



Doubly optimized: reinforcement learning agent & neural network

06/05/2022 Porsche Engineering has combined two artificial intelligence methods to efficiently optimize a side skirt: A reinforcement learning agent finds the best variant of the component – supported by a neural network that greatly accelerates the individual optimization loops.

They are inconspicuous lifesavers: The two side skirts protect the occupants in case a vehicle leaves the road and has a side impact with a tree, for example. In all-electric vehicles, they have another important task – protecting the high-voltage battery, which is housed between the left and right side skirts and must not be excessively deformed in the event of an accident. Developers invest a corresponding amount of work in optimizing the crash behavior of the component.

They have a clear idea of the optimal crash behavior of a side skirt: The deformation (intrusion) should have as steady a progression as possible without too many local fluctuations. This relationship between external force during the crash and side-skirt intrusion can be represented by an ideal curve. The actual behavior of the component should come as close as possible to this ideal curve.

During development, simulations are usually used for such optimizations, because prototypes for crash tests are expensive. In the widely used finite element method (FEM), the real vehicle or component is usually broken down into small squares ("elements") for crash simulations, the corners of which are called "nodes." The behavior in the event of a collision can be calculated from the reaction of the individual elements to external forces – and the more elements and nodes the FEM model contains, the more precisely it can be calculated. "However, higher accuracy is also reflected in greater computational effort," explains Johannes Pfahler, computational engineer at Porsche Engineering. "FEM simulations require high-performance computers, and for a complete vehicle model, a single crash calculation can take a day or more."

This was also true for the crash simulation of the side skirt in an all-electric vehicle. The component originated from a pre-development project at Porsche AG. To optimize the component, Porsche Engineering engineers were able to vary 14 wall thicknesses and compare the crash behavior of each variant with the ideal curve. If the computational engineers had used an FEM overall vehicle model with 9.6 million elements and 6.6 million nodes for this purpose, a computer with 128 microprocessor cores (CPUs) would have been busy for almost 20 hours – just to calculate the crash behavior of a single skirt variant. A simplified FEM partial model consisting of a side skirt, an underbody, and a battery construction (1.6 million elements and 1.9 million nodes) would still have occupied a computer with 64 CPUs for just over four hours per simulation. So a more efficient method was required to examine hundreds or thousands of skirt variants with different wall thicknesses with regard to their individual crash behavior.

A case for the RL agent

This is where Dr. Joachim Schaper, Senior Manager of AI and Big Data at Porsche Engineering, and the company's other AI experts came into play. Their idea: combining the two AI methods of reinforcement learning (RL) and neural networks (NN), it should be possible to find the optimal skirt variant with a reasonable computational effort. In reinforcement learning, an algorithm (called an "agent") interacts with its environment and constantly learns through feedback – by being rewarded with bonus points for actions that lead to a good result and penalized with deductions for failures. Its goal in this case was to find the variant with the highest possible score over time. To do this, the RL agent had to get closer and closer to the ideal crash behavior of the side skirt by varying the 14 wall thicknesses.

How well it succeeded in doing this in each individual case could in principle be determined by an FEM-based simulation of the crash behavior of each proposed variant – if the computational effort and time required for this were not far too high. This is significantly reduced by the second AI method used: The feedback for the RL agent is not provided by a simulation, but by a neural network. "We previously used simulation data to train it to predict the crash behavior of the side skirt as an output using the wall thicknesses as input," explains Schaper. "This takes only seconds for each side skirt variant, as opposed to hours for a conventional simulation. Yet we don't have to make any sacrifices with regard to feedback accuracy with the RL agent."

With this combination of the two AI methods, the skirt optimization could get underway. The 14 wall thicknesses for the first loop were not chosen at random, but came from Dr. Philipp Kellner, specialist in body predevelopment at Porsche AG. "As an expert, he knew from many years of experience what a good combination might be," says Pfahler. "So the reinforcement learning agent was well set up right from the start." The neural network then compared the crash behavior of this particular side skirt variant with the ideal progression and fed the deviation as feedback to the agent, which in turn delivered the next combination of 14 wall thicknesses to the neural network. This process was repeated several thousand times, and with each iteration the AI team of the RL agent and neural network came a little closer to its goal of finding the optimal side skirt.

This is also where the power of reinforcement learning became apparent. "We thought for a long time about which AI method might work best for this task," Schaper reports. "The advantage of reinforcement learning: The agent remembers the history of its attempts. So it knows what has worked well so far and what hasn't – an important prerequisite for developing a good strategy for the efficient optimization of the side skirt." His confidence in the RL approach proved justified: After several thousand iterations, the two AI methods produced a side-skirt variant that sufficiently matched the ideal curve. The goal was achieved – in a fraction of the time that a conventional approach would have required.

Learning from experience

In total, the developers needed about 23 days of computing time to reach their goal. By far the greatest effort went into the 548 crash simulations used to train and validate the neural network. They were based on a further simplified FEM model of the side skirt, consisting of only about 132,000 elements and 129,000 nodes. Each crash could thus be calculated for each variant in one hour – on a normal workstation with four CPUs. "If we had searched for the optimal side skirt using traditional simulations among 548 variants, a high-performance computer would have been busy for 96 days," says Pfahler. "In addition, we might have missed the optimal wall thicknesses if we had looked for them in the wrong part of the parameter range – which would have entailed further costly simulations."

The optimization project started at the end of 2020 with an analysis of the current state of the art. By the first quarter of 2021, the training and validation dataset was already available for the neural network, which was trained in the second quarter. At the same time, the AI experts at Porsche Engineering developed the RL agent and then had it execute the side skirt optimization. They are currently working on using the successful pairing of the two AI processes for further optimizations. "For example, we want to find out whether the pre-trained RL agent needs to develop a completely new strategy for a different crash or component, or whether it can reuse the old one," says Pfahler.

More efficient implementation as a target

In the future, the RL agent will be used to find optimization potential for front and rear crashes. Schaper,

Pfahler and their colleagues also want to investigate how the method can be implemented even more efficiently. "We trained the neural network for side-skirt optimization with 300 datasets, which allowed it to yield very precise results," Schaper reports. "That's an unusually large amount of data, though – usually only ten to 20 simulation results are available for training." The focus is now on asking questions like: Can the neural network still provide accurate information about crash behavior even with significantly less training data? Where is the ideal point of compromise between training effort and precision? What error rate of the neural network can be tolerated? "What is already certain is that we can use AI to optimize components," Pfahler says. "Now it's a matter of improving the process and integrating it into series development."

The AI experts at Porsche Engineering have already learned from the PERL project that the combination of reinforcement learning and neural networks can yield valuable insights: With RL, they succeeded in automatically finding an optimal application strategy for engines that can be used universally for models with different designs and displacements as well as with different turbocharging systems. "There are many parallels between PERL and the current project: In both cases, for example, it is a matter of finding the optimum constellation in a large search area," explains Schaper. "And it has been proven again: Reinforcement learning is the *crème de la crème* in terms of AI for optimization tasks."

In brief

By combining the two AI methods of reinforcement learning and neural network, Porsche Engineering was able to optimize a side skirt very efficiently without the use of high-performance computers. Now the investigation will explore how the method itself can be further improved and in what other areas it could be used.

Info

Text first published in the Porsche Engineering Magazine, issue 1/2022.

Author: Christian Buck

Illustrations: Jurij Chrubasik

Copyright: All images, videos and audio files published in this article are subject to copyright. Reproduction or repetition in whole or in part is not permitted without the written consent of Dr. Ing. h.c. F. Porsche AG is not permitted. Please contact newsroom@porsche.com for further information.

**MEDIA
ENQUIRIES**



Sandro Kälin

Head of Communications Porsche Schweiz AG

+41 41 487 91 16

sandro.kaelin@porsche.ch

Link Collection

Link to this article

https://newsroom.porsche.com/fr_CH/2022/innovation/porsche-engineering-doubly-optimized-reinforcement-learning-agent-neural-network-28254.html

External Links

<https://www.porscheengineering.com/peg/en/>