

Using neural networks to quickly approximate the results of computationally-intensive simulations

A collaborative research effort finds that design teams can now use neural networks to evaluate fluid dynamics issues and dramatically increase the number of design options they consider.

ABSTRACT

This paper explores a collaborative research effort that demonstrated the power of artificial intelligence to stand in for the traditional simulation and modeling workloads that typically run in high performance computing systems. In this project, our research team determined that the use of trained neural networks can allow design teams to quickly evaluate the viability of many different design alternatives in comparison to running computationally-intensive simulations for each alternative under consideration.

Dell Technologies HPC & AI Innovation Lab
April 2021

TABLE OF CONTENTS

EXECUTIVE SUMMARY	1
FASTER FLUID SIMULATIONS	1
BUILDING THE MODEL.	2
PUTTING RATTLER TO THE TEST	2
KEY RESULTS	3
LEARN MORE	3

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying and distribution of any software described in this publication require an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, Dell Technologies, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, Intel Logo are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be the property of their respective owners.

Dell Technologies believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Published in the USA 04/21.

EXECUTIVE SUMMARY

Computer-driven fluid-dynamics simulations are an essential tool in the design of cars, trucks, ships, airplanes and other forms of conveyance. Using simulation techniques, design teams can now test different configuration options in software to optimize their designs to streamline the flow of air or water around their products and minimize the associated resistance.

The use of these computerized simulations has brought a great leap forward from the days of building physical models and testing them in labs. However, this data-driven approach to engineering does come with certain limitations. One of these drawbacks is the time required to run complex simulations in a high performance computing (HPC) lab. Even with supercomputing-class HPC systems, a single simulation might take a day or more to run. The difficulty here is that with many nonlinear systems, traditional methods of calculation often carry a heavy computational burden due to complex dynamics, physical constraints or multi-scale behavior.

To address this challenge, our team of researchers from Carnegie Mellon University, Los Alamos National Laboratory, Massachusetts Institute of Technology, the University of Massachusetts and Dell Technologies explored a novel methodology for forecasting complex nonlinear spatio-temporal dynamics through purely data-driven methods. This work introduced a machine learning method for studying turbulent flow dynamics using high-fidelity simulation data.

In our proposed approach, we showed how design and engineering teams can replace traditional, computationally-intensive numerical simulations with fast machine learning models. We demonstrated that while there is some loss in fidelity using a machine learning approach, our models are suitable for many engineering problems. In particular, our study found that our neural network models are an ideal candidate to include in workflows that require rapid fluid dynamics calculations for design optimization.

While our present work pertains to a particular dataset, the results indicate a promising pathway to address challenges with current numerical techniques for evaluating non-linear spatio-temporal dynamics.

FASTER FLUID SIMULATIONS

In this investigation, our research team evaluated how to most efficiently place lidar arrays on autonomous cars, especially battery-powered cars. In particular, we sought to determine the optimal positioning or placement for the lidar arrays in terms of airflow.

To make a determination like this, design teams need to simulate the airflow around the vehicle, which is a computationally-intensive process with workloads that might run for an entire day in a high-end HPC system. While design teams would ideally run hundreds or even thousands of simulations to determine the optimal placement for a lidar array, the computational time requirements make this number of iterations highly impractical, if not impossible.

To remove this barrier, and to enable faster fluid dynamics simulations and more design iterations, our team explored the use of machine learning techniques to approximate the air flow around a vehicle and calculate a drag coefficient without having to do direct simulation. This data-driven approach would allow design teams to do a coarse approximation to see if they are in the right ballpark in terms of the proposed placement of the lidar array, or any other component, and do it on thousands or tens of thousands of different configurations in a relatively short amount of time. Instead of needing a day to run a simulation and calculate the drag coefficient around the car, they might need only a few seconds to approximate it using a neural network.

BUILDING THE MODEL

To build a machine learning model that negates the need for direct simulation, we began, somewhat ironically, by using simulation to generate a dataset based on simulation. In other words, we are not only using artificial intelligence to replace simulation, but we are using simulation to train our AI model.

While the technical details are beyond the scope of this paper, the simple view is that we had several different objects in a fluid flow, and we simulated the drag coefficient around those objects. We then used that information as training data for the neural network. This allowed us to use very high-fidelity, high-realism simulations as the training data for this neural network, and then get those approximations up to very high accuracy levels, at least under certain conditions. We could then, going forward in time-critical simulations, replace the simulation component with the artificial neural network.

To carry out this work, we used a type of neural network called a neural ordinary differential equations solver, or neural ODE solver. This tool is a neural network designed to solve differential equations. We began the process by doing some encoding at the front end to change the frame space. We then ran the initial conditions through the neural ODE solver and subsequently returned the result to the original reference space to see the current state of the system.

We started with the neural ODE solver because that already exists in the literature, and because it addresses the same type of problem as fluid simulation — fluid dynamics is a partial differential equations solve. So we started with this existing technology, added capabilities to it and retrained it on the data that we wanted to actually approximate. And that gave us the ability to use these approximation tools as a stand-in for traditional simulation.

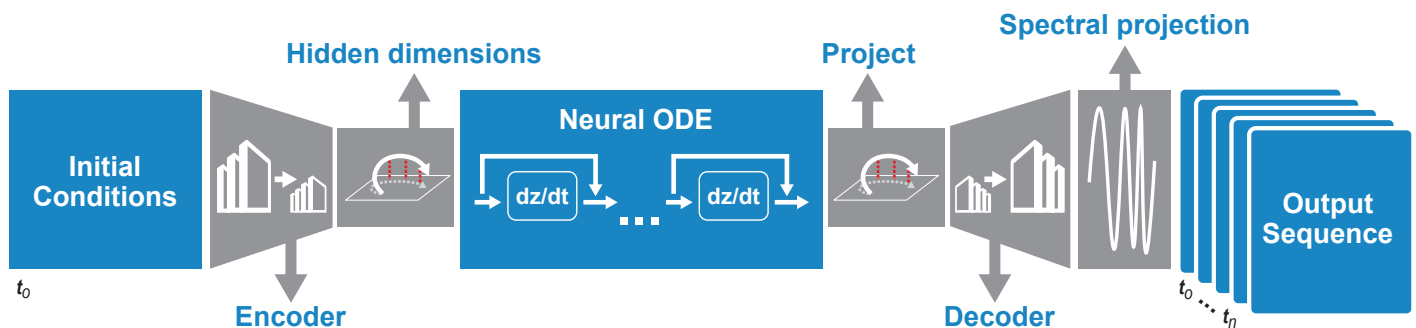


Figure 1. Schematic of the overall model architecture. Initial conditions are given and encoded into a latent space. Augmented channels are concatenated and the dynamics are forecasted through the neural ODE approximator. The sequence is decoded and the divergence-free condition is enforced through the last layer.

For our simulations that calculated the drag coefficient around car bodies, it could take several hours to compute a final answer on multiple GPU-accelerated nodes. Once we finished training the neural network approximation model, we could get that answer within a few seconds. So we achieved a dramatic reduction in time to solution.

PUTTING RATTLER TO THE TEST

We did our initial simulation work and trained our model on the Rattler supercomputer in the Dell Technologies HPC & AI Innovation Lab in Austin, Texas. Rattler, based on Dell EMC PowerEdge C4140 servers, is a GPU-accelerated cluster that is a frequent host for supercomputing collaborations. This cluster, which is the result of a partnership between Dell Technologies and NVIDIA, is designed for extreme scalability by leveraging NVIDIA V100 GPUs with NVIDIA® NVLink® high-speed interconnect technology and NVIDIA InfiniBand® networking.

The Dell Technologies HPC & AI Innovation Lab

The Dell Technologies HPC & AI Innovation Lab encompasses a 13,000 square foot data center in Austin, Texas. This data center, which is devoted to advancing high performance computing and artificial intelligence, houses thousands of servers and a wide range of storage and network systems. Bringing together HPC operational excellence and expertise, the lab is staffed by a dedicated group of computer scientists, engineers and subject matter experts who actively partner and collaborate with customers and other members of the HPC community. [Learn more about the lab.](#)



Rattler was a key resource for running the initial simulations that generated our training data and then for training our model. But once the model was trained, the computational requirements became very small. At that point, the model was ready to be used as an inference engine. For that work, a single GPU or even a very powerful CPU might be all that is needed to get the same types of answers that would have taken an entire supercomputer were it not for the neural network. This is one of the big benefits of neural network-based approximation techniques.

KEY RESULTS

In the investigations in the Dell Technologies HPC & AI Innovation Lab, our research team determined that it is viable to use a neural network to quickly approximate the results that would otherwise be achieved with more precise but extremely time-consuming computer simulation techniques. We further showed that in cases where an approximate result is sufficient, we can greatly reduce the time required to simulate the dynamics of different design options and then narrow down those options to the ones that appear to be most viable.

Our results suggest that design teams can now use neural networks and machine learning techniques to dramatically increase the number of design options they consider. With a model like the one we developed in our project, a design team might be able to go from considering a hundred different design configurations for an object to considering thousands or tens of thousands of options in a reasonable amount of time.

Ultimately, this is a case where approximation techniques and machine learning models are poised to drive innovation and optimization in a design space that has traditionally relied heavily on computationally-intensive numerical simulations.

LEARN MORE

For a deeper and more technical view of this research initiative, see our team's paper "[Learning non-linear spatio-temporal dynamics with convolutional Neural ODEs.](#)" And for a visually oriented overview of the project, see the related [video presentation.](#)

To learn more, visit delltechnologies.com/innovationlab.