

PROEKSPERT²⁵

2018

HELLO WORLD



ON DATA,
DEVELOPMENT
AND DESIGN

A world without waste,
whether digital, talent,
or time.

PROEKSPERT

We help organizations discover and build digital businesses, products, and services.

Our services:



Digital business discovery



Digital product creation



Data science



Cyber security



25
years



4
offices



160+
people



1140
projects



226
customers

Industries:



Intelligent
Manufacturing



Smart
Environments



Smart
Logistics



Telecom



FinTech

Trends and technologies:

IoT + industrial IoT + intelligent things + connectivity + automation + cloud & migration + cybersecurity
digital services and platforms + augmented reality + artificial intelligence and machine learning
intelligent apps + 3D printing + computer vision + data science + micro services + automated testing
voice recognition + conversational UI + robotics + real time analytics + shop floor optimization
smart manufacturing + predictive maintenance



A New World from a Quarter Century of Experience

This year Proekspert celebrates our 25th year working in industrial automation and information technology. Together with our customers we have created many solutions that make our world better in many ways: safer, more efficient, more comfortable – and even cleaner.

We have a quarter century's worth of experience in smart devices and IT, but can we apply this expertise to new technologies – such as AI? This issue of Hello World is about just that: how we've implemented AI in manufacturing, and how we transfer expertise from IT to AI – plus lots of other war stories from the trenches of IT.

What will the world look like?

Based on our experience in the world of IT and smart devices, we see a new world that will be complemented by intelligent systems. People's everyday lives will become thoroughly mixed and blended with technology. There will be a

fundamental change in how we work, do business and spend our free time – on that we all agree. Some offer a dystopian vision where robots take our jobs. There is also the theory that AI creators are monkeys with grenades – dangerous to everyone.

Proekspert stands for a world where robots and intelligent systems take care of the routine parts of our work (you know, the parts you like the least), allowing us to concentrate on creativity, non-standard decisions, ethics, self-development, and the enjoyment of life.

Our future world depends on companies who create future technologies, their vision and skill, their understanding of (cyber)security, and their ethics, too. That is why we see cybersecurity as a natural part of future technologies, and why we in this issue of Hello World ponder the role of ethics – as well as the ethical code of AI creators – in creating this new world.

Our dream: a world
without digital waste

We see companies fail in creating digital solutions that are useful, simple to use, efficient in delivering value, and of high quality. Quite often the end-user is not even a part of the product creation, the product created based on some “expert opinion” of what the user needs. Some organizations abstain entirely from doing things in new digital ways. This is due to ignorance or the fear of failure. This results in digital waste – the waste of time, money, talent, and potential.

Proekspert's vision is a world without digital waste – the total elimination of inefficient ways of doing things through the creation of new and better human-centric digital solutions. If design is human-centric it's the product of people who concentrated on the why of creation – not the what or the how.

Lack of Technological Cooperation Causes Digital Waste



Terry London | PO of simFactory
terry.london@proekspert.com

Proekspert's mission is to reduce the digital waste. As a manufacturing digitalization and software design company, we see many possibilities for improving the efficiency of production through simple cooperation and informal knowledge sharing.

Digital waste includes anything from simple email spam to huge unstructured data collection that nobody is capable of chewing through. Despite the fact that most of digital waste is generated automatically by computer programs and robots, the main reasons behind it are humans and our habits.

From a management point of view we can attribute digital waste to three major causes: being too formal, being too conservative, and being too secretive.

mode to fix the consequences. And when we are tired of fixing our broken things, the band-aid made out of complex process documentation seems the most tempting thing to cover the problem – just to move on to value-adding work again.

Being overly formal and therefore superficial not only causes digital waste, but in manufacturing it transforms into excess processing, a waste of skills, and the lack of timely information caused by motion and waiting waste.

And even if the focus is partially local, keeping all the know-how proprietary and protecting it as a trade secret isn't such a great idea. Why? Because markets outside of the region tend to look at the overall quality and fame of regional production and do not focus as much on sorting out which company's products are great and which are not. It is more reasonable to work together to build better products and expand your private business contracts through that practice. In some cases several smaller companies could probably unite to gain the capability of fulfilling larger contracts.

A great example is the Raufoss Industrial Park where they honor trust and sharing between people and even declare so publicly as a part of the regional culture: "By collaborating on targeted R&D projects the companies will maintain and strengthen their market positions."

Being overly secretive itself may not cause as much digital waste but it still transforms directly into the same known wastes of lean manufacturing as introduced in the bureaucracy topic. And that brings us to the next point.

Unwillingness to adapt newer working processes and technologies

Probably there is no company which wouldn't like to use the most efficient processes and technologies. But to gain something you have to change your habits. Sometimes we see the signs that there is too much confidence in manufacturing SMEs – that old working methods are still sustainable for next three to five years, in which time the company will be ready to adopt a modern solution.

Examples of these methods are production process tracking and operational analysis of the manufacturing shop floor. In many cases, traceability data collection is solved with a simple pen and paper, or the data are typed into the ERP manually by an operator. Of course the speed and accuracy of analysis of this kind of data leaves much



From a manufacturing point of view, these address at least four out of eight known wastes of lean manufacturing.

Bureaucracy and formality vs. the solutions-driven approach

One obvious reason for digital waste is the presence of overwhelming formalities where documents are created just because they are ordered, but nobody actually understands their real purpose. This is something everyone has experienced at least once in their work life. This habit results in tons of files nobody can read through and adds no real value to an organization.

Even if the initial reason for introducing excessive formalities might have been caused by a wish to address customers' concerns or process-quality issues, probably at some point we just forgot to deal with the root causes and started filling formalities just because it seemed the most logical and safest way to continue at the moment. It happens sometimes when we're fire-fighting to solve a unexpected problem fast and then move into repair

The desire to look at the technological edge as a competitive advantage

Inventing everything by ourselves and avoiding the sharing of know-how seems to give us a competitive advantage. Nevertheless there has to be a balance: restrict too much what your engineers can discuss with their peers, and you will kill their desire to innovate.

I admit that keeping production processes confidential is reasonable most of the time, especially in highly-competitive domains where it's only matter of minutes before a stolen design is recreated in another factory in another country. But when we look at a smaller manufacturing cluster in a given region, small- and medium-sized enterprises (SMEs) with 50 to 250 workers and up to 50 million euros turnover, there seems no point to focus on keeping technological secrets between neighbouring enterprises. This is especially true when the companies are focused on export markets.

to be desired – we practically lose our chance to learn from the data during the production process and operate in darkness with very low visibility.

Here's what we suggest: Adopt exactly opposite mindset. Do not wait three years but start immediately with digitalisation. Choose a couple simple but often-used routine processes and find a partner to automate those – you'll have a quick win and learn something too.

Part of the unwillingness to abandon old habits comes from being too secretive. Being conservative is not a bad thing, but the myopic habit of avoiding the sharing of technological know-how between engineers eliminates the possibility of learning from others' mistakes and discoveries, innovation included. And that directly affects our ability to recognize and adopt newer and better processes and technologies.

The smaller a company the fewer resources it has. Improving manufacturing through technological processes, new software, and hardware is a relatively investment-heavy task. Add here the fact that validating new software or hardware is time consuming as well. But making a hasty choice may lock a company into an inadequate solution for at least the next five years, if not much more.

As we have experienced as software system development experts, being too conservative causes lots of underused or even unusable software code – the digital waste. Imagine the company's vital software system that has been built in-house by only one developer over many years without involving an external partner or advisor – a relatively typical case. This kind of solution is a big risk for a company, because most of the technical know-how and complexity is most likely uncharted, and the stability of the system depends on only one person in company. This kind of risks can be mitigated by using a reliable long term development partner or by using a solution that is shared and maintained together with other similar companies.

Research and development projects that have odds of absorbing a huge amount of money before maturing to a profitable business.


And, if an increased understanding of efficient competition can allow firms to see some areas of R&D investment as not being of competitive value, or as offering only “redundant” competitive advantage, they might find it profitable to cooperate in more aspects of know-how development.

When companies are cooperating to achieve the performance improvements through adopting new technologies, the possibility of succeeding is higher as more collective experience is involved. Of course cooperation means compromises, and processes must be unified between different parties, but that can be only good when we can lower our expenses on software and technology.


Eliminate waste by cooperating with your fellow “competitors”

In addition to Raufoss there are other good examples of cooperation in northern manufacturing clusters. The Estonian Electronics Industries Association members, for example, who participate together at fairs, organize know-how sharing events and visits between companies to demonstrate process- and technological improvements to fellow members. Certain regional companies even share their manufacturing shop-floor resources to level peak times for larger orders.

As a result of constant development of technologies the amount of unusable data and insufficient IT-solutions grows constantly. It cripples our capability to produce products effectively with minimal waste and most importantly, it wastes the most valuable thing people have – their time.







Would You Trust AI with Your Life?



Peeter Meos
Data Science Lead
peeter.meos@proekspert.com



Katariina Roosipuu
Knowledge Sharing Lead
katariina.roosipuu@proekspert.com

The inevitability of new ethical questions arising from advancements in technology presents many researchers and scientists with quandaries. These range from weighing the effects of medical drug trials on consenting subjects, to the trade-off of free and open internet access for security and safety online.

Artificial intelligence has created an entirely new range of ethical questions that challenge the very fabric of our society, and threaten to undermine or trivialise laws that have stood for hundreds, if not thousands of years.

In technology we trust

The more we entrust key aspects of society and our daily lives to technology, the more we will rely on it to make critical decisions about things ranging from what track to play on the living room audio system, to whom to save in a car crash. These two examples sit awkwardly together, but the most awkward thing here is that the outcome of both these two decisions rely on pretty much the same AI process.

Specialists at Proekspert spend their days (and most likely nights) pondering these moral questions. As a company that champions AI technology we have a duty to examine the moral and ethical minefield that surrounds our work.

Who is responsible for the outcome? The creator programs the outcome, and no more. AI is a decision made from a bank of data, but some decisions and situations will require far too much consideration for any human to be able to program them.

The trolley problem

The favourite AI conundrum is the “trolley problem.” In a crash, do you save the driver? Or the paying customer who bought the ticket? Is it acceptable for a product to decide to kill them both in order to save a child? Ethically, you should save the child, right? And how would the AI even know? There are too many factors at play: speed, intention, demographics, and the culturally-defined values of each individual. Because of this problems like these, the industry is going through a period of introspection. And so it should.

When engineering AI you have no social mechanism: you have only robotic laws. An artificially intelligent machine isn’t a legal entity (yet) and therefore is technically

above the law. So where does the blame lie if a machine kills? With the programmer? With the owner? Or with the person killed? Decision mechanics within the AI world have to process more data to become more accurate, working within a set of rules entered by a well-meaning creator. AI also creates its own rules, gleaned from the environment; an intelligent learned behaviour spawned within a set of carefully derived parameters.

Ultimate responsibility

The truth is the creator is ultimately to blame, even if it’s just by default. AI cannot make ethical decisions unless it is programmed to do so, and even if parameters or ethical constraints are laid down by teams of hundreds, it’s near impossible to predict the kind of decision that will be made in every single situation. The only thing that can be guaranteed is the parameter within which the AI can work, and more importantly the lines it cannot cross. But in a situation like our example above, where in a tragic event someone must die, the burden ultimately falls at the feet of those without the in-the-moment knowledge or live information from the event, so essential in any human split-second decision.

It’s not all doom and gloom, far from it. We see AI as a tool to improve the human experience, not threaten it. The role of AI in the elimination of waste, the streamlining of processes and the handling of simple yet monotonous decisions within an industrial setting for example, is proving invaluable, and more importantly carries little of the moral chaos that more headline grabbing applications do.

Release from monotony

AI frees up time, it allows us to devote more time to creativity, to be more productive and less constrained by routine. By releasing us from monotony, we can unlock our skills. It may be able to automate our work, but instead of putting us out of a job it allows us the freedom to do a better job, making even routine work pleasant as

it takes away the boring part. We become more self-determined, independent and free to look for new solutions, set new standards and assume an optimised role within our work.

In this sense the ethical questions are less dramatic and more attainable, but still essential. As one person is given greater power through the use of AI, they need a heightened ethical code to follow. When people are on the line, their education and personal development become the central ethical question, their value laid bare.

No matter how routine the role may be, such as a production line or repetitive process, there will always be the need for tough decisions, and these can as yet only be taken by a person. It’s when we accept this that the dilemmas surrounding more experimental applications of AI become naked and exposed, their progress impossible to fathom within an ethical framework and relying on technology with emotional limitations that struggles with the simplest of human decisions. The advantages of AI are clear, but its limitations can’t be ignored, and we may need a seismic shift in the way we as humans see the value of our own lives for any significant progress to be made, though it’s this kind of shift that may signal the beginning of the end for us.

On the Transferability of Software Engineering Knowledge



Andrus Kuus | Master of None
andrus.kuus@proekspert.com

Proekspert has 25 years of experience in software development. Creating a team of data scientists in our company we learned the hard way that some communication practices software developers use daily may not be so obvious for machine learning experts.

Recent years have witnessed a new rise in data science packaged as machine learning and artificial intelligence. This time the revolution has finally broken out of academia and research labs to the applied business sector. We see practical ML-based applications bettering traditional algorithmic solutions in their ability to tackle a wide array of problems.

The talent shortage
is here to stay

We should still acknowledge that AI- and ML-based products are in the embryonic stage. Early adopters – tech giants with deep pockets – are hoovering up the talent and pouring lots of money into practical application. The majority of organizations, however, are only taking their first steps in the field.

According to a 2017 survey by Kaggle, close to half of data science practitioners see the “lack of data science talent” to be one of the most important barriers in their work. As experience from the IT sector

shows, this problem is not diminishing. According to a recent MemSQL and O'Reilly Media survey, 88 percent of respondents said their companies already have, or have plans to, implement AI and ML technologies. Soon this lack of talent will become even more serious.

The current situation in data science appears to be very similar to the problems faced by companies during the 1990s and 2000s where a wide array of new problems that could be solved with software created a huge demand for software developers. The world soon understood that the deficit of software engineering talent can be dealt with only when the principles of the software development are improved.

Not an engineering problem


Until the 1990s, software development had been viewed as similar to production. Subsequent decades proved this analogy not to be the best. There are few serious problems with engineering and production in software development – or at least the

technical problems are much easier to solve. The harder problem to tackle, in my opinion, has been communication.


I believe communication in a broad context is the main reason successful teams perform better than less successful ones. I do not mean only communication used to understand the problems to be solved, but rather the communication of knowledge and progress inside a team and the practices that allow you to communicate with your “future self.”

Our experience with building data science teams shows that results-oriented people who are not used to working in teams often require training in learning to communicate. The most important change in mentality when shifting from one-man projects to teamwork is the understanding that most of the work you do impacts others on your team or people who will join your team in the future.

HOME



• BONUS




PERIOD




• POSS •

GUEST



BONUS •


FOULS



SHOT CLOCK



FOULS



Lessons from software development

Here are some practices originating from software development that our ML team has found useful in their everyday work.

Reduce your projects' cognitive load.

Work in software engineering and data science means constant learning. To reduce that load it makes sense to have consistent structure in your projects. This includes both the process and how code, documents, and data are organized. Agree on the formats within your team (or organization) and refrain from changing them too often. We've found Cookiecutter Data Science to be a good source for best practices on work item organization for data science projects.

Make experiment set-up easy

Often setting up the environment for an experiment is complicated. Sure, over time you've installed all imaginable python libraries and the scripts run in your machine without problems. But have you

considered how much time you'll spend to set up your environment when you get new new computer? How much time will you need to spend to get "the new guy" up and running? There are good tools available to log and recreate the environment. Conda or Virtualenv for example. Depending on the technical skills of your team, packaging your experiment environment to a Docker image could be extremely effective.

Script your pipeline

Most experiments have to deal with data acquisition, cleaning, splitting and feature definition. Script these transformations. Manipulating the data manually each time new data becomes available, or in case you want to change the split or features could become very time consuming in a long run. When there are multiple steps in your data preparation process the scripts also help to reduce the risk of human error. It may even be worth pipelining these scripts when the operations take a long time to complete so you can run them unattended. You may use simple shell script, makefile or a dedicated framework (like Airflow) for longer-term projects.

Use a single source of truth for your code

I couldn't imagine a software team nowadays who would not use version control to manage their files. The same should become standard for data science teams. In addition to a convenient way of safekeeping and sharing your code within the team, it allows safe experimentation – multiple branches of development can exist in parallel.

Log and store your results

Versus software engineering, data science includes more experimentation. The paths you take are often dead ends in a maze and it makes sense to map where you've been. The best comments in source code often describe alternative solutions which might seem tempting to take but have proven to be less effective or incorrect. It is best practice to document architectural decisions implemented in software with the reasoning why this particular approach was taken. The same applies to ML.





Create tasks with a clear Definition of Done criterion

What's the goal of your current task? Have you agreed on what it takes it to finish? In some cases, for example in EDA, it may be hard to define what is a task and what are the tangible results you expect. In such cases, logging what's been done and timeboxing activity will help keep you out of an infinite loop. Based on the findings better understanding of the goals may arrive.

Visualize your progress

One cannot overestimate the value of knowing what others in your team are working on. Usually, organizations employing data scientists have working task management systems. Maybe it's an online system; maybe a scrum- or kanban board at your office. Visualizing tasks and progress allows the team to find out if somebody is stuck on a task and needs help.

Communicate your progress and intermediate insights

Agile practices in software development have proven that shortening the communication loop between the team and stakeholders is crucial to building a good product. The same is true for data science. Demo your intermediate results and discuss your insights with domain experts and other stakeholders to validate the paths you're taking. This ensures you're not wasting time doing the wrong thing and increases the likelihood of errors being spotted early.

Years of software development have given us simple yet useful practices that facilitate learning, knowledge sharing, and communication, allowing teams to achieve better results. These practices may be unknown to people with other backgrounds. Building on those practices and putting your scarce data science talent to its best use will help bring to life the awesome power of ML and AI applications much faster.



Adventures of a Data Detective



Tanel Peet | Data Detective
tanel.peet@proekspert.com

How frog sounds, birdsong, and other real-life exploits of a data scientist resulted in a tool called Kogu.

Being a data scientist often feels like being a detective who needs to find out how and why something happened. The difference is that data detective's job can have much bigger impact compared to police detectives: the data detective has the superpower to predict the future.

The coolest job in the world

The combination of challenging problems, innovation and the potential big impact of the work makes it one of the coolest jobs in the world. The awesomeness of being a data detective arouses the Instant Gratification Monkey in the Present Data Detective. These rewards are brought on by the manifestation of creative and innovative work in a form of new discoveries or improvements in model accuracy.

For hobby projects done for yourself, the hunger for quick wins might not be a problem. But if you're working in a team or have a long-term project on your hands, the quick high you get from new discoveries may cause problems for the Future Data Detective. Just like the universe tends to move towards chaos and disorder, so will your data science projects if you're not willing to put in some extra energy.

Frog sounds and problem of reproducibility

The first data science project I did was to help a colleague in classifying frog sounds using Deep Learning. My first detective

work went quite well and resulted in a conference paper. I was very eager to earn the doses of gratification the project had to offer me, which resulted in messy structure and code. It didn't seem to be a problem, as it was quite a short solo project, so I could also remember all the experiments I did and hyper-parameter sets I tried. Present Data Detective was happy.

Fast-forward three months and I'm happily working on a new project with my first scientific publication to my credit. I received an email asking me to share code for the published paper. I found the folder where I kept my code, ran the algorithm, and got results which weren't as good as in the conference paper. It took me a half hour to go through the code and understand what it did. I was angry at the Past Data Detective for not using Version Control System (VCS) and not writing Clean Code.

It turned out that Past Data Detective had introduced changes to the code for testing a couple of new ideas and getting another fix of instant gratification after submitting the paper. As the process of reviewing and publishing papers took several months, I had moved on and forgotten the changes, leaving Present Data Detective in a sticky situation. It took me several days before I could restore the code so that it corresponded to the structure and results described in paper.

Birdsongs and learning from failures

My next adventure was to classify birds by their songs, which was a much harder task. But this time I was more prepared: I wrote cleaner code, my project had better structure, and I used VCS quite often. It took a bit more effort and fighting with the Instant Gratification Monkey, but still I ran into unexpected trouble.

There was a long summer between the end of my internship and the time I started to write my thesis. At the end of summer, I could easily find the code and I was happy how readable it was. However, there was a problem with the way I logged results of my different experiments. I used Google Spreadsheets to hold some values about preprocessing and training parameters, but I also used free-form text comments. The problem was that the experiments weren't connected to the code, and for most of the experiments I couldn't recall what the Past Data Detective's comments meant. I wasn't sure what I'd already tried, which resulted in several days of computing power gone to waste (one training took about eight hours).

Since I had learned from the past failures, I wrote a script which would save the logs and all the necessary files, including the code files used for the experiment, to a new directory. After the training I could then open a log file, get the results and put these in the spreadsheet.





When I went to work for Proekspert, I discovered I was not alone in my problem. Other data scientists were also conflicted between getting fast results and having a manageable project. We decided to start automating the annoying parts of our detective work so we could concentrate on the things that really mattered to us. This drive towards reducing energy and time spent on gaining order in data science projects has resulted in a tool we call Kogu.

Kogu


Kogu--an Estonian word translating as “whole” or “entire” in English – is a tool for managing data science experiments. It helps structure projects, giving each a standard structure, allowing users to quickly find relevant data, reports, source code and figures from a project.

Kogu enables the management of a variety of versions of code and it can be used to link source code with experiments, giving us reproducible experiments. The results of each experiment can be logged together with metadata, like interactive figures, tags and comments. This helps users determine what the Past Data Detective did and it can also be used to check how fellow data detectives are doing at the moment.

On its surface Kogu is just another piece of software, but the philosophy behind it is wider. It introduces best practices and processes to fellow data detectives, so we can work on things what really matters, while keeping projects sustainable.

K x g u

AFTER PROEKSPERT'S BIRD LANGUAGE DECODER WE TOOK A REAL CHALLENGE:
A CAT LANGUAGE DECODER.
NOW I CAN FINALLY UNDERSTAND WHAT MY CAT TRIES TO TELL ME.
... WAS NOT WORTH IT.



Be the first to get Kogu. Type 'kogu invite <email>'
\$ kogu invite

Unknown Unknowns and Industry 4.0



Peeter Meos | Data Science Lead
peeter.meos@proekspert.com

Small and medium-sized companies (SMEs) don't know what they don't know and that's okay. Because nobody else knows either.

Everyone's talking about the hype cycle for generic IoT platforms for the manufacturing business. Everyone seems to be doing it, too, and there is a palpable fear of being left behind. However, after a little investigation, the situation is naturally more nuanced.

SMEs and risk aversion

There are industry giants like Siemens leading the way with strong implementations of their digital strategies. But the true backbone of the manufacturing world, SMEs, are believing the larger narrative yet simultaneously trying to avoid the risks inherent in turning their businesses towards the still partly-uncharted waters of digitalisation.

When you take a deeper look into digitalisation platforms on offer, one collides with the real life practicalities of combining legacy technology with modern platforms and practices. Without exceptions, all real-life solutions in manufacturing require a high degree of customisation that most current platforms have a tendency to ignore. Integration is expensive, deals with legacy tech, and most likely is not a scalable business.

Bridging the gap

The gap between real life implementation and the vision of digitalisation is yet to be bridged efficiently. The gap itself is two dimensional, spanning organisational maturity and technological readiness. In order to address both, we advocate a step-by-step approach: moving from creating incremental data collection to establishing situational awareness to advanced analytics tools.

The reason behind this approach is our belief that experience matters, and that before taking on complicated tasks it's necessary to get the fundamentals right. It is not efficient to deploy advanced predictive models before an organisation has experience in data collection, data quality, and visualisation.

For SMEs with fairly limited resources, the true challenge lies in how to jump start the digitalisation of a traditional organisation with efforts that are low risk, have a shallow learning curve, and provide the possibility to fail fast. Since many SMEs feel they have reached a stable level of efficiency they are risk averse to implementing drastic changes, changes

that have a positive effect that will be difficult to assess beforehand.

Accept the experimental nature

An approach that seems to make the most sense is to accept the experimental nature of this journey and start with quick, easy and cheap proofs of concept that validate a customer's specific value proposition.

It may very well be that the traditional way of organising the business is already optimal. This approach provides the least painful method of either proving or disproving that. Start with a quick business discovery workshop, design proofs of concept and prototypes, adjust them endlessly, and find your way in customising the digital landscape to fit your business. In the end you'll have built a solid support structure that carries the organisation into the future in a robust and flexible fashion. As we love to say in the IT world: Experiment, fail fast, iterate.



AI Comes to HVAC



Tähve Lõpp | Business Developer
tahve.lopp@proekspert.com

Tähve Lõpp, Proekspert Business Developer, shares our experience in creating AI-based HVAC control technology that reduces energy consumption ca 10%.

Proekspert in cooperation with marine industry powerhouse Rheinhold & Mahla has developed an HVAC control policy technology based on artificial intelligence (AI) that reduces energy consumption by around 10% while maintaining optimal customer comfort.

The solution can be integrated to existing HVAC controls with minimal installation costs and can start producing value in a couple of months. The software learns from historical data, with its estimated benefit increasing over time. Maximal gains can be expected after one year when the full weather cycle has been observed.

AI in HVAC – centrally controlled intelligent temperature control

Current HVAC control systems rely on a sensible yet simple policies technology. It is common to partition all possible weather modes into either winter or summer, both having their own specifications. The underlying software has the capability to utilize elaborate programmable settings at setpoints for room and air handling units. However, in practice, the setpoints are changed by the crew in a haphazard manner. A central coordinating agent that could account for the effects of setpoint changes to overall energy consumption or occupant comfort is therefore lacking. This approach can be problematic since in practice no two ships are identical both in terms of equipment and thermodynamics. By definition, any calculation based on an a priori design condition falls off target most of the time. An optimal system would require accurate calculations in a continuum of design conditions, weighted by the probability of the condition occurring.

Streamlined data-driven logic

Control systems in modern vessels rely on intricate sensor networks. It is common for HVAC modules to record temperature, humidity, carbon dioxide and pressure in almost every room. Meanwhile, data on

the operating characteristics of various devices is gathered concurrently. This can include but is not limited to fan speeds, valve positions, damper states, pump speeds, refrigerant pressures and general device load levels. All of these readings are typically gathered in localized programmable logic controllers (PLCs), which are today capable of sophisticated control decisions. A central HVAC server usually concludes the highest control abstraction level, being able to monitor and control the entire vessel's operations. This sets the stage for potentially complex control policies, which account for external parameters such as the weather but also for the interdependencies between the hundreds of devices on board.

One step ahead with accuracy

Machine learning models such as neural networks and gradient boosted tree ensembles can approximate functions across a variety of domains with a precision unmatched by theoretical simulations. There are even mathematical proofs stating that neural networks are universal function approximators, meaning that they can learn any arbitrary function. This means that any of the vastly complex differential equations governing thermodynamics and electricity can be simply learned from the data. A trained model can then be used to make predictions or simulations, a kind of a what-if scenario analysis to find optimal control decisions. In the future, the model can also be applied as an adaptive control system which re-tunes PID controller parameters where conditions differ from the original (e.g. due to wear and tear on the devices), enabling more precise microcontrol and allowing for further optimization gains.

Tried and tested


Proekspert worked on a pilot with two RoPax class ferries. Using R&M's domain knowledge and Proekspert's technical experience in data science, a solution was devised that gathers data remotely from the vessels, uses trained models to simulate

different control decisions and sends the optimal decision to a local PLC, which communicates the command to a device.

Despite there being only a few months of operational data, the prediction models already achieve a great deal of accuracy. The temperature in arbitrary rooms can be predicted with a median error of merely 0.2C. The energy consumption of large devices such as chillers can be clearly linked to weather and internal configurations such as chilled water outlet temperatures. Using these predictive models as inputs, a simple control policy was devised that outperforms the current system by a margin of 10% of HVAC energy consumption. The control decisions (such as room temperature setpoint) are performed autonomously – not requiring human intervention. However, if the need arises, the setpoints can be overridden by the crew.

In addition to seamless optimized control, the data-driven methodology proved useful in detecting HVAC faults and misconfigurations. Non-catastrophic faults such as sensor failures or ineffective setpoints are difficult to detect unless there is direct negative customer feedback. However, from the data these can be immediately identified. By monitoring a key set of specific HVAC metrics (such as chilled water inlet-outlet delta T), a real-time overview of the system's operating characteristics can be established.

Overall, the pilot confirms the hypothesis that in large interconnected systems with many components, artificial intelligence can be used to improve control decisions by making it possible to enhance human-imposed rules with adaptive suggestions that have been learned from thousands of historical datapoints.



Danfoss and Proekspert: 15 Years of Industrial Automation Leadership

This year Proekspert celebrated 15 years of collaboration with Danfoss, the production automation leader based in Denmark.

“We took an equity position in Proekspert in 2003 because of similar values,” Søren E. Nielsen, Vice President Global Project Office for Danfoss Drives, told guests at an event to celebrate the two companies’ 15-year relationship. “I’m very pleased that our first experience with Proekspert led to fantastic, long-term cooperation.”

Danfoss is a global leader in industrial automation, and throughout the relationship Proekspert acquired skills in the same field, beginning by developing embedded software tools for variable frequency converters and evolving

to the point where Proekspert became a development centre for Danfoss.

Proekspert has supported Danfoss by developing production automation software for Kronos beer, so prevalent that every fourth beer in the world is created using this software. Proekspert software also is present in the Danfoss frequency converters developed to ventilate car tunnels in Monaco, ensuring air quality and safety in the event of fire. The Danfoss products made for nuclear reactor also uses Proekspert software to cool the reactors.

Ownership structure has changed over the years, and in 2015 Estonian minority shareholders became the sole owners of Proekspert as a result of a management buyout, yet the partnership with Danfoss remains close.

Though the ownership is now in the hands of Estonians, Søren Nielsen still considers the Proekspert team to be colleagues. “Together we have both the history and the experience. We see Proekspert as part of the Danfoss family.”



Søren E. Nielsen from Danfoss Drives, and Marko Sverdlík from Proekspert.



Rui Martins Costa from Proekspert.

Seamless Payment Solutions for Public Transport Networks



Robert Filic | Data Scientist
robert.filic@proekspert.com

Imagine having a convenient, transparent and secure way of getting around the cities of the world. Free of passes, tickets and coins. No wrangling with price lists and timetables. Well, there's no need to imagine: the technology is already here!

A seamless transport system needs a seamless payment solution – The technology is here, and according to leading figures in the city planning sphere, it needs to roll out fast! We take a look at the benefits of a contactless payment solution for users, municipalities and city infrastructures worldwide.

The commute is an essential part of our daily lives; city planners and transport strategists need to be thinking about how to make it not only easier and more bearable, but enjoyable too. Gilles Vesco, former Vice-Mayor of Lyon and VP of Grand City of Lyon, believes: “The goal is to rebalance the public space and create a city for people, [with] less pollution, less noise and less stress.” And it’s with this idea in mind that cities around the world are looking to the future.

Not just the traditional

When we think of getting around a city, most of us might naturally turn to the traditional transport infrastructure. The more tech savvy amongst us will consider taxi-on-demand services like Uber and Taxify, or apps that exist to facilitate car sharing. The more adventurous will often go under their own steam, either walking, cycling or using modern alternatives like the hoverboard, favoured by younger people.

The truth is however, that often the most convenient way to get around a large city is by using an already established public transport system, that’s why it’s still the most popular option. But is the experience as seamless and convenient as it could be? The simple answer is that it all depends on where you live.

Contactless makes sense in so many ways

Most developed cities in 2017 are developing their transport infrastructures around the ‘integrated multi-modal transportation system’. The foundation of this system being the ability to effortlessly move between key points of a city with minimal inconvenience and maximum

speed. You land at the airport and can take a tram to the centre of the city, then the metro or bus to your home. The challenge for such systems is collecting payment. Many cities have reloadable travel cards, day tickets or good old-fashioned conductors. These solutions all require one or more extra steps of effort, which in turn is perceived as inconvenience and actually more of a challenge for the user themselves.

A simple solution that’s tried and tested

London, England. During the 2012 Olympics Transport for London pioneered a system that many users of the network had wanted for years, but previously the tech hadn’t been up to speed. The system allowed users to pay for journeys with their contactless bank debit cards, and has now been permanently rolled out across

the network and offered as one of TFL’s standard payment options. By using the same card each time commuters can take advantage of daily and weekly caps and offers. To date there have been 1.8 million journeys made using contactless payments, from over 100 countries which in total accounts for 40% of TFL’s total amount of trips made since the implementation of the system.

Enter Proekspert

Together with our partners at Ridango, we’re developing a contactless payment solution that can be integrated into an existing infrastructure, or installed from scratch as a standalone system. Integration and operational costs can be reduced for networks using older technology, with significant savings on wages for the oldest of all technology, the conductors and ticket sellers.





Benefits for both municipal operators and their users alike

No longer will a weary business traveller have to go on the hunt for a ticket office or machine when landing in a new city. They will simply pull out their contactless debit card and touch-in when boarding the train, tram or bus from the airport and throughout their multi-modal journey to their destination. The appropriate fee will be deducted immediately; the most effort needed may be a short registration form on the city's transport app or website, but generally this would not be required*.

Speed is a benefit for both the user and operator, less than 500 milliseconds to make a transaction means fewer queues and a more efficient service. As the secure technology works as an offline data authentication solution there is no need to process transactions online in real time, and there is no requirement for the user to enter a PIN, both key to saving valuable time for all.

From the point of view of the municipalities, costs are reduced and the need for large capital investments to implement the technology can be avoided by integrating into existing fare management smart card

systems. Operational savings come from a reduction of activities, this relates mainly to the issuing of new smart cards as well as the maintenance of their infrastructure. As users appreciate these small yet significant advancements – especially those that make the experience of using public transport more pleasant – operators that integrate this technology into their networks will see their reputation benefit exponentially.

Focus on making people's lives better and you can't go wrong

Cities are essentially collections of people. People working, living, loving and having

fun. Any technology that takes away the barriers, that allows us to get on with just being, this is the tech that's essential to us. They say it's the technology you don't notice that's the most valuable, the tech that just works, and this has never been such a poignant observation. The moment technology puts an extra stage into an everyday routine is the moment it becomes irrelevant – something we here at Proekspert strive to remember.



Idea Validation Done Right Means Product With Fans, Not Just Users

A surprising number of businesses are just one small change away from greatness.



Kaja Kruus
Service Designer
kaja.kruus@proekspert.com



Vaiko Hansson
Mindset Coach
vaiko.hansson@proekspert.com

You're head of innovation at your company. You're proud of the awesome things your teams do. Your projects aren't only innovative, but they produce a dazzling ROI.

But not all of your projects. Sometimes you get mediocre results. Sometimes there's a core element missing from the way you do product development.

It starts with a gap

Josh Lovejoy from Google Design says "If you aren't aligned with a human need, you're just going to build a very powerful system to address a very small—or perhaps nonexistent—problem."

But it seems that it hasn't always been like this at Google. Google started as a tech company, and in 2013 they launched Google Glass. Remember those weird looking glasses that looked like they came from a Sci-Fi film -- a bad Sci-Fi film?

The product was very innovative, seemed like the next big thing, but it wasn't solving any human-centric problems. On the contrary - it raised new problems like security and privacy. What was the gap they were filling?

This happens when you are solely focused on the technical solution, but have forgotten the human need. Google is a company that can afford that kind of publicity and hard lessons. But can you?

Get to know your product's users like they were your best friend

First think about your best friend for a moment. Try to answer the following questions before reading further: what kind of life does she live? What excites her? What worries her? What are her ambitions and deepest worries? Who is she as a person? And now think about your product's user. How well can you answer these same questions about her? In other words, do you really know her?

How do you build products and services that are successful and valued by users? It's done by knowing about people, really knowing about them.





Ok, but how?

The most common and successful way is to interact with people in the field; but you need to get out there, it's not happening in your comfy office space or behind the computer.

There are many companies that talk to their users. They ask how they like the product or service, what's missing and what's overwhelming. They even ask their users for ideas. This is fine, but you are not getting to the core. You're missing out on the real need, and this is where the opportunity is to be found.

Forget about the product when talking to your users. Concentrate on human beings and get to know them. Develop a deeper understanding of their context, their everyday life. Then you won't miss out on valuable information.

How to get to this level of understanding?

There are couple of things that can radically change the result you get from your user interviews.

1. Validate the problem that your product will solve, not the product idea itself.

We had a case where a group of top professionals with deep domain knowledge had a product idea. They relied on market trends and research and also interviewed people. But they were validating their idea, not the problem. The result was that after the product launch, nobody wanted to use it. The product solved no actual problem.

Usually people are polite, and during the interview they probably won't tell you that they'll never use your service because they don't need it. They'll tell you, "Yeah, it sounds interesting. It is something I might use," but they're thinking about how to politely finish the interview and hit the nearest Starbucks.

Your idea is actually just another hypothesis. Write them all down and act as a detective trying to get the most objective information out of people without giving clues to what you're actually talking about. In fact, try to forget about your product yourself as much as possible. Talk about them, their lives. Really try to understand them.

2. When you get new and very valuable information, act on it. Even if it proves you've been wrong before. Don't hide your head in the sand.

Here's a story from the early product development phase. From the first round of interviews, we learned that the users will not use the service if the installation process is difficult and takes time. Still, we pushed the test version out (because we'd set the deadline beforehand), but the alpha users were frustrated and we lost them.

People tend to fall in love in their ideas, so when contradicting information comes from interviews it's easy to ignore. Our

suggestion: Don't. You'll end up with a much better product in the end and the satisfaction and love you feel towards your product will be many times greater.

3. Validate the priority of the problem.

There's another angle that needs to be remembered: priority.

Once we designed a solution to a problem we had successfully validated. It existed and generated extra expenses for our client. But getting a sale dragged on and on. We figured that there must be something that we had missed, so we dug deeper with additional interviews. We discovered that even though the problem was real, it wasn't the top problem in the list. There were many other problems that kept the client up at night. So we went back to the beginning with problem discovery.

You can validate that a problem exists, but is it something that can be categorised as "A nice thing to have, but not worth reaching for my wallet!"? Or is it a "Oh my, I totally need this; just take my money!"-type problem?

You also need to know where in the priority list the problem lies. Imagine there's a list of activities that a person is focused on and able to cover, if your problem doesn't fit into that list you're going to have a hard time making it into a business, unless you enjoy being in the friendzone.

We know the feeling of excitement when you've done interviews, you discover the problem and you're ready to ideate and prototype. And after some time (and investment) you realise that the problem you were solving is not in the priority list! It doesn't even come close!

Awe and ROI

Problem discovery isn't easy. In fact it is probably one of the hardest things in product development. But it is also one of the most important ones. Doing it wrong will most likely result in a product nobody wants. Do it right and you may get a chance to be the object of your colleagues' awe -- and enjoy a healthy ROI to boot.



Cyber Security: Why It Shouldn't Be an Add-on



Proekspert's Cyber Security team

Tech's standard approach – develop first, protect later
– is both wrong and dangerous.

In an era of cyberattacks and ransomware we rely on platforms, frameworks, and cloud-providers to ensure the security of our information. But can we trust these solutions?

There's a saying that a powerful lock on a paper door does not make a cardboard house secure. With this in mind it's high time we review our approach to cybersecurity.

Can one size fit all?

Countless security companies will attempt to sell you a one-size-fits-all solution. They offer security as a service, high-quality penetration testing, or log-analysis for finding outlying incidents and anomalies. These may seem compelling but don't rely on them to work on their own. Because security is about people, behavior, and the processes and standards we put in place. Software can assist, but it can't do the work for you.

This is not to say that security tools, services, or penetration tests are meaningless. This is to say that there are more security measures a company can and should take.

Of course people and companies aren't perfect. Circumstances change quickly in the world today, and it's understandable that no one ever seems to have the time to document, plan, or consider theoretical security aspects if features aren't ready, and the MVP isn't delivered.

MVPs without security

Security and privacy requirements are often left out of the MVP. Only when

design and development are complete do we consider talking to the security team to ask for their final sign off.

But this approach is fundamentally flawed. It is us asking the security team to put their most secure lock on our paper door. Any final-phase changes would delay the project and be extremely costly.

Security by design

Security by design means that security and privacy are planned, never an afterthought in the final stage of development. In practise it means onboarding security analysts as early on in projects as the whiteboard-drawing phase. Security by design means that changes and preventive measures are easy to make, and far more cost effective.

You'll never be able to eliminate all software bugs before a release, but you can catch many errors before the final version goes out and penetration testing begins. Practise threat modelling in the development process, and identify weaknesses and attack vectors in a systematic way.

If your company hasn't yet implemented any kind of a Security Development Lifecycle, we'd recommend you start by experimenting with the Microsoft model, and then adjust it to your company's needs:

To examine other best practices of how to create a custom SDL, we would suggest that you also take a check out.

After implementing a suitable SDL, and integrating security analysis within earlier stages of your design process, there are

still quite a few measures you can take to ensure security. Here are a few:

Train your developers regularly to ensure that the majority have a comprehensive understanding of the most common attacks, defence mechanisms, etc. This may be time consuming, but as long as attacks change, so must protection mechanisms, and it's important your team be on top of the latest trends.

Make sure you have a dedicated security team who can assist in design and architecture, conduct penetration testing, give general advice, and work on research that will help product teams deliver secure software.

Try to implement a plan where you're prepared against attacks, where risks are analyzed and prioritized. If something goes wrong, you can assess potential loss and know what you need to do on the protective side. Create a system where some services can be taken down in case of critical incidents, without causing disruption to business. At the same time, mission-critical systems need to have the best possible security you can implement.

When it comes to software, there are no 100-percent guarantees against hacking. Critical bugs are bound to come around, and they may not even happen in your code, but in library- or service code that you are using. It is simply impossible to mitigate everything, and that's why you need to consider these risks in your threat model, acknowledge them, and prepare for their eventual occurrence – no matter how unlikely.



North Star: AI Lands in Tallinn



Sayan Pathak



Travis Oliphant



Maxim Orlovsky



Pierre Lupi-Chen

North Star AI, a long-awaited artificial intelligence and machine learning conference, brings some of the best minds in those fields to the heart of Estonia. This is the largest machine learning conference in the region, and we at Proekspert are proud to be part of the organising team.

North Star AI is a gathering of hundreds of developers, tech leads, data analysts, technical founders, and others who are passionate about the subject and its future. It's a great opportunity to hear from some of the world's leading influencers.

North Star AI focuses on the exchange of real-life experiences and case studies—from one practitioner to another. Speakers are developers and analysts who'll focus on benefits they can bring to their peers without an explicit commercial interest.

The schedule is packed
with great speakers:

North Star keynote speaker, Sayan Pathak, principle machine learning scientist at the Microsoft AI Research Center in Redmond, Washington, and faculty member at the University of Washington, will be speaking about scalable deep learning at Microsoft.

North Star is proud to present Travis Oliphant, the founder and Director of Anaconda, Inc. Travis is also the creator of SciPy, NumPy and Numba. He'll speak about scaling and optimizing Python on modern hardware infrastructure.

We're also pleased to welcome AI and Blockchain expert Maxim Orlovsky as one of our speakers. He is founder and visionary at the Pandora Boxchain project, a medical doctor turned neuroscientist now holds several posts in the area of technology-oriented knowledge-intensive business. He is also winner of the National Prize in Science and Technologies (Ukraine), along with a number of international scientific Soros Foundation prizes. Maxim will explain how AI benefits from blockchain and game theory with scalable censorship-resistant consensus protocol and give technical examples of his work in the field.

Anaconda CTO Peter Wang will speak about using Python for data science systems design. Peter has software design and development experience across many areas, including 3D graphics, geophysics, financial risk modeling, large data simulation and visualisation, and medical imaging.

Pierre Lupi-Chen, Data Scientist & Engineer at Realeyes will talk about

measuring how emotions drive buying behaviour using computer vision and ML.

Antti Tarvainen from Curious AI will share his long-term experience in semi-supervised learning.

Other speakers include Siim Sikkut, CIO for the government of Estonia; Ahti Heinla, CEO, CTO, and co-founder at Starship Technologies; Jaan Altosaar, Physics PhD student at Princeton; Mattias Nilsson, AI Researcher at FOI and the Swedish Defence Research Agency; Hendrik Luuk, CTO at AlphaBlues; and Peter Prettenhofer, data scientist at DataRobot.

Andrus Kuus, software analyst at Proekspert, will speak on the topic of streamlining data science projects, with a focus on how to unify, integrate and bridge teamwork and developmental processes to beneficial effect.

Everyone working in the field will benefit from the range of opinion and experience present at the conference, as well as the opportunity to meet and network with the speakers. Proekspert supports North Star in order to connect Estonian data science to the global scene, nurturing the local data science and artificial intelligence communities.

NORTH STAR AI
MACHINE INTELLIGENCE CONFERENCE FOR DEVELOPERS

Preparing Kids for AI



Helen Lillep | Customer Experience Lead
helen.lillep@proekspert.com

How Proekspert helps younger generations embrace robotics and make the world ready for AI.

We believe in the importance of understanding one's tools; a painter should know how to wield his brush, and a surgeon must know how to handle her scalpel. In the dawning age of robotics we must understand our smart tools – which are totally unique to our time. And because the true age of AI will belong to future generations, Proekspert considers it a responsibility to promote robotics education.

Proekspert is an active participant and continual backer of Robotex, which recently applied to the Guinness World Records to be listed as the largest robotics competition in the world. With more than 1,300 robots, hundreds of teams, and tens of thousands of participants, the event has had an undeniable impact on robotics education in the Baltics. Our employees are always eager to participate in the competition, helping

Robotex teams make their robots smarter. Proekspert has launched the traditional Beautiful Code competition at Robotex to young developers the opportunity to showcase some of their best and most innovative work.

We also support student robotics teams in competitions and hackathons, offering expert mentorship. We're actively involved with a variety of robotics



clubs across Estonia, and we provide work-shadowing opportunities to the young club members so they may get a first-hand look at how Proekspert develops smart products related to data science, including machine learning and artificial intelligence.

We've found that the benefits of this mentor relationship run both ways: we help kids find their potential in robotics and IT, and the children inspire us with their enthusiasm. Also, employees who

support student teams, mentor young talent, and play around with crazy ideas during hackathons, find that it feeds their own thinking and work in interesting ways. We learn by teaching.

One of our most recent projects was to establish a robotics club for school children on the island of Saaremaa. Proekspert employees run the club, where emphasis is on teaching kids how to build robots and expand their analytical thinking.

The more we work with young minds, the more we are motivated in our own work, and the more we're convinced that the future is bright. With a little effort, the experts of today can pass on their knowledge, preparing a new generation to develop even more innovative tools.

If you're inspired and would like to join forces to help prepare kids for a world with AI, get in touch:

helen.lillep@proekspert.com.



Doctoral candidate Kadri-Liis Kusmin mentors young people interested in robotics.

Proekspert Moments



The next generation
robotics pro-s from
Saaremaa.

Proekspert's jury at Robotex,
the largest robotics
competition in the world.



Tähve Lõpp, Proekspert's
Business Developer,
mentors at the Smart
Industry Hackathon.

Proekspert hosts Spinner
project interns from Narva.
They look inspired, don't they?



EUROPEAN EXCELLENCE AWARDS IN HR



GERMAN-ESTONIAN BUSINESS AWARD



EUROPEAN BUSINESS AWARDS



ENTREPRENEURSHIP AWARD



Join Us!

+372 651 8700
info@proekspert.com
proekspert.com

