hierarchical Temporal Memory (HTM). This is a form of hierarchical neural networks which focus on the temporal aspect and try to reproduce fundamental mechanisms of the neocortex.

This is something that is also being pursued by the Human Brain Project (HBP), a major project of the European Commission. The project was initiated by Henry Markram, Professor at the ETH Lausanne, and originally intended to build a computer model of the human brain by 2023. So far they have only managed to reproduce a tiny piece of the cerebral cortex of a rodent. But if the researchers were able to discover key regularities when building the model, the original aim could once again become the focus.

While research on strong artificial intelligence is still in its infancy, enormous progress has without doubt been made in recent years in the area of machine learning with the rediscovery of the significance of neural networks. Today, virtually every individual task, in isolation, can be completed more efficiently by a machine than by a human being. Machine learning also enables the introduction of new business models and is having a disruptive effect in many areas. As our examples show, machine learning is already being used in many areas today. But there is still a lot of untapped potential and a lot to do until we have fully exploited all the technical possibilities we have available to us today.

Matthias Loepfe

Matthias Loepfe, who has a degree in electrical engineering, significantly shaped the development of AdNovum in the early years as Technical Lead and later CTO and co-owner. In 2003 he sold his share and then focused on cyber crime investigation and digital forensics for over ten years. In 2016 he returned to AdNovum and now as Head of AdNovum Incubator and together with his team is researching the suitability for daily use and disruptive strength of innovative technology. Not possible? No such thing. Matthias Loepfe is passionately dedicated to finding the most elegant solution. He proved his skills in craftsmanship when he was younger by rebuilding VW buses. Otherwise he likes being in the great outdoors.

Matthias Loepfe: Explores new technologies.

Machine learning – still lots of potential

Cognitive solutions with machine learning are becoming part of everyday life. They are taking care of routine tasks and support us when it comes to efficiently completing complex tasks. But not only that: They also enable business model innovation. IT consultants Nina Zurbuchen and Zsolt Czinkan talk about the current state of play.

In the 1970s, artificial intelligence was a trend for a while before people forgot all about it. How do you account for the fact that it is back at the forefront of people's minds?

NZ: There are several reasons for that: First of all, machine learning algorithms have been further developed and today we know much more about which algorithm can be used where. And then of course hardware now has more power. This means that even large amounts of data can be processed very quickly. Another factor is the rediscovery of neural networks. People were experimenting with neural networks back in the 1940s but at that time there simply was not enough computing power. Today, neural networks are experiencing their renaissance in the form of deep learning.

ZC: Another aspect is that today data is captured and managed systematically. Everything is networked with everything else, all systems produce data and log files and store data. These massive amounts of data can be evaluated and also used to train algorithms.

> All systems produce data and log files. This data can be used to train algorithms.

How do you train algorithms?

NZ: Data is often spread over various systems and it first has to be standardized before we can use it. Then we analyze the data using algorithms and create a model. How that is done depends on whether you know the expected results.

If you know the results, this is referred to as supervised learning. In this case you use some of the data set as test data. This makes it possible for you to test how reliable the model works. If you don't know the results, the data is analyzed and clustered and the relations between the data are interpreted. The aim is to detect deviating behavior. For example, this is how you can discover that someone is using a stolen identity to transfer money.

What cognitive approaches do your customers use?

ZC: Our customers use cognitive solutions, for example, to optimize the customer advisory process. Algorithms compile information about customers from different sources and then link them to present customer advisors with an integrated overview.

NZ: A lot of companies are just starting to use cognitive solutions. There is still a lot of potential here.

Where do you see the greatest potential for cognitive solutions?

ZC: Wherever people have to work with massive amounts of data. Let's take a look, for example, at the processing of damage claims in an insurance company. Cognitive solutions can single simple cases out and process them to take the pressure off the damage experts so that they can devote their time and attention to more complex tasks.

What other advantages are there of automatically processing large amounts of data?

ZC: For example, it helps detect fraud. If you have large amounts of data and a high volume of damage claims, you can detect certain patterns and recognize cases which deviate from these patterns. These cases are then examined by the damage expert. It is not a question of the machine making a decision for the person but of supporting and accelerating the decision process.

What are the prerequisites for companies to be able to use cognitive solutions?

NZ: On the one hand the availability of data in the required quality, and, on the other, the willingness of the company to use such solutions. Time also plays a role. When you see the result at the end, you are always thrilled. But it takes a certain amount of time to get that far. That is why we suggest to our customers they begin with a proof of concept (PoC) mapping a reduced use case. This helps you to get results fast and provides the basis for you to decide how the project should progress. Very often, a PoC can give rise to interesting "sideline products" which can also be used for other use cases. That is why it is a good idea not to define the project scope until after the PoC.

ZC: Here it is important to understand that there is a big difference between traditional software engineering projects and machine learning projects. A company makes available a certain budget, resources and time for traditional projects. And the project has to be carried out within the confines of this. When it is finished, it will be used for five to ten years, often with minimal adaptations. In our case, the entire process is iterative. We keep looking at the machine learning solution to fine-tune and adapt it. Partly because the data used as the basis for the model can change.

WE KEEP LOOKING AT THE MACHINE LEARNING SOLUTION TO FINE-TUNE AND ADAPT IT.

What hurdles have to be overcome when introducing cognitive solutions?

NZ: For a company to be able to use machine learning, it has to know which data is available. In addition, it is good to know what the purpose of using machine learning is. We check whether this is possible with the existing data and ensure that we can offer the right tools. Sometimes, certain questions can be addressed using methods simpler than machine learning. This is why it is always a good idea to take a look at all possible solutions and then opt for one of them.

ZC: There are actually cases in which machine learning cannot be used because there is not sufficient data or the data that is available is not the right data.

NZ: Another hurdle that has to be faced is the fear employees have of losing their job. Everyone knows that every industrial revolution led to a certain number of professions disappearing. And that is why they say: "If a machine can do my job, they won't need me anymore." We have to explain to people that it is worth their while to get to know the new possibilities. The tools will help them take care of their work more efficiently.

ZC: That's right. Being afraid of losing your job is an interesting point. I feel that some job profiles will disappear and others will crop up. This often brings with it some organizational consequences and you have to have a certain degree of flexibility to deal with these changes. A company has to be prepared to be open to this change process.

What effect does using cognitive solutions have on business?

NZ: In some cases, cognitive solutions can lead to entirely new business models, something which is referred to as business model innovation. A well-known example of this is a tool maker working on an international scale. For example, the company offers their customers service subscriptions instead of drills, as the customer ultimately does not want drills, but holes.

So the customer no longer buys a machine but holes?

NZ: That's right; through a service subscription the customer has access to the entire machine catalog at all times. Machine learning supports the new model in different ways: for example, it allows for flexible pricing, supports resource planning or predicts when a machine will have to be replaced. Thanks to the service subscription, the customer always has exactly the right machine at hand and does not have to maintain an inventory of machines himself.

However, this does not mean that personal service is no longer important today, as the example of a car importer shows. This importer has a large call center as only a small proportion of its customers use the electronic ordering possibilities and around 8,000 garages prefer to order by telephone. Calculations show that the around 40 call center employees take 1,700 calls per day and require around 2.5 minutes for each one.

Would a machine learning solution not pay off in this case too?

NZ: That is exactly what the importer was wondering, whether a machine could take the calls just as quickly as an employee – or whether they should change over completely to a new business model. Now the call center employees don't just sell spare parts during the phone calls but also offer additional services. And this is where the machine learning software comes in: It tells the call center employee how probable it is that a customer will buy an additional service, for example order a liquid. And in fact it is the case that every second caller purchases this kind of product.

ZC: This example clearly shows that it is often a question of generating additional value with existing resources rather than cutting existing resources or jobs.

People say, "Digital insights are the new currency of business." How do you see that?

ZC: Companies with strong competition can only stand out from the crowd by becoming digital and gaining new insight from data, giving them a lead.

Nina Zurbuchen and Zsolt Czinkan: use cognitive solutions in customer projects.

Companies today do actually have a lot of data at their disposal, but can't use the data properly. What can they do in such a case?

ZC: A lot of companies collect data and use it to create what are referred to as data lakes, central databases which store masses of data. Now they want to use this data and generate assets from it. In this case, data science or data mining as it used to be called can help.

Do companies do that themselves or do they need specialists to help them?

ZC: On the one hand they need experts who know all about data science and algorithms. On the other, they need people within the company who know the data – in other words understand both the content and context of the data. Once a company has detected a certain use potential, it requires IT specialists to implement specific measures.

> To stand out from the crowd, companies need to become digital and gain new insight from data.

To what extent do cognitive technologies reinforce the trend toward monopolization?

ZC: If someone creates such considerable added value for the customers that they then change to the supplier's platform and this supplier then grows so much that all others become insignificant, this of course can result in a monopoly situation.

NZ: Although if we take a look at the development of large retail shops, we can see that boutiques are nevertheless still around. Highly specialized online shops are unlikely to die out because they cater to specific customers.

ZC: Amazon is becoming ever stronger in retail business in Germany. A lot of people find this development very worrying. We really have to hope that some strong competitors enter the market.

NZ: Or the market will be regulated by the competition authorities.

Will the company not just change to another country in that case?

NZ: Conceivable, yes. But there are laws to be taken into consideration everywhere. For example, there is quite a lot that Swiss law does not allow. It could be that companies have data that they are not actually allowed to evaluate.

So in other words the laws are an obstacle for your projects ...

NZ: I would not say that. After all laws are there to prevent unjustified and dangerous data evaluation – something we wholeheartedly support. What's more, the risks are not to be underestimated and most companies have very little if any experience with using cognitive solutions. And it is also very difficult to find experts in this field. That is why we recommend starting off with something small and then extending the solution step by step.

> Laws are there to prevent unjustified and dangerous data evaluation.

ZC: Most companies are just starting out with cognitive solutions. They are primarily concerned with collecting the necessary information. The next step is then the evaluation, connection, accumulation, etc. of the data. Of those that already have data, there are many that are sitting on hundreds of data pots that are not connected to each other. Another point is data protection. Data that is collected for a specific purpose cannot just be used for something different, and that is particularly the case if we are talking about personal data. ■

Nina Zurbuchen

Nina Zurbuchen, who has a degree in Business Information Technology and in Business Administration, has been at AdNovum since 2015. As a consultant, she spent several years dealing with strategy, governance and compliance as well as risk and process analysis, before she became involved with the developments in the area of cognitive solutions. Together with her customers, she develops solutions best suited to their needs and if so desired supervises their implementation. Blockchain is another subject she is driving forward. She likes spending her free time with her family, playing tennis or reading.

Zsolt Czinkán

Zsolt Czinkan, MSc, is an information technology engineer and has been working at AdNovum since 2012. As principal consultant he advises key accounts on IT innovation, organization and implementation. Since 2017, he has been responsible for AdNovum's consulting practice for the insurance sector. He enjoys spending his free time most of all with his family in the Swiss mountains or traveling.

Better results thanks to cognitive solutions

Customized cognitive solutions combine data analysis, intelligent search and machine learning. They allow companies to gain information and insights from structured and unstructured data that leads to more efficiency, faster response as well as better customer service and business results.

The increased and automated use of existing data is best done in small steps. In this way, an organization can focus on the most profitable cases and learn continuously.

For companies that want to generate more benefits from their data, AdNovum provides the following services:

- Potential analysis for cognitive solutions (1–2 days)
- Introduction and market overview of cognitive solutions for decision makers (1 day)
- Use case workshops for cognitive solutions (1 days)
- Data analysis: Where are the values hidden in the data? Formulation of hypothesis (usually 1 week)
- Review of hypothesis by means of prototypes (2–8 weeks)
- Implementation of customized cognitive solutions
- Continuous optimization and maintenance of cognitive solutions

Our interviewees Nina Zurbuchen and Zsolt Czinkan will be happy to answer your questions or provide more detailed information.

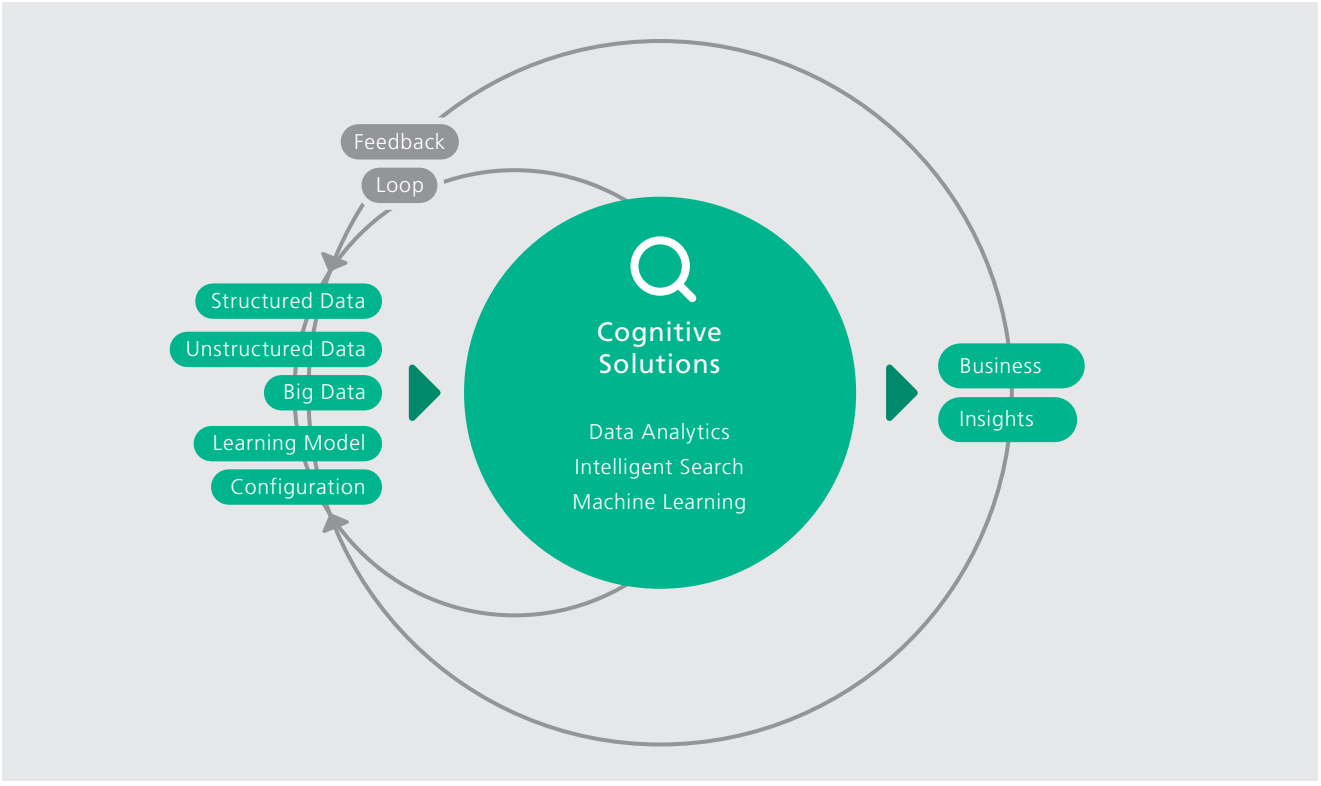

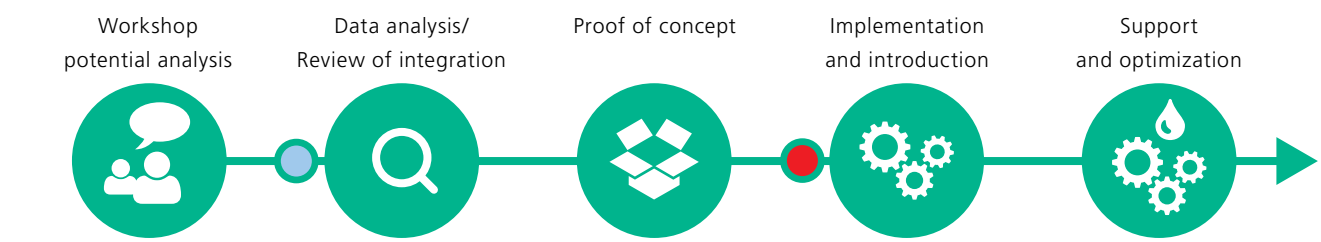

A step-by-step approach to a cognitive solution.

The needle in the haystack

Digital channels make life simpler, but they also make detecting cyberattacks as difficult as finding a needle in a haystack. Fortunately, machine learning optimizes security across various levels, making such a task that much more achievable.

By Hartmut Keil and Aldo Rodenhäuser

Using machine learning (ML) it is possible to find the needle in the haystack – often before it pricks you. In the cyber security environment, the needle represents a nonlegitimate event such as the initiation of a transaction by an assailant. The haystack represents the millions of legitimate events, whether they involve a user accessing a system online or they are generated from internal processes. Machine learning has the potential to redefine the balance between an assailant and a defender.

> Machine learning has the potential to redefine the balance between assailant and defender.

Imagine the following situation: A customer phones his bank and reports a transaction that he supposedly did not execute himself. Said transaction took place on a digital channel.

This is where the work of the digital forensic scientist starts. The first thing the scientist does is to collect all available and relevant information, which is usually a manual and often lengthy process that involves asking of questions such as:

- At what time and from which place did the customer/assailant use the digital channel?
- Which device did the customer/assailant use? Does the communication pattern of the device used correspond to known patterns?
- Does the sequence of queries correspond to a legitimate, established pattern? Is it typical for this user?
- At what time and from which place did the customer issue the queries? Who were the recipients and what were the amounts to date?

In most cases, only some of this information is available. Either because it was not recorded by the systems or because it was recorded but has since been deleted. We are, after all, talking about enormous quantities of data here. And on the application side, it was perhaps not clear during development, which data would be relevant and thus worthy of recording by the system.

The digital forensic scientist's second step is to correlate the data collected from different systems and communication layers with one another and extract a pattern. This takes place manually in most cases.

In a last step, the forensic scientist is now capable of determining whether the transaction in question was executed by the legitimate user or whether he/she is indeed dealing with an attack and, if so, what the nature of the attack is.

The aim of the analysis is essentially to estimate to what degree of probability the transaction corresponds to a previously observed pattern.

It would seem logical to build a system to automate this process. The system should detect in real time whether a transaction is initiated by a legitimate user or not. This would have two advantages: It would reduce the expensive analyses and, at the same time, increase security.

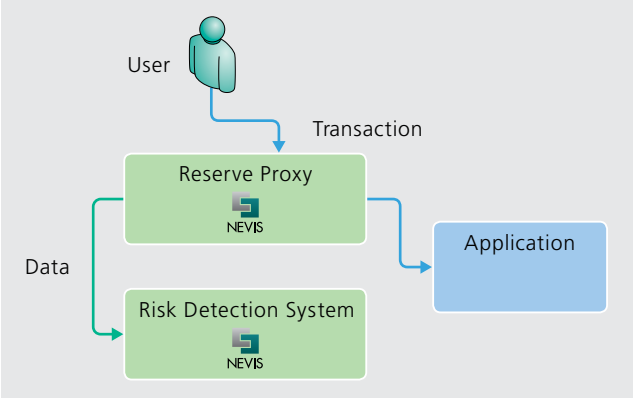

User transaction protected by risk detection system.

Risk detection system replaces manual analysis

The system to be built is called a risk detection system. The system's first task is to collect data. If a reverse proxy is used, it can send all requests and the corresponding data (TCP/IP messages, SSL records, etc.) to the risk detection system.

The surrounding system landscape is not affected and is transparent for the application. The next steps of pattern extraction and the detection of nonlegitimate transactions take place in the risk detection system.

The effectiveness of such a system depends very much on the information available to it. The following examples illustrate this:

- The transaction in question is run by malware on the customer device. This can be detected only because of anomalies in the connection establishment.
- The transaction is executed by a member of the customer's family after the person has authenticated him-/herself with the valid access data. This case can only be detected with biometric information.

The data generally available can roughly be divided into the following categories:

- Device and network information: all information that is influenced by the user's device. In the case of a web application this includes, for example, the browser used, the operating system or the hardware model.
- Context information: the information determined by the physical context such as location data (geo location) and time.
- Biometric information: all information which is determined solely by the user and/or his/her behavior, for example the way the user uses the computer keyboard or mouse.
- Applicational information: transaction data and patterns.

Aldo Rodenhäuser and Hartmut Keil: Their main focus is on security.

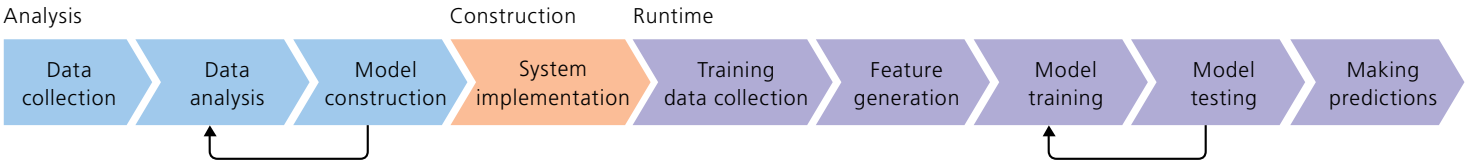

The individual steps of risk detection.

Defining ML aims

We now want to describe the job of our risk detection system in the context of ML in more detail.

First of all it collects data, referred to below as training data. Then, or parallel to that, it generates features from this "raw" data. These features are used as input for the model. Secondly, the model is trained, something we referred to above as pattern extraction. A model usually contains two aspects: certain domain-specific statements which the model has to infer and static assumptions about the data or features which have to be fulfilled.

THE MODEL COLLECTS DATA AND GENERATES FEATURES FROM THE DATA SO THAT THE MODEL CAN BE TRAINED.

This training phase basically consists of optimizing certain features of the model so that it concurs as best as possible with the training data.

The last step is to use the model to decide whether a transaction is legitimate or not. This forecast is referred to as a prediction. At this point we have to define in more detail which prediction our risk detection system or its model should make. Our model should decide whether a transaction is legitimate, in other words whether it corresponds to the pattern of legitimate transactions, or whether a transaction is not legitimate, in other words does not correspond to the pattern of legitimate transactions.

This specification is necessary if we want to phrase the task in ML terminology. A distinction is made between the following tasks in ML:

- The prediction of whether a transaction has the label "legitimate" or «not legitimate» is referred to as the classification.
- The prediction of whether a transaction is similar to the known transactions of the training data or not is referred to as novelty detection.

How can this distinction be explained and how can it be motivated? In classification it is the model›s task to detect whether a fruit is an apple or a pear. The training data has around the same number of apples as it does pears. In novelty detection it is the model's task to detect whether a fruit is an apple or not. The training data contains only apples.

With this background we can now discuss the basic task of our risk detection system or its model: We know that in the haystack, i.e. the training data, there are virtually no needles or that we cannot detect the ones that are there. This is why the model's task has to be novelty detection. What is important now is that our model can deal with this situation. It has to be capable of minimizing the influence of the outliers (needles) in the training data to the prediction (this characteristic is an example of the above-mentioned statistical assumptions of our model).

The features used and the model are the central aspect of our system: Which features are to be generated from the data and which domain-specific and statistical conditions must our model fulfill? These questions should be clarified in an analysis phase, generally using data which has already been collected.

IN EVERY WEB-BASED APPLICATION THE USER'S IP ADDRESS AND THE TIME OF TRANSACTION ARE KNOWN.

Example of context information

Context information is part of the data which is generally available. For example, in every web-based application the IP address of the user and the time of a transaction are known.

The IP address as such provides only very little information. Using special geo location databases, which usually have to be paid for, it is possible to generate features from the IP address, such as country, region, city, geographic degree of latitude and longitude.

When it comes to the time, you have to decide on one of two possibilities: the local time of the sender of the transaction (user) or the local time of the recipient (server on which the application runs). Other features can be generated from the time, such as second, minute, hour, day, month and year.

Which time and geo location features are used depends on the specific case. The following example shows how domain knowledge influences the choice of features and the model: Let's take an application which sales reps of a multinational company use solely for business purposes. The sales reps also

work in places which are not registered completely in the geo location database used. In this case it is obvious to use just the country as geo location feature. For the time features, day and hour of the local time of the user are sensible choices. Furthermore we know that a user always uses the application at a similar time. The selected model must imply this domain knowledge, i.e. the independence of time and geo location features.

What happens now if our model does not imply this assumption? It would possibly learn from the training data that a user in a specific country does not use the application on certain days. This, however, is only a result of the quantity of training data not being sufficiently large.

And what happens if we use other features? Let's take the city as a further geo location feature. We know that the city is not contained in the geo location databases for all IP addresses. For the model this means that it also has to be able to deal with incomplete data – which limits the quantity of suitable models.

Our risk detection system can now use the selected model to decide whether a transaction is legitimate or not. A user's transaction is legitimate if the time and geo location features correspond to the learned pattern. If this is not the case, the transaction is not legitimate.

THE IDEA OF IDENTIFYING A BROWSER UNIQUELY IS NOT NEW.

Example of device and network information

Like context information, device and network information are generally available. The idea of identifying the browser uniquely is not new. Using the JavaScript of new HTML5 functionalities, it is possible to glean more and more information from the browser and thus be able to identify it virtually uniquely. The disadvantage: This approach only works if the user's device is a browser.

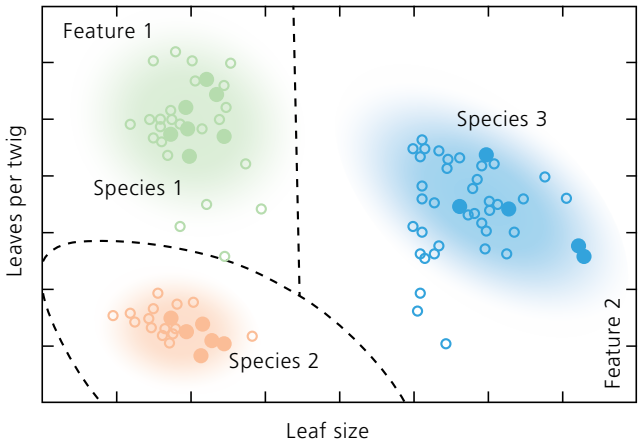

Browser identification via features.

A more general possibility of identifying the browser or the user's device is to learn certain patterns in the structure of the TCP/IP and SSL connection. This approach is based on the fact that parts of communication protocols such as TCP/IP and SSL allow a certain degree of freedom in implementation. This is how it is possible, for example, to distinguish two SSL implementations using the session ID. The specification of SSL requires merely a unique session ID of a maximum length. The appearance of this session ID is a detail of the individual implementation.

IN E-BANKING, THE MODEL LEARNS THE PATTERN OF THE CONNECTION ESTABLISHMENT.

As with the context information, you have to find out which features are to be generated from the TCP/IP and SSL data packages. This requires precise knowledge of the protocols.

But unlike the case of context information, there is, however, virtually no other domain knowledge here to find a model. This situation reveals the power of ML: Over the last decades, very general and flexible models have been constructed. One category of models (support vector machines) has been successfully implemented for the classification of devices.

An unknown data point is now classified by the decision boundary: The label is predicted according to the side of the decision boundary it lies on.

Support vector machine models are suitable not only for classification but also for novelty detection. A further important aspect: They fulfill our requirement to be able to deal with outliers (needles) in the training data (haystack). This type of model is referred to as soft-margin support vector machine.

There are several possible approaches of how our risk detection system can use this identification of the user device. Take a look at the following two examples:

The application is the order system of a stock exchange. "Users" are the systems of the associated banks. In this situation it is sensible for our model to learn the pattern of how each individual system connected establishes the TCP/IP and SSL connection. A transaction is then recognized as not being legitimate when the connection establishment does not correspond to this pattern.

In a bank's e-banking, the model learns the pattern of the TCP/IP and SSL connection establishment of all browsers used. It is not the case of one model per user being trained as in the first example. A transaction that has been initiated by malware, for example, is detected as not legitimate as the connection

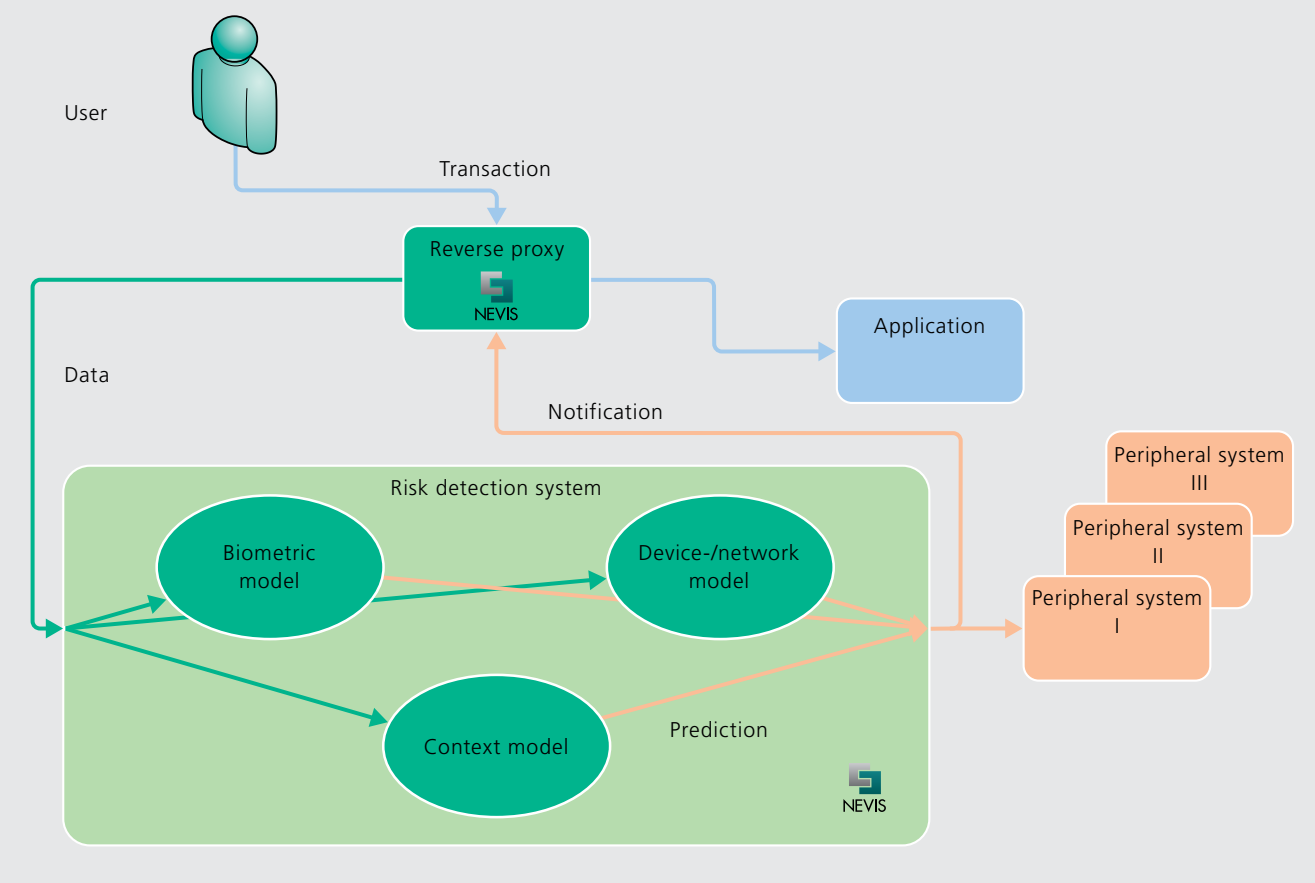

Architecture of a risk detection system.

establishment does not correspond to any of the browsers used. Here, the approach used in the first example would also be possible but the massive number of users would make the quantity of training data to be saved immense.

Putting it all together

Using examples, we have shown how ML can be applied to learn whether a transaction is legitimate or not from context and device/network information. ML approaches have also already been established for biometric and applicational information.

For our risk detection system to fulfill its task perfectly and provide protection for as many types of application as possible, it seems only natural to combine all these approaches. A central risk detection system thus always has an integrative character: It provides information for the different approaches, prevents attacks at the early stage on the reverse proxy and notifies peripheral systems.

This makes full use of all the available information to detect attacks. For the assailant, the complexity and required knowledge thus increase incredibly: He not only has to steal access information from his victim, he also has to look like the victim, behave like the victim and be at the same places at the same times.

Hartmut Keil

Hartmut Keil, MSc in Physics/CAS ETH Visual Computing, joined AdNovum in 2000 as a software engineer. For several years he was involved in the development of middleware products for e-banking and e-government applications. Currently, he focuses on the use of machine learning in the area of security. In his free time, the nondigital world is on the agenda, where he builds or repairs things with his children.

Aldo Rodenhäuser

Aldo Rodenhäuser is Head of Security Consulting at AdNovum and has more than ten years of experience as an IT security analyst and advisor. In his work, he focuses on cyber security, identity and access management as well as mobile security. On a regular basis he provides advice to leading global banks and government institutions in Asia and Europe on their security strategies and architectures as well as their organizational and technological risks. His private program includes expeditions to foreign countries, where he enjoys nature and culinary highlights.

From binary yes/no to continuous authentication

The rules of online identity verification change when you combine a dynamic risk assessment and machine learning solutions. Behavioral biometrics as the missing link to achieve a cognitive security layer.

By Ingo Deutschmann

As humans, we are taught that we are the weakest link when it comes to IT security. Many attempts have been made to remove the human factor from the security equation, but no one has succeeded. If we look at the security we're used to in our devices and services, it is based on thinking from the 1970s, where a binary yes or no at login made more sense. In our always-on culture, that kind of thinking is no longer adequate. Equally, adding extra steps can be a good way to boost security, but also gets in the way of user experience. It is ironic, then, that the human factor, the so-called "weakest link", can be the solution to the security challenge, simply by humans behaving normally.

Big data machine learning biometrics

To find the beginning of the story we need to travel all the way to the Arctic Circle in northern Sweden to Luleå University of Technology. In 2006, an undergraduate behavioral biometrics project, with help from the university's innovation team, resulted in three students founding the spin-off BehavioSec. The idea was uncomplicated while the technology, the algorithm were not. Would normal end user interaction with a device or keyboard be enough to verify the identity of a human being? Are we that unique?

Behavioral biometrics

Human gestures can be repeated in ways that may look similar to the naked eye, however, when they are measured by a behavioral algorithm, they look totally distinct. The way a person holds, swipes, or types on a screen or keyboard is a source of data for user authentication and verification.

Behavioral biometrics technology doesn't measure just one gesture, but a whole range of data inputs, with a high level of accuracy and precision, and can do so throughout a user session. This new capability to be able to continuously authenticate an end user, not just at login, is intriguing to a wide range of organizations, as they see a solution that can protect against account takeover, identity theft, and even internal fraud.

The power of choice

The modern end user of today has high expectations on user friendliness, and they know that they are in a power position to get what they want. Whenever end users are offered a choice they will act with brutal decisiveness: One BehavioSec client operates

About Ingo Deutschmann

A security professional with more than 15 years of experience in development, consulting and product services, Ingo Deutschmann is Business Development Director DACH at BehavioSec. With his former background as General Manager Germany at Gemplus, he will add to the team his knowledge in security software development as well as his experience from the Swedish start-up company Celo Communications and German DEH GmbH, where he was responsible for R&D operations. Ingo was codeveloper of the hardware antivirus solution ExVira. He is a mathematician from the University of Jena, and holds worldwide patents for a smart card reader.

as an identity provider for banks with a combined user base of 7 million. When they started offering strong authentication with a mobile app supported by our behavioral biometrics technology, they saw an exponential growth in usage from 3–4 transactions a month to 20–25.

This highlights the potential for user experience successes, and how the disruption of financial services is already in progress.

Pre analysis

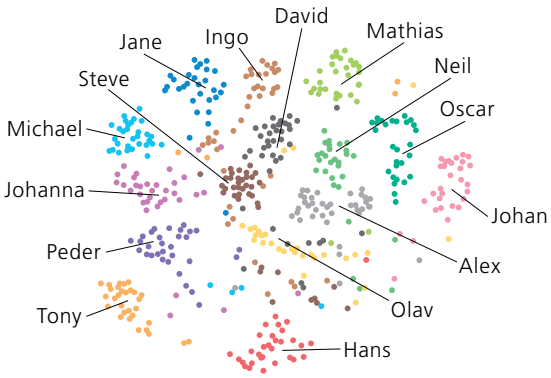

BehavioSec analysis

The BehavioSec scatter plots above before and after BehavioSec algorithm analysis. To the left: clusters of processed end user data from 15 people typing one same password where each dot represents one session. The end user behavior profile cluster is the result of the transformation of 22 dimensions that are simplified into 2 dimensions.

Risk-based authentication

Product, customer and end user experience teams are continuously working to decrease friction, in order to meet the high expectations of busy, multitasking users. Adaptive, dynamic, layered security helps you to create authentication processes that align with these expectations.

Fewer than 30% of us log out of our accounts when we're finished using a service. Our mobile apps are especially vulnerable, even more so now that social media services also act as identity providers and will soon be entering the payments space. Security needs aren't all the same, even within these individual services. For example, checking your bank balance is not as risky as carrying out a large transfer or changing account details.

Get the right level of security, at the right time

BehavioSec analyzes every session from start to finish, continuously profiling behavioral patterns. The system creates a profile match score based on a range of factors, by comparing it to stored results. Is this person typing as they normally do? Are they in a recognized location, using their usual device? This is monitored throughout the session, so that security is an ongoing process, not just a step.

From yes or no to "if-this-then-that"

As a user interacts during a session, the similarity score is fed into your risk engine, and your security or fraud team determines what happens next. If the score is high, the system allows the user through. If it's not high enough, that's when you can add further steps, using the other layers in your system. If the score is very low, your system can log the user out completely.

If it is not you, then who?

BehavioSec has already proven to successfully verify that it is the right person requesting access. The holy grail of fraud prevention is to be able to transform end user behavior to narrow down the group of known users who are the prime suspect for a potential fraud. This is accomplished by efficient machine learning capabilities and applying artificial intelligence to user profiles. Our user profile's level of sophistication enables BehavioSec to find the needle in the haystack. Consequently, more session intelligence will increase the chances for it to be a customer-friendly and secure user journey. By transforming existing user interaction behavior into an additional layer of security you have created a cognitive security solution that will considerably improve the security posture. ■

Imprint

Publisher:
AdNovum Informatik AG
Corporate Communication
Roentgenstrasse 22, 8005 Zurich
Phone +41 44 272 6111
E-mail info@adnovum.ch
Subscription: www.adnovum.ch/notitia
www.adnovum.ch

Responsibility and editing:
Andrea Duttwiler
Feedback: notitia@adnovum.ch

Design and realization:
Comuniq, Zurich

Photography:
akg, Berlin, Gerry Nitsch, Zurich, Fabian Unternährer, Berne

Printed on Balance Pure

Seeking the next big idea?

Let us help – from idea to software.

With AdNovum as your software partner, your business idea will soon be ready for the market. We support you from the initial draft of your idea to the launch of a new product, making you a pioneer and leader in the market. Team up with us and watch your idea become the perfect digital solution. AdNovum, Roentgenstrasse 22, 8005 Zurich, Switzerland, phone +41 44 272 61 11, www.adnovum.ch.

AdNovum