CHAPTER 4

Artificial Intelligence (AI)

4) RWE Forecast Combination Model: A model to predict renewables

Dominik Felske, RWE

RWE is a German energy company active in electricity generation, building storage systems and energy trading. After the reallocation of its asset base with former competitor E.ON, the company specializes in generation. Its renewables business is expanding massively and is adapting to meet the new challenges that emerge within the existing, fossil fuel-based business.

This case depicts how incremental rather than disruptive innovations can be integrated into core processes, thereby serving as a role model and enhancing financial performance while maintaining the overall organizational configuration.

Background: We need a model to improve our weather forecasts

I am responsible for the short-term commercialization of the renewables portfolio. We have to nominate and forecast our short-term production of renewables one day before delivery to the power exchange, and every 15 to 30 minutes before delivery. Since solar and wind are intermittent renewable technologies, we depend on actual sun radiation and actual wind. Other than conventional generation technologies such as gas or coal, solar and wind cannot be steered. Therefore, we have to deal with both a certain random factor and structural forecast errors.

What did we do? We looked at various weather forecast providers who offer wind forecast on a granular level for wind farms. In order to minimize errors, they use different weather models and algorithms for different countries and regions. This is their main Unique Selling Proposition. The problem for us was how to identify the best forecast. What we needed was an algorithm helping us

How to cite this book chapter:
Burger, C. and Weinmann, J. 2024. *Leveraging Digital Innovation: Lessons for Implementation.* Pp. 41–70. London: Ubiquity Press. DOI: https://doi.org/10.5334/bda.d. License: CC BY-NC 4.0

to find this out. Since we are in very short-term transactions, we don't have a lot of time to think. Thus we had to include intelligence into the algorithm. In addition, the model would have to run automatically and purely data-driven as there is no one who can decide every 15 minutes what would be the best forecast for more than 100 wind farms.

We decide to apply AI to improve our forecasts

Forecast errors have been a problem since renewable technologies exist. You always try to use the best technology available. But now renewables take bigger commercial risks, margins get tighter, price swings and potential costs increase.

With the emergence of Artificial Intelligence and digitization we decided to find out how we could apply AI tools to forecasts. We hired a Master's student from the university of Duisburg-Essen. In his thesis, he developed a concept how to apply AI to choose the best forecasts by combining different forecasts – a concept we refined with him. After that he started to test his concept and implemented it in a monitoring tool.

Currently we have 5 weather forecasts running in parallel and are able to check the forecast for a wind farm in the course of the next half hours and the next day. Then we do a split by percentages and can say, for example, that the combination of 20 percent of this forecast, 30 percent of this, 40 percent of this, and 10 percent of the other one will be the best forecast and minimize the expected forecast error for a wind farm at this point in time.

If you have a reliable weather forecast you will make money

Forecasts are a cost factor. You have to pay if there is an error and you are not able to deliver the energy you have predicted. These costs are significant and can make a difference in markets with tight margins. So the forecast combination model can be used to further enhance the revenues in the short-term marketing of renewables.

We reduce costs, and if you are doing it correctly and have a more reliable forecast than others, you will profit. If you help to stabilize the system, the system operator will remunerate you for your input.

The main challenge is to reduce forecast errors

Errors differ depending on country and locations. You have forecast errors about 10 to 12 percent, which sounds high, but is actually very low, as they can go up to 40, 50 percent. Of course, you try to reduce forecast errors as a risk mitigation strategy. If you rely on the wrong forecast, you either under-produce or over-produce. In most countries, under productions means that you have to buy back the volumes you have not produced. With more intermittent renewables production, intra-day or imbalanced prices are getting more volatile. In some markets you pay €1,000 per megawatt hour. If you are forced to buy back volumes on that price level, even a small forecast deviation can have an impact regarding profit and loss.

We develop the model internally

Since our model required the handling of huge amounts of real-time data, we collaborated with our IT department. They are responsible to maintain and organize our assets and have all the data available from the various wind farms, but the software, or the tool itself was developed by us. And we started from an abstract concept, which we developed for our use cases.

As trading in real-time data processing is a very sensible topic you have to be confident to use it, since there is always the risk of losing money. Therefore, we did a lot of dry runs and back-testing before we implemented our model for a single wind turbine. In case the algorithm would not perform or we would have problems with data processing, the harm would be limited.

Our learning costs were the opportunity costs of having used external forecast providers instead of our combination model. However, given our step-by step approach these costs were low. We were lucky, from others we know that using the wrong forecast provider can easily result in costs of €100,000 in just one or two days.

We implemented our model with the help of Agile project management. We did several sprints with other departments using internal resources only. We met every 1 to 2 weeks to jointly review the accuracy of the algorithm prototype and to decide on very actionable implementation steps with short feedback loops.

Only in the beginning we had the «external» support of our Master's student who was coached by his professor, whose research area is the connection of the energy market and AI. We met with him several times to receive conceptual input, but we never established a dedicated organizational unit. We brought our student in, who started in my team as an external analyst. By now he is fully integrated in our team. I think you should always incorporate the person who developed a model into your operational processes.

The development costs of the algorithm comprised the man hours of IT, who had to put the functionalities of the model into our automated data flow for the wholesale market, plus those of the analysts in my team to further enhance the algorithm. In total, we have invested low to mid six digits EUR.

Our rate of errors is down to 10 percent

What is AI other than a self-learning algorithm? There is an objective function for the algorithm, and there are certain rules and boundaries. Our algorithm had to learn from certain situations we had in the past, recognize and remember these situations and draw conclusions regarding a new forecast. Currently we use an online machine learning algorithm – if it is active, we use it every half hour.

Let's take Italy as an example. Italy is very hilly and has a long coast. Normally, it is a low wind country. There are forecast providers who are experts for low-wind situations, others for high-wind situations, hilly areas, or sunny regions. Our algorithm combines these factors and understands, okay, we now have a temperature of x degrees, in a low-wind situation for wind-farm x. Let's use the combination of previous forecasts for the next forecast.

So far our algorithm is able – just by a simple combination of previous forecasts – to beat the best forecast providers in the market. At the moment we are improving our algorithm further and expect it to reduce errors to 10 percent which is very good when we talk about forecasts.

We use portfolio theory

When it comes to forecasting for utilities, you are always faced with the question whether you want to do it internally or externally. Most utilities came to the conclusion that in-house forecasting means high fixed costs. External forecasts are cheap. It makes more sense to buy them, find out how best to combine them and benchmark them against each other. There are many weather forecast providers, typically small companies with 20, 50, or 100 people. You could buy one of them or even five, but then you may lose when a different company has the next innovation. We do an annual benchmarking and pick the best forecast providers for us.

One interesting detail: All weather forecasters rely on the same global weather models and are connected. And as we know from portfolio theory, it may be wise to add one source – in our case a forecaster – who on average may not have the best forecast, but is not related to the others. That's why we have a forecaster who relies on a different weather model than the others. This enables us to strengthen our forecast combination. It's just a matter of statistics.

Company-wide we are now a positive example of digitization

By now our model is well-perceived within RWE. It is one of the examples to show that we are becoming more digital. People know that we have a «Forecast Combination Model», but only a few understand what it does. Our internal marketing helped to present our model to our board – as a positive example of digitization – which then helped us to increase our research budget.

From Europe to the U.S.

At the moment, we use our model for Europe, but we are rolling it out for the U.S. as well. However, we are still in our testing phase trying to reduce forecast errors. Next we'll be looking at AI in trading. There the new buzzword is «algo trading.» We need to find out how to improve our revenue at a certain time combining forecasts with pricing. If we can observe certain patterns in the market, we can use them for trading. There are patterns showing that 6 or 12 hours ahead there will be a high-price situation, which would mean a high risk in case we under-produce and thus would be forced to have costly buy-backs. If we see that confidence intervals are widening or discover a certain wind- or solar-feed in in different market zones, the system could be long or short. We could also take price signals into account and include them into the algorithm. But this is on our agenda for the next 12 to 18 months.

Ramping up

We need to become more confident that the combination of forecasts and our learning algorithm will actually perform across different countries and conditions.

Our key challenge right now means to ramp up our model and to know when we are comfortable enough to remove restrictions from our algorithm so that it has a higher degree of freedom to optimize and self-learn which will reduce the number of our interventions. After that we'll see what happened.

Currently we are monitoring our model on a daily basis. If there are high-risk events or events where the forecast did not do its job, we have the option to intervene and switch from the Forecast Combination Model to the simple forecast we used before.

The interview was conducted in April 2020.

Dominik Felske
Head of Commercial Asset Optimization Renewables,
RWE Supply & Trading

Dominik has more than ten years of experience in management consulting and as executive within the energy sector with focus on renewables and commercial topics, and multi-stakeholder management.

Before he took over his current position in June 2022, he was Head of Commercial Optimization CE & APAC and Head of Commercial Analysis at RWE Renewables, working at the forefront of the energy transition by commercializing new and existing renewable energy projects across many geographies. Until 2019, Dominik was heading the Commercial Analysis at E.ON Climate & Renewables and responsible for the Carve-Out of the renewables business unit in the context of E.ON/RWE transaction.

(Continued)

Before joining the renewable business in 2016, Dominik did his MBA at the European School of Management & Technology (ESMT) with a focus on innovation and sustainability in Berlin.

Furthermore, Dominik has worked as management consultant at E.ON Inhouse Consulting focusing on business development and performance improvement and as economist in the energy department at the German cartel office. Dominik holds a Bachelor's & Master's degree in Economics from University of Mannheim & Cologne.

5) Saint Paul Escola de Negócios: Individualizing online learning and making it affordable for all people

Adriano Mussa and Bruna Losada Pereira, Saint Paul

Saint Paul Escola de Negócios is one of the most innovative Brazilian business schools, located in São Paulo, with different branches to enhance learning and generate positive impacts on the market and society.

The business school has the mission to transform the world through the training of ethical, creative, innovative and change agents, by offering pioneering programs for executives that generate a positive impact on society, leading the industry in creating trends and promoting integration between the market and the academic world.

This case demonstrates how an ambidextrous setting can be successfully established if the institutional framework of the organization substantially diverges from the strategic requirements of the digital innovation.

Background: We met Watson and discovered ways to reinvent education in Brazil

We started our AI project in 2016, after I attended an IBM conference where Watson, the IBM supercomputer, was introduced. During the IBM presentation everything seemed easy and practical. We realized that AI - or some fields of AI – could be used in education, especially when it came to revising and reinforcing our students' learning.

We talked to IBM, told them we saw the opportunities Watson offered, but did not yet know how precisely we were to use it. We made it clear that we needed to learn more about the technology and its applicability for our business school. After all, we cannot afford to harm the reputation we are building. We have 15,000 students per year, which is good. However, Brazil is a huge country with a population of 220 million, most of whom are too poor to pay BRL 6,000 for high quality education. We hoped that AI would enable us to reinvent education in Brazil and make it affordable. At that time AI was available as a SaaS (software as a service) model, thus we decided to found the digital platform LIT as a startup. Our School was and is doing well, so we thought, let us be creative, and if we make mistakes, we'll learn from them, just as we'll learn from the mistakes other business schools made when they started their online courses. The main thing was to use the opportunity the market offered.

As soon as we understood what we could do with Watson, we started thinking of our faculty, their knowledge, and their classroom experience. They knew how our students' mind worked and how they learned. What we had to find out was how to take this knowledge and combine it with AI.

We developed «Paul» as a new way of learning

A student can learn reading a book, watching a video class, attending an academic session. We are providing a new way of learning, which starts by talking to our supercomputer named Paul. In the process we have developed, our best professors transmit their knowledge to Paul. Then Paul will transmit this knowledge to our students using our chat bots. So the final feature is a chat bot, but, of course, there is a neural network below the surface.

The best part of this AI journey is not just that our professors teach Paul, and Paul teaches the students. The best part is that when our students are talking to Paul, asking questions, they are teaching Paul their ways of learning, which we transmit back to our faculty. In fact, we train Paul continuously, and Paul is improving every day. Since we have a large student body, we can iterate. We can start to sync up the cycle. This is the AI part below the surface, which is very important for us.

Our expertise and curatorship are assets

Paul does not teach whole classes. We divide what we are teaching to Paul according to our classes. We taught Paul five of our best courses, i.e., Basic Accounting, Financial Analysis, Innovation, Creativity, and Entrepreneurship. These courses had to be split into small parts. We developed a specific algorithm, which we built with IBM. We had a lot of challenges. For instance, I can teach Paul, what is an interest rate? But what kind of interest rate? Are we talking about interest rates that the ministry of economics of a country defines, the interest rate that we use in a valuation process, or the one we use when a credit card company charges us?

By now our AI platform LIT has a number of other learning units as well, such as video classes, case studies, a forum. Paul is an additional option. The advantage of Paul is that you can talk to him anytime, anywhere. Say, you have a meeting this afternoon, and your boss announces that one of the topics will be the company's EBITDA, but unfortunately, you have no idea what EBITDA is. Even worse, you cannot wait for the third lesson of our Financial Analysis course to learn about EBITDA. However, you can talk to Paul and ask: «Paul, what is EBITDA?» Paul will find a professor who knows the answer. They will explain what EBITDA stands for and suggest that you familiarize yourself with depreciation and amortization. The student will see the field – the neural network – and can start their conversation with Paul. One has to know a lot before understanding what EBITDA is, but after the student has acquired some basic knowledge, they can go deeper into the topic and learn, for example, why we use EBITDA in a valuation process or in compensation.

Curatorship is a key aspect of our digital platform. This platform is a core element for all the reasons mentioned above, but, from a student's point of view, you could still wonder, «Well, what's the difference if I ask Paul or go to Google and ask the same question?» Here our curatorship comes in. First of all, using Paul you'll know that he'll give you the correct answers, as it has been curated by trusted faculty members who have a solid academic background. Secondly, you are not just getting an answer and that's it. Instead you can enter a comprehensive network of various concepts that are related to your question. Or, we could say, that you start a learning process – a real thought process – which makes Paul so much more valuable than just entering a question into a search machine and receive answers that are isolated from their larger context.

Another advantage of Paul is that he is much faster than reading a book. Often you have to read an entire book to find the answers you were looking for, and still might discover that it's not quite what you wanted as the learning process when reading a book is not always built in a linear way. Overall we can state that Paul is extremely good for our students.

The difficulty of creating a new habit

By now about 40 percent of our students are using Paul frequently. If a student does it once, it's rare that they'll stop. In fact, people who use it, use it a lot.

For most of our students, AI is a new concept and explaining it to them properly is a challenge. After all, our students are adults who have studied and learned a lot, but never talked to a robot in order to learn. Therefore, we have to create a new habit, which is difficult.

The 40 percent of our students who are talking to Paul to learn, use him mostly when they have an assignment or before a test. Others ask questions that are not related to their coursework. They may be studying the budgeting process but won't ask about the budgeting process. Instead they may go into topics that they inadvertently thought were related to their course.

Students using Paul have a grade average that is two points higher than students who don't use Paul. Maybe it is not only because of Paul. It could very well be that students using Paul are really dedicated to their study. We are dealing with this question to achieve more clarity.

From error and trial to success

At the beginning, our team consisted of Bruna and myself. We knew that we had to start small. One reason was that the world is not synchronized. Another was, that at the beginning, most of our faculty were afraid to transfer their knowledge to a robot, which we understood. However, by now – after 2–3 years – almost 70 percent of our faculty are working with Paul; but in the beginning, it was hard.

Only Bruna and myself were aware of the fact that we would enter a trial and error period and had thus to accept that we would make mistakes. For months we worked very hard, had to throw away our results and start again. We knew that no member of our faculty would be willing to spend time on a project like that. Also we needed total confidentiality. We were developing an algorithm. Not developing from scratch, but customizing an algorithm with IBM. Bruna and I are professionals in business, education, and finance, therefore, we needed the help of IBM professionals.

Eventually, our CIO joined us to integrate what we and IBM had created with LIT using the APIs (application programming interface) needed to connect all parts.

After we had succeeded to map the algorithm – which basically means building a map of knowledge that can be processed by programming and Artificial Intelligence – a large number of our faculty started to be involved and ended up being fully engaged teaching and re-teaching a robot.

Onboarding others went better than we expected

We didn't force anything. Most of our faculty said, «That's a beautiful project. I can create my content in LIT … but as to AI, I need to understand a bit more.» In the end they wanted to be part of the project. By now we have a number of professors waiting to participate and teach Paul new content. Many professors are conservative when it comes to new technologies, therefore we gave them time to understand Paul and answered the questions they had.

A key factor of success was that our academic team is very innovative and young in spirit. This was necessary in order to get the faculty involved. Fortunately, our key people are open to innovative ways of doing things, contrary to a large part of the educational industry. It is not judgmental to say that it is a fact. However, we were convinced that in order to make Paul work, to become early adopters, we could rely on the innovative spirit at our School. In the end, though, feeding knowledge into a robot depends on the involvement of teachers. So after a while, after the initial barriers had been overcome, our faculty came on board.

Students need to learn to ask Paul questions

We think that Paul is a success. IBM used Saint Paul for a global case study in education, as Paul was the first «professor» using IBM Watson. In 2019, we were number 1 in McKinsey's digital maturity ranking for Brazil. This external recognition is very important, as it shows that we are on the right track. However, we are just starting. So far we have taught five courses to Paul, even though we have more than 100 to offer. We are still trying new things and will face more challenges. But we know that we will participate in the future of education.

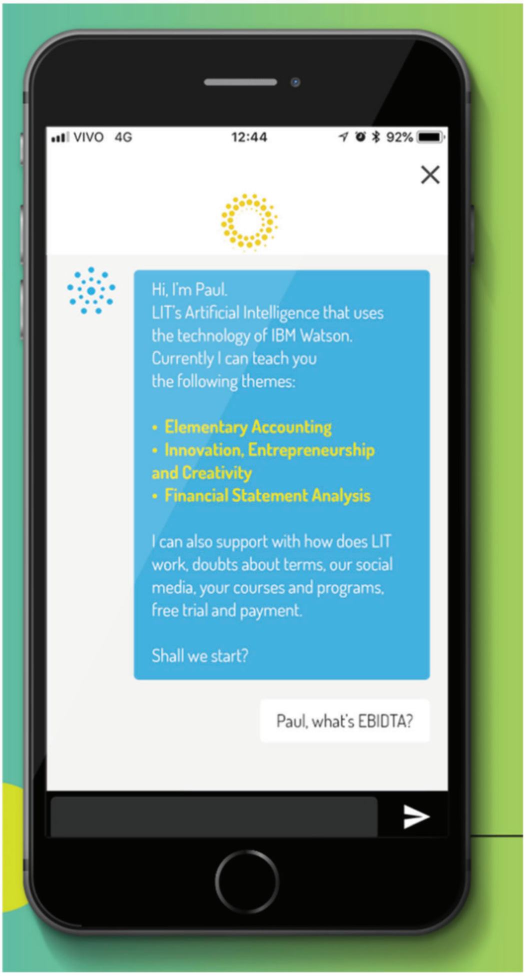

Figure 7: Screenshot of a conversation with AI-based learning tool «Paul».
Source: Losada Pereira and Mussa (2022).

We still need to make more people use LIT. We have to make sure that neither faculty nor students are afraid to use AI. It's not easy to change habits. In addition, we still need funds to scale up the offer. As with any AI project it was scary at the beginning. It was a shock when we analyzed the time sheet for the first course: We had invested 1,500 hours only on our side, IBM probably even more. But we were working on this project every month. Now we can teach a new course to Paul using only 16 hours of a faculty member.

However, given that we offer 100 courses, there is still a lot of work ahead. Our artificial neural networks are becoming more and more complex; when we add new content, we have many overlaps and different points regarding a topic. Take the example of «interest rates» mentioned above.

However, our main challenge is communication, since we have to convince our students of the benefits using Paul. This means to help create new habits, which is a major task.

At the moment, if a student is to use Paul, they have to start a conversation – and ideally be interested in continuing this conversation. This requires the drive to communicate with Paul and to formulate a question. Even if it's not the right question, you have to ask something. Since Paul is not a real person, you don't have to be ashamed or embarrassed of what you ask, still a student may be reluctant and ask themselves «Oh my God, what should I ask?» This is a tiny little detail, but we have to deal with it. The student has to be comfortable when interacting with Paul even when they don't know yet what they don't know.

EOCCS is the first online certification of EFMD. We have become certified, but the members of the group evaluating us, concluded that Brazilian students don't interact enough and don't like to ask questions. They were right. We are waiting for the professor to tell us what they know. It is part of our culture and we have to deal with it. However, we won't be able to solve this problem within a year. It may take a generation to change this cultural attitude.

We split school and platform operation

I am sure that many of our people perceive both LIT and Paul as very innovative, whereas others may still see it as a waste of time. Therefore, we decided to separate LIT and the School. LIT still provides its services to Saint Paul, and vice versa, but people now either work for LIT or Saint Paul.

We needed to do that since the philosophy of LIT and Saint Paul differs. Saint Paul is a traditional business school, whereas a learning platform such as LIT is a disruptive element. It is based on AI which needs a large number of interactions and A/B testing. In addition, it is a subscription service. These concepts cannot be applied to Saint Paul.

Most of our employees supported us when we trained and retrained Paul. They helped us to map different ways of asking the same question. They were our first users and our primary source of testing. And they saw and appreciated the results. Nevertheless, we decided to separate the operations. We could not have the same people working for both, the philosophy and the culture of the two are too different.

Paul will be given a voice

We have started new features, such as using AI to personalize the learning process. For this purpose, our students go through the Big Five personality assessment and we work with three outputs. The first one is the Big Five

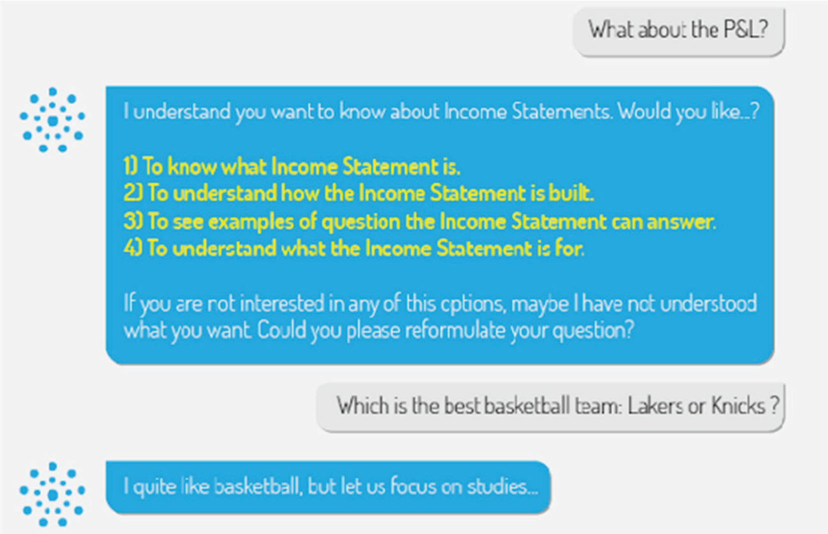

Figure 8: Screenshot of a conversation with AI-based learning tool «Paul».
Source: Losada Pereira and Mussa (2022).

report. For the second output we connect a student's personality traits with our learning objectives. The third output means to select the best learning methods for them, i.e., videos, podcasts, cases study sections, texts, etc.

However, the main advantage of AI is hidden below the surface: The LIT platform is proprietary, this means that every second we receive data students are generating. We are using data analytics to enhance the learning process, finding patterns of behavior, discovering problems. For example, we discovered that on LIT extroverted people learn less than introverted people. So we encourage the extroverts to participate and to join the group discussions. Furthermore, we discovered that if a student does not start a LIT-course during the first 2.5 weeks after subscription, they will probably cancel the subscription. Therefore, we need to help our students when they select a LIT course.

Also, AI presents a fantastic opportunity to support our faculty. Already now they have a high level of understanding and realize how AI can influence and transform our business. AI is part of our MBA program, executive education programs, and BA programs, and we are proud to offer this advanced technology to our students, our alumni and our community.

We certainly will continue testing new AI possibilities. So far Paul communicates in writing and does not yet talk, but we are in an advanced stage giving him a voice. At the moment Watson still has problems dealing with Portuguese. Therefore, we are developing a voice for Paul, giving him the specs IBM applies while enhancing Watson's Portuguese.

We are also starting with a sentiment analysis to help our faculty answer student questions quickly. We have an SLA (service level agreement) to respond within 40–48 hours to questions. For a question such as «What is EBITDA?», 40–48 hours response time is enough. However, when our students complain about something and there is a danger of them becoming disaffected, they need to receive an immediate answer. Therefore, we are testing sentiment analysis with IBM to accelerate our responses in these cases.

Students, faculty, and companies perceive Saint Paul as an innovative business school

Concerning benefit, it was important for us to have the public perceive Saint Paul as an up-to-date, technologically advanced business school. We thus were very pleased when we won the digital maturity premium award from McKinsey & Company in Brazil. It showed us how positively the public regards our use of AI.

It was not easy to shift from being a typical business school to a truly innovative institution, using a technology that is not part of our core business. We certainly benefited from introducing AI to Saint Paul, as we are now perceived as an innovative School by all members of our industry – not only our students, but the companies in Brazil and the Brazilian public as well.

This is the advantage that we discussed with our board members. If you start a project as we did, you traditionally think about budgeting, financial viability, ROI. If we had just focused on these aspects, we would have given up after two or three months. However, we agreed to shift our focus. In the long term, we would like to break even, but in the short and middle term, we are concentrating on the exponential curve regarding our users.

The interview was conducted in February 2020.

Prof. Bruna Losada Pereira, PhD
Deputy Dean of Saint Paul, and one of the team leaders who implemented Artificial Intelligence in Saint Paul's distance learning

Bruna Losada is a consultant at McKinsey & Co, professor, speaker, mentor and passionate about entrepreneurship and finance. She has published the book «Finance for Startups», which was the result of her post-doctorate studies at Columbia University (NY). She lectures at MBA Programs and innovation forums, such as the Consulate General of Brazil in NY, Coworking Programs, Conferences, innovation hubs, among others. In addition, she holds a PhD and a master's degree from FEA-USP, and in the past she held positions of Deputy Dean and COO at Saint Paul Business School.

(Continued)

Prof. Adriano Mussa, PhD
Partner, Dean and Academic & Artificial Intelligence Director
at Saint Paul Business School and LIT

Adriano did his PhD in Business Administration from University of Sao Paulo, followed by postdoctoral studies in Artificial Intelligence from Columbia University, NYC.

For more than 15 years, he has been a professor of finance and AI for Saint Paul's MBAs and Post-MBAs programs. He was a professor at other major Brazilian business schools such as FIA, Insper, HSM and Fundação Dom Cabral, receiving more than 30 awards as honored teacher / best teacher of the programs in which he taught.

He is the author of scientific articles published in periodicals and annals of congresses specialized in Finance, AI and Education, including the book «Inteligência Artificial: mitos e verdades».

6) Lengoo: Automating individualized translations

Christopher Kränzler, Lengoo

Lengoo is a Berlin-based company comprising of engineers and language enthusiasts focusing on professional translations. With their HALOS framework, they custom-train neural machine translation models and pair them with skilled and experienced translators.

Being an early adopter of the GenAI algorithms, Lengoo was a global pioneer of customized adaptation of machine learning models that enable their customers to move from human translators to IT-based tools. In the age of Large Language Models, they retain a competitive advantage via their sophisticated algorithms and ultra-customization of machine translation models.

This case not only portrays the perspective of a startup disrupting an established market, but it also describes a pathway to take a key stakeholder group – the translators – on-board.

Background: How we started

We started the company while we were still at university. Besides studying, I was working for an IT consultancy as a localization manager whose job it was to hire and lead external freelance translators. At that time these processes were handled manually. We took this as the starting point for our company and built a platform to automate the translation process.

At Lengoo customers can upload their documents on our platform. The documents will be automatically analyzed regarding language and word count. The customer selects the target language and receives a fixed price for the translation. They click «order« and we will send the text to a translator. After the translation has been completed, both document and invoice will be sent to the customer.

Initially, though, there was one part that we were not able to automate with rule-based programming, i.e., matching translation and translator. The problem was that in most cases the translator had to be a native speaker and a subject matter expert as well. If you have a legal document to be translated from German into English, you need a native speaker of English who both knows the terminology and the form legal documents require. This level of complexity was too high for a classical rule-based software.

Therefore, we looked into machine learning. Essentially, we analyzed how a human project manager would approach this problem. This meant to analyze both document and content, and find out what type of text it is. If it is an employment contract you match it with all past translations on your platform to find the translator who has done similar translations and received good ratings from the proofreaders or the quality assurers. The more documents a specific expert has translated, the more experience they have, the better you can match them to new jobs. That was the first step how we got into machine learning. We approached it with process automation.

Becoming faster and better than others

From there we moved to machine translating. When you deal with translations and machine learning, it's not far-fetched to look at providers such as Google Translate, Microsoft Translate, and others who deliver generic machine translations. However, these technologies don't work when it comes to professional translations.

We defined the quality of our machine translation-technology based on how fast a professional translator could process a raw translation into a high-quality translation to the extent that an independent professional translator would not know whether it was done by a human or a machine.

We found out that, theoretically, if someone like you and I were to translate a document, we would come up with a speed of 300–350 words per hour. If we were to use Google Translate or any of the other generic machine translation tools, we would reach 550–600 words per hour.

Reality looks different, though. There are efficiency-gaining technologies on the market. For translations this would be a technology called «translation memory», which is widely used. It basically says if you have translated a certain sentence before, the second time it appears in any document you don't have to re-translate, you can just reuse it. Think about manuals, where you have a large number of sentences that are always the same. If your translation tool automatically inserts the previous translation for you, it accelerates the speed of your translation – but you again will come up with approx. 600 words per hour. This meant that generic machine translation solutions didn't yield efficiency gains or improvements for us.

The question was, how to move the technology of machine translation into a setting where the machine translation would be so good that it actually helped the professional translator? That's how we came up with our customization technology for machine translation models. Since we work with our customers directly, we have their past translation data and can use these data to automatically train customer-specific or use case-specific machine translation models by selecting the right data depending on what type of translation needs to be done.

Machine learning models behave the same way as humans. If you have a generalist, who is very good they'll have answers no matter what question you ask. However, in most cases the answer will not be detailed or in-depth enough that it can solve your problem. So you will go to an expert.

The same happens with machine translation models. The models used by Google Translate and others have been trained on very broad data sets including law, finance, mechanical engineering, etc. This enables them to translate whatever you paste into the textbox in real-time.

In contrast, we only give the machine a specific type of data so that it can specialize in this area. This way you reach a quality level that's roughly three times better than the results you get from generic models. This also meant that our translators could work roughly three times faster and deliver 1,500 words per hour, which is a significant increase compared to the 500–600 words of existing technologies.

We automated the entire training process behind our models, because you don't want to have an army of data scientists to do this for every customer, for every department, for every use case.

Our ultra-customization of machine translation models

I wrote the initial machine algorithm as part of a university project in the U.S. Luckily, by now, we have an in-house team that is much better equipped to deal with research topics. But we have always done everything in-house. After our first funding round in 2017, we applied for a European Union funding program to develop our technology, which we called «ultra-customization of machine translation models.» Funding was approved and we received about 1 million euros that we used to build up our research team. Initially, they operated separated from the rest of the company, figuring out how we could get the project up and running. Once we had proven that it worked, we focused on further development. We spoke to other companies that are active in the field of machine learning. We spoke to companies that develop machine learning applications and to companies that have successfully integrated machine learning into their business processes. That was the most important factor: Integrating machine learning into your entire business, not just within your IT team. Embedding machine learning technology in your own IT team is fundamental, but for most of existing machine learning solutions, the human input is still inevitable. In our case, this input comes from the translators with whom we work. Our technology makes them more efficient, but does not aim at replacing them. In fact, this would be impossible as of today.

Putting the end user first is paramount. When you start to develop machine learning applications, you must keep the people in mind that will be using the technology. When you move your idea from research into the real world and put your technology into production, you have to make sure that you include your user in every step of the way. In our case, we assembled a product team with people from our IT team, research team, and a translator. Additionally, we built our own software to interact with machine-translated content, a tool that is designed specifically to post-edit machine-translated output. We are working very closely with our translators and continuously collect feedback in order to make sure that the software really makes their work easier, not harder or more complicated.

We still need humans in the loop

Let's take a translation order worth EUR 1,000 of a traditional translation service provider. Typically about EUR 450 would go to the translator, EUR 100 to the proofreader, and EUR 350 would be used to cover the cost of project management. Project management here includes the people who deal with a client's requirements and communicate with the translator. The remaining EUR 100 are the margin for the translation service provider.

By automating these processes with rule-based applications, including the automation of the translator allocation, we were able to reduce the cost of project management to almost Zero. The translator allocation is a major component of project management, since here human expertise comes into play. The longer you work as a translation service provider, the better you know the translators, find the right translator, thus increase the quality of a translation and make the customer happy.

By turning this part of the process into a machine learning model, we were able to reduce the cost of translations for our customers by about 30 percent. However, this part of the process automation is peripheral.

Far more interesting is the translation part. When we start working with a new customer and the customer is able to provide us with sufficient past translation data, we can save the customer money. Our value proposition is that we deliver the same translation quality as their previous providers but faster and at a lower price. So from a sales perspective, this is working very well.

Let us look at the learning component of machine translation models.

Unlike generic machine translation models, we use a human-in-the-loop approach. Every translation produced by a custom machine translation model is proofread and polished by an expert. Thus our output is constantly being improved. This also means that we have a constant pipeline of highest quality training data available from which the machine translation models can pick. Using this data to permanently retrain our models, the machine translation models become better and better and the translation quality continues to improve. For some of our customers who have been working with us for more than a year, we already have become six times faster than traditional providers and can process up to 3,000 words per hour. So this setup is working and adds real value.

Our advice: never implement machine learning in your core business first

Unlike many other industries, the translation industry has already had a fair share of experience with machine learning technology. The rise of generic machine translation, such as Google Translate, has caught major attention, even though the response was mainly negative. And rightly so: Early versions of Google Translate delivered very low quality. Some translation service providers implemented post-editing flows to poor machine-translated output, which did not help much. Among language experts, all this has led to a low acceptance rate of the technology.

We were in the advantageous situation that machine learning technology had already been in use in our company for peripheral processes – such as the allocation of translation jobs to translators. Hence our translators had been familiar with the technology and open to using it further.

I would advise everyone, who plans to get into machine learning, to never start implementing it in their core process, but in a peripheral process first. Winning acceptance for a technology is a fundamental success factor and it is a lot easier to achieve this if you start with a process that is not essential to the user's everyday performance but ultimately has a positive impact on them.

The algorithm we use in allocating translation jobs to translators helps us to find specialists. It tracks the evaluations of translations performed in specific subject areas and calculates a score that indicates how well a translator fits a specific translation job. This score is the most important figure in identifying the most qualified translator for a given job. We provide our translators with a profile website, showcasing their respective languages and subject areas. This website is their tool for self-marketing, especially since it is based on data, not on mere words. It shows exactly how many words each translator has translated in which field and how their performance has been evaluated by proofreaders. Hence, they receive more jobs in their field and can become still better and faster. And if you are able to translate faster than others, then you earn more. Our translators have all seen the positive impact machine learning can have and are open towards it.

So one might ask, why do our translations cost «only» 50 percent less than those of other providers? We believe that it is absolutely necessary that every stakeholder will benefit. That's why we make sure that in all projects where clients provide us with sufficient language data for the custom-training of our machine translation model, the translator ends up with a higher hourly wage. I think that's how you should apply efficiency gaining technologies. For us, this is working very well.

We didn't do any social media promotion. But we have a community management team that makes sure that everybody understands that we apply our technology not to replace humans but to make them more efficient and increase their income.

Human experts are essential but scarce. In other words, there are not enough skilled human translators on the market to cover the demand. Thus it is a necessity to use machine translation to deal with the demand due to increasing globalization. With fewer people being trained as translators, you need this technology to fill the gap.

The next steps of the machine

The product team, which includes product design, machine learning, and IT, makes up for half of the company. The other half is evenly split between marketing, sales, finance and operations. Operations delivers customer support and liaises with our translators and proofreaders.

The next logical step would be to automate proofreading as well. What does this mean? I think it's best explained looking at the tool we built for post-editing: You have the original text, the translation and the person correcting words or grammatical constructions, etc. to produce a perfect final text. We track all of the actions that it takes to make the machine-translated text perfect. Now we plan to use the data we're accumulating to train models that support proofreading. We can use the data from the correction of the machine translation to show the translator where to go to so that they don't even have to look at the entire document anymore, just at those sentences where the machine isn't sure whether it did a perfect job or not.

Further, we are currently working on making the interaction with the machine better. You know the function that, for example, Google's email service Gmail offers: it predicts what the next word in a sentence will look like while you are typing. This technology is called «predictive typing» or in the case of translations «adaptive and interactive machine translation.» Here we want to make use of the data of all our customers. Currently, we are working in silos as customers do not want to share the information of their translations with others. Still, there is a meta-level, where we can learn how to automatically build a model from one customer's data and use it to build a translation model for a new customer without violating confidentiality. «Transfer learning» is the keyword, and that's where we are going to expand our technology.

Right now we can only apply our technology to large enterprises. But we would like to make it available for small and medium-sized businesses as well. In order to achieve that, we need to transfer learning from larger companies to smaller ones and see if the learnings we have already made with large clients, can also benefit smaller ones.

The question of intellectual property

The data is always owned by the customer. They are paying us for our work and for the right to use the translation. Our customers are also the sole owners of all past translation data.

Concerning the specific machine translation models that we are training based on this data, things get a bit more complicated. Technically it is a machine that trains a machine and we own the machine. But it doesn't matter so much since we work in silos, as I just said. We want to learn based on the translations for a specific customer. We want to train the model on how this specific company expresses itself in certain cases. Which terminology does it use? Which style? Formal or informal way of addressing? These kinds of things are company-specific, even department and use case-specific.

We translate anything you would be able to think of

The possible number of application areas is infinite. If you are producing a technical medical device and are located in Sweden but your target market is East Africa, you may want to translate your specifications into Swahili. If this Swedish company had enough past translations done, we can build a nearly perfect model for Swedish to Swahili.

So much as to the machine translation part of the process. As I said earlier, we don't deliver just raw machine-translations, but include humans. The value we deliver means that our final texts are ready for publication. So we have to consider our available pool of experts. We are currently focusing on Germany, Austria, and Switzerland. The translation demands across all industries in these countries are very similar, the majority being translations from German to English, followed by German to French. Recently, we also have seen an increase in the demand for translations into Arabic and Chinese.

As to subject matters, we translate anything you would be able to think of that falls into the realm of professional translation. Essentially, I would say you can break it down into two types of customers: We have customers that either produce a physical or service product. In the case of physical products, about half of the translations will be marketing texts, the other half technical documentation. Service-oriented businesses demand translations of marketing texts and customer support documentation. About 10 percent of our translations are financial and legal documents.

Our model is our competitive advantage

There are large translation service providers such as Lionbridge, TransPerfect, Acolad, among many others. Ten, maybe even 20 of them occupy most of our market. Then there are smaller translation agencies, who are mostly specified, but would not be able to handle the demand of large companies. We are dealing with the incumbents. However, most of them have not focused on machine translations. In terms of technology, we are one to two years ahead of them.

There is another startup in the U.S. which is similar to us. They built a tool for translators to increase their efficiency and just recently switched to becoming a full-service provider offering translations to end customers.

We could apply for a patent, but our technology is evolving so fast and we are continuously adjusting it and improving our framework that by the time we would get the approval for a patent, it might not be relevant anymore. And even if we wanted to get everything patented, we simply do not have the financial resources to fight for it. So our approach regarding our technology is very simple: We are trying to always be one step ahead of others and are continuously improving our model so that nobody can catch up. This is a very common approach now in the startup scene. The process to obtain a patent is simply too slow to cope with the rapid development of technology, especially in machine learning.

Looking into the future

We just closed our large growth round last September[5]. We have hired a chief sales officer. We are building up our salesforce and will start internationalizing next month. We expect to triple our growth rate this year and will hopefully do the same each coming year. Our board is convinced that the best way to continue would be to go for an IPO.

The large incumbents we mentioned are struggling and won't have the money to buy us. As to companies such as Google and the large tech giants, we don't see them acquiring us either, since we have a human component in our technology. And we will need this component for a long time if not forever to do the final proofreading. This is an overhead these companies don't want to have.

Of course, we have backup plans and interesting exit candidates. IT consultancies and digital consultancies are by now less interested to partner with startups to provide services, but rather build up their own product portfolios. They have realized that technological tools are becoming more self-explanatory. So they start buying solutions. This would be a possibility for us, but our main goal is the IPO.

The interview was conducted in March 2020.

Christopher Kränzler
Co-founder and CEO, Lengoo

As founder and managing director of the AI company Lengoo, Christopher Kränzler has made it his mission to shape the future of enterprise translation. Lengoo has developed an EU-funded machine translation technology based on a highly innovative training approach for Artificial Intelligence. The AI-supported professional translations combine the precision of human creativity with the huge advantages of Artificial Intelligence. Due to the high degree of automation in project management and the translation itself, Lengoo's technology can significantly reduce the costs of professional translations.

Christopher Kränzler holds a Bachelor's degree in Industrial Engineering from the University of Karlsruhe (KIT) and a Master's degree in Data Science from Columbia University New York.

Since June 2019, Christopher Kränzler has been a member of the main board of Bitkom e.V.

[5] *Note from the editor: The interview was conducted in March 2020.*

7) Peregrine Technologies and Allianz: Autonomous driving and the integration of visual data

A dialogue between Stefan Sellschopp, Allianz Partner, and Jorit Schmelzle, Peregrine Technologies

Peregrine Technologies provides new solutions in computer vision and machine learning for a comprehensive understanding of the physical world in real time. The team maximizes the value of existing perception systems, for example in modern vehicles, and adds its own solutions for classification, data fusion, and scene analysis. This real-time visual context allows customers to rethink products and services in the domains of smart infrastructure, insurance, autonomous driving, as well as security. Instead of developing in-house expertise, German insurance company Allianz chose Peregrine to collaborate on autonomous driving.

Under certain circumstances, even an incumbent player in the industry prefers outsourcing the expertise on niche applications of disruptive technologies. This case reveals the benefits of tapping into the larger innovation ecosystem and building on the expertise of external suppliers.

Background: How we started

Stefan Sellschopp, Allianz SE: Allianz is a Dax-listed insurance and service company. Our department is working on B2B2C topics and providing solutions that our business partners will sell to their customers. I am in Automotive, where I work in Connected Cars and Autonomous Driving. There we deal with risk evaluation and an improved first notification of loss. If someone has an accident or the car breaks down, we help the customer settle the claim.

Autonomous Driving is related to Connected Cars, but it has its own specificities. Since 2017, we have been collaborating with various companies worldwide in order to provide them with insurance packages for each autonomous vehicle sold. For example, EasyMile is a French software company that sells bots for small buses that are driving on defined routes. In Germany, the Bavarian town of Bad Birnbach is one of these locations. In 2019, Deutsche Bahn launched a bus service from the train station to a community center. The bus covers two kilometers on a public road. It drives autonomously but still has a safety driver. Should anything go wrong, we handle the damage.

Interesting for us is the following: What kind of technology is in the vehicle? What kind of sensors are there? What is the safety level? What is its purpose? What are its surroundings? What is the traffic situation where the vehicle is deployed? In order to answer these questions, specialists go on-site and check out every curve, traffic light, traffic signal, signs, other objects, and so on. When we have the answers, we evaluate the risks and decide whether we want to insure the vehicle in question or not. And if yes, we determine the premium.

However, autonomous driving is still new, and we are suffering from a lack of experience, which makes it hard to evaluate risk. Therefore, I contacted Peregrine Technologies. The idea was to understand traffic situations and the many ways they can change. In order to find this out, we usually assess the road where the autonomous vehicle drives. But conditions on this road might change over time. After one year or two or three years, traffic flows may be different. Pedestrian walkways or the environment, such as traffic lights, may have been added.

Therefore, we had to be able to automatically detect changes and judge the risks involved. We had to compare both the vehicles and the different situations, thereby finding a solid way of measuring risk. We also wanted to be able to understand whether there is a risky situation before an accident happens. We wanted to measure both the distance and time to collision.

With the help of Peregrine, we are able to solve these problems. We use the data onboard a vehicle, crunch the data, recognize the objects that are in its way, and decide if these are critical. The resulting information is sent to a server.

Which is also to say that now we can properly evaluate the data we receive. We know which insights we can get and which tools and products we need in order to make the data usable. Currently, we are in the process of turning our knowledge into a product, which we will offer to our customers from the autonomous vehicles side.

Peregrine Technologies as supplier

Jorit Schmelzle, Peregrine Technologies: After my studies in physics, I collected industry experience. My last job was at an insurance company in Switzerland. From physics I had knowledge in machine learning, computation, and statistical physics. Eventually, I discovered that the insurance industry has an enormous amount of data that they do not use and that there was room for innovation.

Steffen Heinrich, one of our co-founders, worked for many years for Volkswagen – at first in the software development division for self-driving cars, then on the group level with the Chief Digital Officer on strategic questions and collaborations for autonomous driving. Our third co-founder, Naja von Schmude, studied computer science just as Steffen did. Naja is familiar with robotic systems. Already during their studies, Steffen and Naja were part of a team that built humanoid robots that played soccer and successfully participated in world cups.

Together we realized that we had been working on powerful technologies that were waiting for the revolution of autonomous vehicles. We wanted to make these technologies accessible to everyone so they could benefit from them now, independent of the vehicle type, manufacturer, or model.

So we founded Peregrine and implemented the perception stack of autonomous vehicles in a very simple way. We use off-the-shelf hardware such as smartphones or tablets and deploy our software on these. People can plug in the hardware behind the windshield and turn on our software, which understands what is happening around the vehicle in real time. The software detects the road, the drivable area, and the corridor. It shows the other objects participating in the traffic scenery. Are these cars? Are these pedestrians, trucks, traffic lights? Last but not least, we try to foresee the intentions of other traffic participants and estimate the relation between relevant traffic participants and objects. We measure how our vehicle reacts to changes in the outdoor environment. In sum, this gives us a holistic view of the inherent risk of each situation.

All of the relevant data is collected at the edge, as it would not be feasible to send the whole video stream and analyze it in the cloud. This would be too expensive – in fact, it might not even be possible. The main analytics need to happen at the edge with the AI that we programmed into our perception stack. However, we play back the relevant information – an anomaly, a risky situation, or something else that is of interest – to a server and make it available to our customers on our platform.

Furthermore, customers may use the platform to access their metadata so that they can understand in greater detail what happened in a given situation. They get the location, they can look at the risk score, the velocities, accelerations, forces, and alternative options.

If you have a large number of data sources and fuse them, say, in Berlin, you will see risky and dangerous situations coming up. But the core – or what is really new and important – is that we understand image material. We gather visual context information from a car's environment. There were systems prior to ours that collected or tried to collect data on driver behavior from the dynamics of a vehicle. But just because you perceive a strong braking maneuver, you cannot judge the driver as long as you do not understand why they behaved that way. Maybe a child ran onto the road and the driver braked to save its life. However, if the driver was drunk and did not see the red traffic light or saw it at the very last moment, the braking might look the same, but the cause is quite different. So video traffic analytics is our unique selling proposition.

We support our customers with our platform based on a three-pillar approach. What happens in the car in real time provides the basic data. What we use of it and send to our server can be used for teaching, an improved understanding of risk, as well as for predictions and optimizations.

If we combine all the kilometers driven by our partners and customers, we have by now surrounded the globe approximately 10 times. And every 2 meters, we analyzed the situation and the surroundings of the vehicle. This resulted in more than 16,000 hours of driving and more than 300 million situational analyses.

Insurance companies can use the information we provide to understand what happened during an accident, how severe it was, to send help if it is needed, to give the user the option to interact with the insurer directly, and to protect themselves against fraud.

As to fraud, we are hearing quite frequently that people rent a car and crash it into their private car knowing that the rented car is well-insured. And then they use the money from the insurance to refurbish their car. This is a problem, both for the car-sharing company and the insurer. In these cases, our system can easily help.

But more important is that we can enable a customer to sell usage-based insurance as a value-added service. With the help of our platform, they learn how their customers drive and can incentivize better driving with premium discounts.

Furthermore, we combine our solution with common telematics that try to assess risk in traffic situations but are based on pure dynamic values. We look at how many of these risky situations actually contain elevated risk levels. We can determine that driver X reacts the wrong way in 30 percent of all cases, but does well in 70 percent. We can say, here something happened that was not good, based on the visual context, for example driving through a red traffic light without braking, which would be missed by traditional telematics in 75 percent of cases. In short, we can show why things are happening, not just that something is happening.

Teaching the system

You need to establish a certain ground truth, similar to when you teach your children. Imagine telling your child, «This is a car. This is a truck. This is a traffic light» and showing the real items to them. We, in contrast, collect images and label them «car, truck, traffic light.» Then we let our AI try to assess the image. We feed an image to our algorithm. Our algorithm makes a guess, «This is a car.» We say, «Yes, you're right» and reward this behavior. If it is wrong, we say, «No, this is wrong. This is a car,» and the algorithm gets punished. By doing this millions of times, our system starts to learn.

Leveling the data

How do we level the data? For example, in the case of a near-accident, we would go into automation mode and then to a specific event. We would download everything related to the event – report and metadata – to understand how the car moved. What kind of objects were there? Then we have the video and the changes over time, such as the average number of objects detected in similar situations or on this stretch of the road.

Privacy

How do we deal with the privacy of other traffic participants? Potentially, we could see a lot of vehicles and people that could be identified in the videos that we store. However, we are not allowed to do that. In order to protect their privacy, we anonymize image data. Before we store anything, number plates are anonymized as well as the faces of drivers and pedestrians. Only in the case of an accident are we allowed to store our raw material.

Take person X who must not be recognized. We can anonymize them by blurring or changing features automatically, but still keep valuable information, such as where was this person looking at each point in time?

Can the technology be applied to different environments?

Can it be repurposed, for example, to identify the biodiversity in a forest, such as plants and wildlife? The answer is yes, even though we would have to retrain the system.

We have a variety of use cases for smart cities and traffic infrastructure, to name only some. If you drive in a city and want to know where you can find a free parking space, we can answer that. For this, we just need to tell our system what we are interested in. The same applies to biodiversity in a forest, even though this is beyond our scope, since we focus on traffic and safety.

Collaboration with original equipment manufacturers

We will be happy to share our data with car manufacturers. The logic we use is based on the logic used by original equipment manufacturers. But they have been using this knowledge to better control their cars. What we are doing is attempting to retrospectively understand what happened while a vehicle was driving. In addition, we want to understand what could have happened. However, we are using standard hardware and standard software, that is, basically the same logic used in the embedded device in the car, even if we use it differently. Still, for vehicle manufacturers, it will be crucial to understand how we judge risk and how they can improve their technology. Therefore, we are advising them on how to do it.

Collaboration in general

Theoretically, we could collaborate with other companies that do what we do. However, the cameras are too different. We want to get as close with the system we train on the data. We want the data to be as close to the data that will come in later and have to be assessed. So there is a difference between what we do and our system compared to, for example, a camera that is mounted on the outside of the vehicle and has the windshield in between. When it rains, there will be droplets on the windshield. There might be dust or dirt. The light might come in from an unfortunate angle. That is why we use our own data and label our own data.

Another problem is that we always have new ideas regarding objects and events that we want to configure. We want to decide what is being recorded, find out which are the thresholds, what is considered risky, and when we want video footage. For example, do we want bikes, or pedestrians, or cars, or a combination of them? Say we find out that there have been a considerable number of accidents with e-scooters lately. But we have never discovered e-scooters and do not know what is going on. In this case we can tell our system, «I want situations that were dangerous and had an e-scooter involved. Collect the data.» And when we have the data, we can see what is typically going on in these situations.

Monetizing the data

Stefan Sellschopp, Allianz SE: This topic is of relevance for fleets, less so for retail. Some fleets, such as delivery services, have a hard time finding insurers. Using a risk-mitigation measure is a way to prove to the insurance provider that the customer who wants to be insured is a responsible driver. At Allianz, we have discussed how to leverage the Peregrine technology to insure fleets, thereby reducing the risk and making them more manageable.

For many insurers in Germany, Austria, Switzerland, and France, fleet insurance is not a profitable business. Insurance companies are spending 7 percent more for claims settlements and administration compared to what they earn in premiums. Usually insurers do not just take the risk, but they also have their specialists assess the risk and advise the customer how to mitigate it. That does not work in fleet insurance. The specialists from the insurance company would need to sit next to the drivers, assess their behavior, and tell them how to do things differently, if necessary, without being able to check on them all the time. However, with the Peregrine system, these controls are now possible.

The interview was conducted in March 2021.

Jorit J. K. Schmelzle
Co-founder, Peregrine Technologies

Jorit Schmelzle is co-founder and CPO of Peregrine Technologies, which uses AI to derive actionable insights from video and other data sources so their customers can better manage their vehicle fleets. These insights can be used for logistics, safety, mobility and other applications. Its two revenue streams include a licensable SDK as well as the sale of its data assets.

Jorit received his Master Of Arts & Science at the Swiss Federal Institute Of Technology (ETH) Zurich and his Bachelor Of Arts & Science at the Indian Institute Of Technology in Bombay.

(Continued)

Stefan Sellschopp

Senior Consultant Connected Car – Claims, Allianz Partners, and co-founder e-REVOLT

Within his role at Allianz Partners, Stefan Sellschopp has been building Minimum Viable Products (MVPs) and turning the successful services into products. With his team, he did market studies to find possible partners. He did crash test with Allianz Center for Technology to find out how the crash detection solutions deliver and then launch the automated first notification of loss solution in the market, with an ambition to roll it out to other markets.

In his other role as a co-founder of e-REVOLT, he is responsible for the networking and connectivity part of the e-R3VOLT EV retrofit solution – upgrading an existing Internal Combustion Engine vehicle to enable a longer and carbon friendly life cycle.

Among many other things, Stefan spent several years on EV projects in Silicon Valley and supported the Audi e-tron IT development.