VORTECH

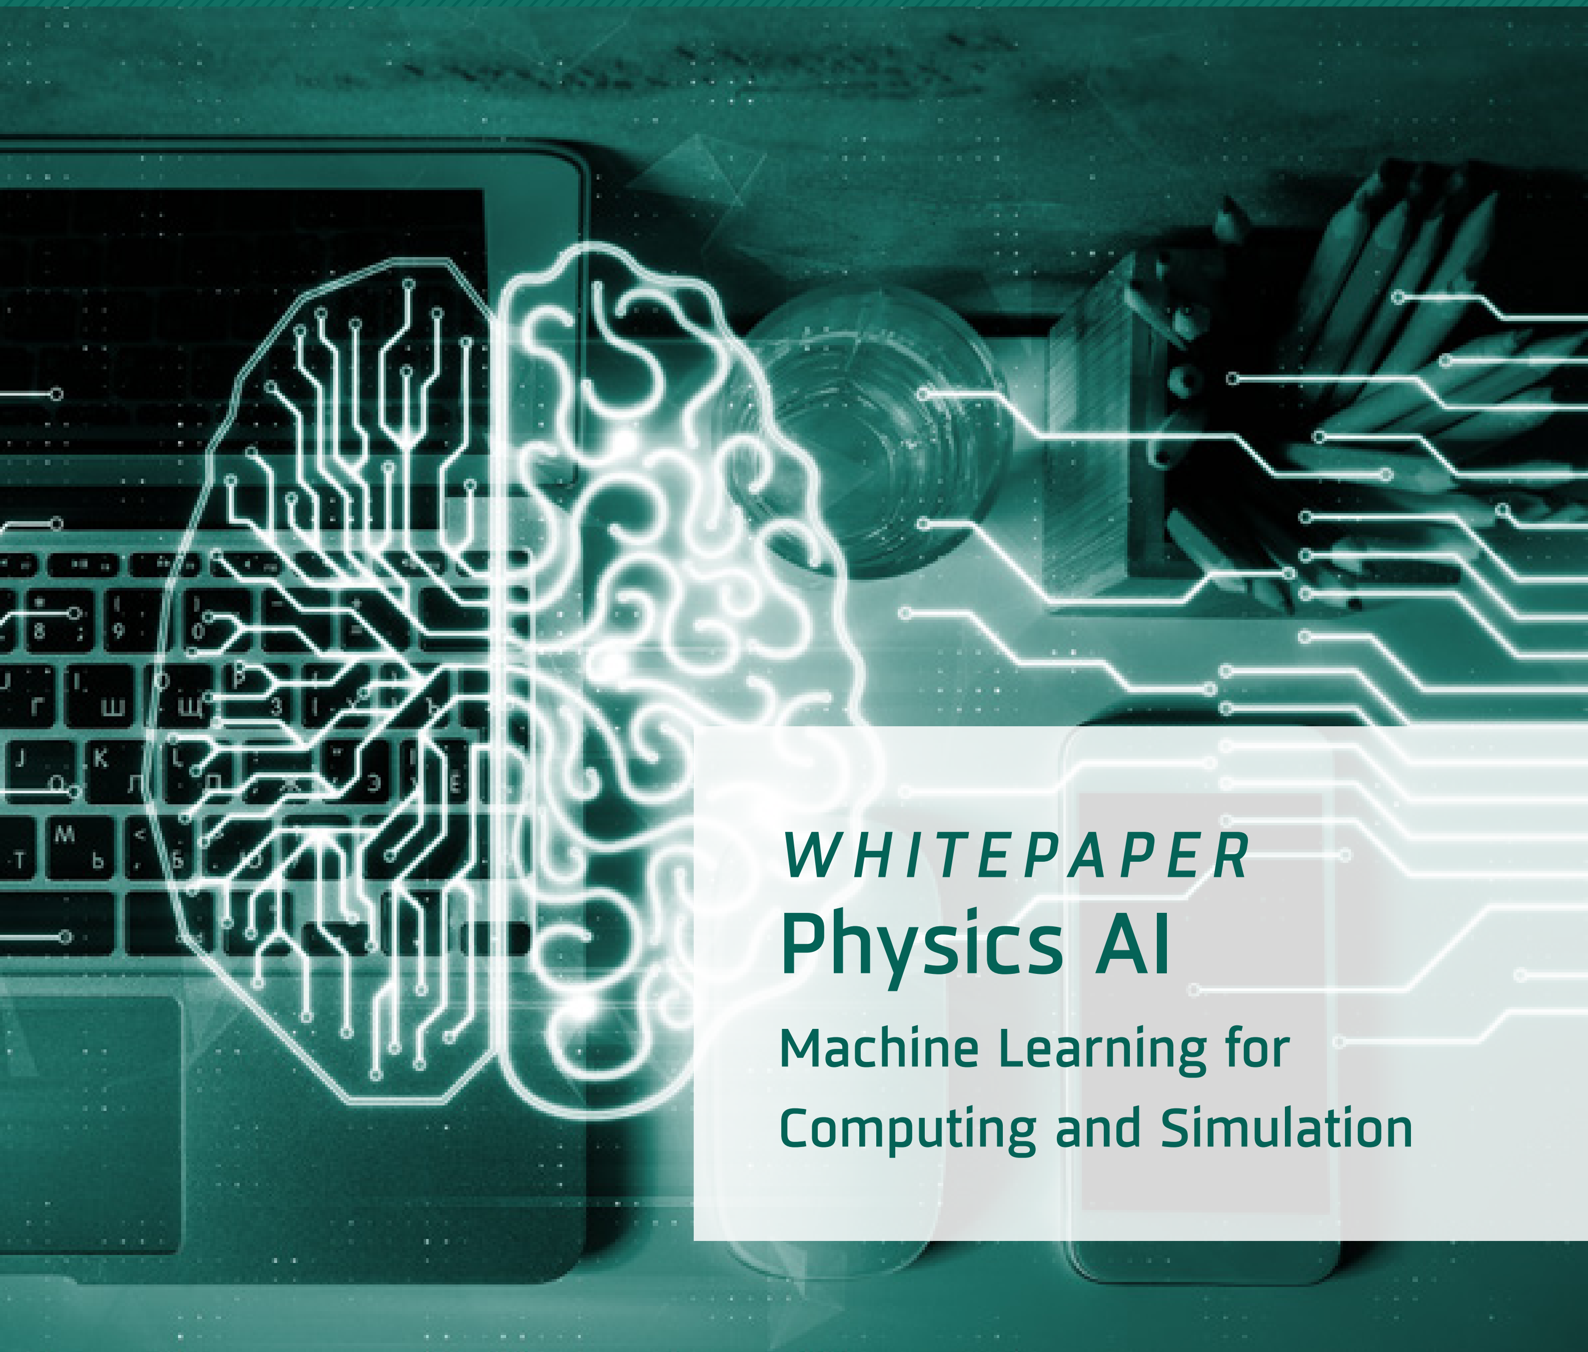

WHITEPAPER

Physics AI

Machine Learning for Computing and Simulation

Introduction

Most people know machine learning from the Large Language Models like ChatGPT. But there is an entirely different type of machine learning that offers tremendous opportunities for the simulation of industrial and environmental processes.

This other type of machine learning, called physics AI or scientific machine learning, uses a combination of traditional mathematical modelling, scientific insights and machine learning techniques. It allows for training reliable machine learning models even with the sparse and noisy data that is typical for many applications in industry and environmental management.

In this whitepaper, we lay out our vision on how this technology will lead to exiting new opportunities, expressed in three main use cases, and how we can help to leverage it.

The Use Cases of Physics AI

There are three main use cases, where physics AI can be transformative with respect to computing and simulation. We will discuss these in turn.

Use case 1: Making Simulations Much Faster

Simulations of complex physical processes are typically very compute intensive. This is problematic when a model must respond fast, like in the case of model-in-the-loop control or in crisis situations. Long runtimes are also an issue when a model must be run many times over, for example in optimization, in scenario studies or for uncertainty quantification.

For these time-critical situations, a machine learning model can be trained to replace a full-blown traditional simulation. Once trained, such a machine learning model, called a surrogate model, will provide answers in a fraction of a second. A nice example can be found in this case study.

However, there is usually not enough observational data from the process that must be simulated to train a model. Some additional data can be generated from runs of the traditional models, but this too will be limited: after all these models were too compute-intensive in the first place. This problem can be mitigated by leveraging the scientific and engineering understanding of the simulated processes. Feeding such understanding into the training phase provides the machine learning with enough extra information to learn a fast and reliable model even from limited data.

An entirely different way to speed up a simulation is by using a machine learning algorithm instead of the central solver that underlies a typical simulation. Inferencing a machine learning algorithm can be much faster than using a traditional solver.


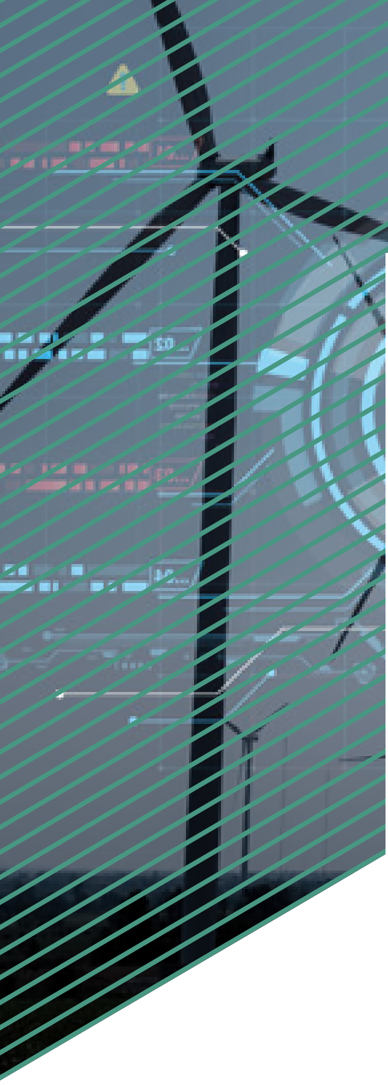

Use Case 2: Closing the Gap between Model and Reality

Models are, by definition, only an approximation of reality. Even the most sophisticated simulation models usually don't fully reproduce observations of a particular process in the real world.

Sometimes this is because parts of the process are not easily captured in nice expressions. For example, there may be a part of the process that is not fully understood but must nevertheless be accounted for. Another well-known example is a process that plays out on multiple scales. In this case, an explicit model can usually be formulated for the large scales, but something must be done to account for the effects of the smaller scales.

In these cases, a closure model can be added to the original model to represent the part of the model that is missing. Such a closure model can be a neural network that is trained on the difference between the model and the observations. If there is some physical understanding of the closure terms, this can be leveraged during training.

Finding the closure model can also be done by learning the expression that describes the difference between the original model and observations. In this case, machine learning is used to select the combination of mathematical terms that best represents the part that is missing from the original model. Some understanding of the nature of the missing part will help the learning process by restricting it to meaningful terms. This approach is particularly elegant as the expression that is learned may in turn provide physical intuition.

The closure model can either be used directly in the simulation or as a postprocessing step.

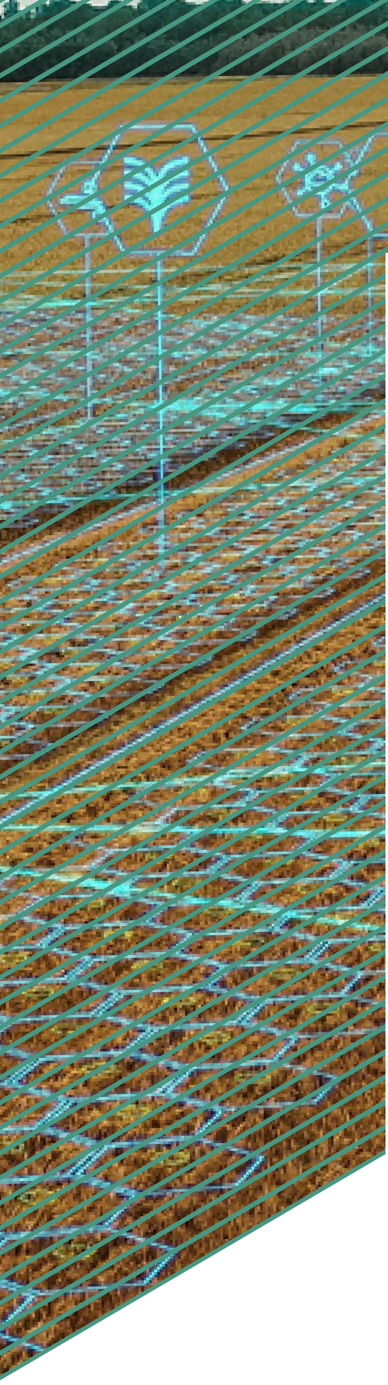

Use Case 3: Learning from Sparse and Messy Observational Data

Today, sensors generate massive streams of data, fast and frequently. But still, these observations are typically relatively sparse in space and time. Sometimes data is missing for a period when communication between the sensor and the data collection platform fails. And sensor data also has other issues like noise and bias.

By leveraging the understanding of the processes behind the sensor observations, it is still possible to create a reliable machine learning model from this patchy data. For example, Physics Informed Neural Networks, or PINNs for short, use physics in training the network to nudge it towards physically consistent behaviour. Also, symmetries in the observed processes, rotational invariance, conservation laws or other structural information can be used to augment the data and provide extra information for machine learning on top of the sensor data.

Machine learning can also be used to train an inverse model, that is: a model that computes backwards from the observations to the full state of the system from which the observations were taken. Traditional models only compute from the input to the output, but with machine learning it is also possible to learn the relation from the output to the input. Or how the input and output of a model relate to the parameters of the model.

This can be used, for example, if you have observations of air quality at some point to determine the likely source of air pollution.

How we help our clients to use machine learning

VORtechs mission is to help its clients develop the best possible software for simulation, optimization and forecasting. Therefore, we are eager to bring these new physics AI methods to the market wherever it brings benefits.

VORtech has always provided the combination of mathematical knowledge, software development skills and a good understanding of processes in physical reality. This makes us uniquely qualified to help our clients apply physics AI and sets us apart from other machine learning consultancies.

We provide our services always in close collaboration with the experts and software developers of our clients. They have all the required knowledge about the application and should be closely involved in everything we do. Not only because that leads to the best solutions but also because our clients are best served if their own employees gain knowledge about these new technologies.

These are the services that we provide:

Consultancy

If you wonder whether these techniques could work for you, we can do an assessment for you. We'll talk with your experts to understand your application, study the software that is already there and then report the findings.

The outcome could be that there are interesting opportunities. But we will also honestly say if we think that physics AI does not offer sufficient benefits or is going to be too costly. After all, our mission is to help our clients create the software that is best for them, not to sell the latest hype.

If there are benefits in physics AI techniques, we can support you in the next step. In close collaboration with your experts, we can work out a plan to introduce the new technology. Such a plan will have go/no-go moments to make sure that risks are properly managed. In choosing the approach, we will carefully consider how important it is to have an explainable model and choose the appropriate machine learning approach.

If you are already using machine learning, we can provide you with a second opinion or an audit. For example, you may wonder how reliable the model is or if it is properly tested. We will report our findings extensively so that you fully understand our advice and take the appropriate actions if needed.

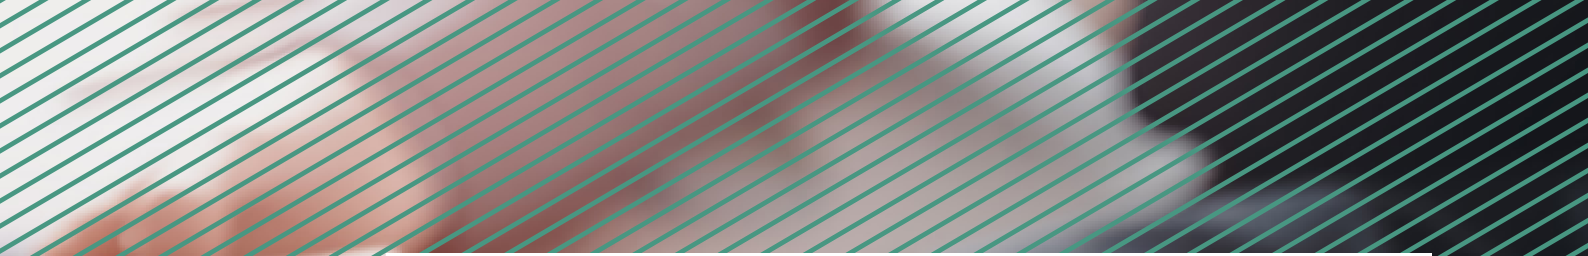

Training fast models to replace a compute-intensive model

Knowing the techniques and their pitfalls, we can support your experts in developing a machine learning model to replace a much more compute-intensive model. We can also develop the model ourselves, but not without guidance from your experts as they are far more knowledgeable about the application than we can ever be.

In any case, all the software and training scripts that we create will become your property. This way, you can recreate or retrain the model without our help.

Improve existing models using machine learning

If you have sensor data and you want to use it to improve your computational model, we are here to help. Note that machine learning is not necessarily the answer. In many cases, more traditional data assimilation methods are quite sufficient, and we can also help you apply these. But if machine learning is indeed appropriate, we can help you train a closure model to account for effects that show up in the sensor data but are not yet in the model.

Train a machine learning model from your sensor data

If you don’t already have a simulation model, but you do have lots of sensor data from some process, we can help you create a machine learning model from this data. As mentioned above, we do require the collaboration with your experts as they likely understand best what the sensor data means and what process is generating the data.

And again, we can either train your experts to build the model, or we can do so ourselves with input from your experts. But even if you let us do all the work, all the software that we write and all the scripts that we use for training will become your property.

Why VORtech?

VORtech has a long history of helping its clients to develop software for simulation, optimization and forecasting. We've worked for major companies and institutes in the Netherlands and abroad. In most cases we build a long-term relationship with our clients, developing and learning together on a basis of trust.

Unlike many newcomers in the field, we bring decades of experience delivering large, high-tech projects. We are experienced in doing such projects in a controlled way. Also, we know more methods than just machine learning. Therefore, we can choose the most appropriate method for your challenge, whether it is machine learning or not.

With over 35 colleagues, many with a PhD in STEM and with extensive experience in developing computational modelling software, we feel confident that we can provide the knowledge, skills and expertise that you need to leverage the potential of physics AI.