



# 2nd Northern Ireland High Performance Computing User Conference

Intelligent Systems Research Centre,  
Ulster University, Magee Campus,  
Derry/Londonderry,

10th November 2022

# Organizing Committee

## CHAIRS:

Prof Damien Coyle

Prof Roger Woods

Mr David Smith

## LOGISTICS:

Mrs Louise Gallagher

Mr Senhui Qiu

Dr Karl McCreadie

## MEMBERS:

Dr Vaughan Purnell

Dr Jose Sanchez Bornot

Dr Luis Fernandez Menchero

Mr James McGroarty

| <b>Conference Programme</b> |  |
|-----------------------------|--|
| <b>09:00</b>                | <b>Arrival and registration</b>  |
| <b>09:30</b>                | <b>Introduction welcome:</b> Damien Coyle  |
| <b>09:45</b>                | <b>Keynote 1:</b> Marcel van Druenen (DELL)<br>Title: A hardware vendor's view on real-life HPC and AI use cases   |
| <b>10:15</b>                | <b>T1:</b> Senhui Qiu (Intelligent System Research Centre, UU)<br>Title: Accelerating optimisation of Deep Predictive Coding Networks with Bi-directional Propagation using the NI-HPC                                   |
| <b>10:30</b>                | <b>T2:</b> Charles Gillan and Stephanie Beck (Centre for Data Science and Scalable Computing, QUB)<br>Title: Application of neural networks and other AI methods to spice fraud detection                                |
| <b>10:45</b>                | <b>Comfort &amp; Tea Break</b>   |
| <b>11:00</b>                | <b>T3:</b> Vahid Saranirad (Intelligent System Research Centre, UU)<br>Title: Bioinspired Artificial Intelligence optimisation using High-Performance Computing  |
| <b>11:15</b>                | <b>T4:</b> Cathal Murray (Intelligent System Research Centre, UU)<br>Title: A comprehensive comparative analysis of state-of-the-art deep learning-based times series prediction algorithms accelerated using the NI-HPC |
| <b>11:30</b>                | <b>T5:</b> Attila Korik (Intelligent System Research Centre, UU)<br>Title: Benefits of a high-performance computing cluster for calibrating brain-computer interface technology  |
| <b>11:45</b>                | <b>T6:</b> Pal Schmitt (Marine research group, QUB)<br>Title: Marine Engineering on NIHPC cluster  |
| <b>12:00</b>                | <b>T7:</b> Evan P. Troendle (Wellcome-Wolfson Institute for Experimental Medicine, QUB)<br>Title: SARS-CoV-2 introductions to the island of Ireland  |

|              |  |
|--------------|--|
| <b>12:15</b> | <b>T8:</b> Jeffrey Johnston (School of Mechanical and Aerospace Engineering, QUB)<br>Title: Large Eddy Simulation of Tilted Wind Turbines for Improving Power Production   |
| <b>12:30</b> | <b>Keynote 2:</b> Venkatesh Kannan (ICHEC)<br>Title: R&D to drive innovation for national high-performance computing and data services.  |
| <b>13:00</b> | <b>Networking Lunch – provided in the ISRC dining area</b>   |
| <b>14:00</b> | <b>T9:</b> Fionntán Callan (Astrophysics Research Centre, QUB)<br>Title: Modelling Stellar Explosions using Monte Carlo Radiative Transfer Simulations in the NIHPC cluster  |
| <b>14:15</b> | <b>T10:</b> Dmitry Makarov (Hydrogen research group, UU)<br>Title: Computational Fluid Dynamics modelling to support the development of hydrogen safety strategies and design  |
| <b>14:30</b> | <b>T11:</b> Abdul-Akim Guseinov (School of Pharmacy and Wellcome-Wolfson Institute for Experimental Medicine, QUB)<br>Title: MD simulations to searching for membrane-facing allosteric sites of GPCRs: Neurotensin receptor |
| <b>14:45</b> | <b>Comfort &amp; Tea Break</b>   |
| <b>15:00</b> | <b>Keynote 3:</b> Kristofer E. Bouchard (UC Berkeley)<br>Title: High-Performance Computing in Neuroscience for Data-Driven Discovery, Integration, and Dissemination   |
| <b>16:00</b> | <b>Discussion panel:</b><br>Building towards a better HPC community – how each user can support the adoption of HPC  |
| <b>16:30</b> | <b>Closing Remarks:</b> Roger Woods  |
| <b>16:45</b> | <b>Tour of ISRC</b>  |

## About the HPC User Conference

In 2019, Queen's University Belfast and Ulster University established the Northern Ireland High-Performance Computing (NIHPC) Facility (<https://www.ni-hpc.ac.uk>) with funding support from EPSRC. It is one of the nine Tier-2 HPC services in the UK and is a powerful facility for researchers to conduct large-scale data analytics, simulation and optimization for some of society's biggest scientific questions and challenges.

The conference brings together both existing and potential users of HPC in Northern Ireland, but will be of particular interest to users who are considering the transition of software and research to the facility to accelerate and improve their research. The current program has three keynote International speakers and a number of talks from users in neuroscience, neurotechnology, precision medicine, pharmacy, hydrogen safety and physics. The HPC user conference is focused on HPC user experience and aims to develop case studies that promote HPC adoption among researchers with challenging analyses who see the benefits of using HPC. Mainly, the present conference aims to highlight

- New user experiences in research software engineering.
- Insights into accelerating code through parallel computing on multiple CPUs and/or GPUs.
- Performance and improvements in applications, particularly those areas new to HPC.

## **KEYNOTES AND ABSTRACTS**

### **Keynote #1: Marcel van Drunen**

**EMEA HPC Technical PreSales, Dell Technologies**

**Title: A hardware vendor's view on real-life HPC and AI use cases.**

**Bio:** Marcel van Drunen has been working in the Dell EMEA HPC&AI PreSales team for fifteen years, being its leader the last five years. This team assists HPC&AI customers in determining their HPC&AI infrastructure strategy, by providing assistance with benchmarks and proof of concepts, and detailed correlation and discussion of the various roadmaps. Marcel studied Applied Mathematics and Computer Science at the Technological University of Delft. Prior to his career at Dell he worked for Unisys as a High-Availability consultant.

**Abstract:** In this keynote Marcel will try to give an overview of the kind of HPC and AI business his team gets involved in on a daily basis, across all of EMEA and across all types of HPC and AI customers. This overview will reveal some unexpected trends, and maybe some recommendations for future users of HPC&AI technology.

## Keynote #2: Venkatesh Kannan

### Irish Centre for High-End Computing (ICHEC)

**Title: R&D to drive innovation for national high-performance computing and data services.**

**Bio:** Venkatesh Kannan works as the Technical Manager at the Irish Centre of High-End Computing (ICHEC), the national centre for high-performance computing (HPC) in Ireland, whose core mission is to deliver HPC-related capabilities and expertise to higher education institutions, enterprises and public sector organisations on behalf of the Irish State. At ICHEC, Venkatesh is responsible for defining and implementing the technological vision, strategy and roadmap of the Centre's activities. This includes developing partnerships with domestic and international organisations in academia, industry and the public sector to undertake publicly and commercially funded computation-driven R&D and projects. Venkatesh also manages and coordinates the Centre's technical activities with all members of staff and computational scientists across offices in Dublin and Galway. ICHEC focuses on HPC and its applications to a wide range of areas, including high-performance data analytics (HPDA), big data management, artificial intelligence, optimisation of computational workflows and applications on a variety of computing platforms including extreme-scale HPC clusters, edge computing devices and quantum computing platforms. These are applied to a wide range of application domains such as environmental sciences, remote sensing, digital healthcare and digital twins. Venkatesh also represents ICHEC in a number of Irish and European Union programmes, activities and projects.

**Abstract:** Computational science and research is increasingly enabled by data and computing platforms. Particularly, the emerging convergence of heterogeneous data and high-performance computing technologies is complemented by the digital continuum in which data processing and analysis workflows that span edge and cluster/cloud based computational platforms and technologies. In the context, it is essential for HPC centres to deliver platforms and services that are at the cutting edge to enable new science

beyond what is capable on traditional monolithic HPC or cloud systems. This talk will present the next planned national HPC platform and services being developed for Ireland by ICHEC supported by our European strategic and R&D activities.

### **Keynote #3: Kristofer E. Bouchard**

**Helen Wills Neuroscience Institute & Redwood Center for Theoretical Neuroscience, UC Berkeley**

**Title: High Performance Computing for Neuroscience.**

**Bio:** Kristofer Bouchard is PI of the Neural Systems and Data Science lab at Lawrence Berkeley National Lab and UC Berkeley and group lead of the Computational Biosciences Group at Lawrence Berkeley National Laboratory (LBNL). He is a staff scientist in the Scientific Data & Biological Systems and Engineering Divisions at LBNL and adjunct professor in the Helen Wills Neuroscience Institute and Redwood Center for Theoretical Neuroscience at UC Berkeley. His research is focused on understanding how distributed neural circuits produce coordinated behaviors and perception. His lab takes a two-pronged approach to this problem by conducting in vivo neuroscience experiments and developing data science tools.

**Abstract:** Neuroscience initiatives aim to develop new technologies and tools to measure and manipulate neuronal circuits. To deal with the massive amounts of data generated by these tools, the authors envision the co-location of open data repositories in standardized formats together with high-performance computing hardware utilizing open source optimized analysis and simulation codes. Here, I will provide vignettes from research in my lab utilizing HPC to address diverse neuroscience problems.

**Talk #1: Senhui Qiu<sup>1</sup>, Saugat Bhattacharyya<sup>1</sup>, Shirin Dora<sup>2</sup>, Damien Coyle<sup>1</sup>,**

**<sup>1</sup>Intelligent System Research Centre, Ulster University**

**<sup>2</sup>Department of Computer Science, Loughborough University**

**Title: Accelerating optimisation of Deep Predictive Coding Networks with Bi-directional Propagation using the NI-HPC.**

**Abstract:** Predictive coding (PC) is an influential biomimetic model that can account for diverse perceptual phenomena of the brain. Unlike traditional deep neural networks (DNNs) using the backpropagation (BP) training algorithm, PC uses a neuroscience-inspired inference learning (IL) algorithm, which means it can use the local error to update weights and inferences in parallel. This unique property enables PC to have the potential to overcome the BP algorithm's inability to update weights in parallel. In addition, PC is a generative model that can achieve classification and reconstruction. However, its practical applications in computer vision are still in the initial stage of exploration. In view of the problems of current DNNs with numerous parameters, high time cost, complex model, and inability to update parameters in parallel, we proposed a novel Deep Bi-directional Predictive Coding (DBPC). Experimental results show that the classification accuracy of convolutional DBPC (99.57%) is better than FIPC3 (98.64%) and other classic DNNs, using fewer parameters. Meanwhile, the original input image can be clearly reconstructed from all representation layers in the proposed model by sharing the weights of the classification task. Our model uses 10-fold Cross-Validation to perform hyperparameter optimization on the MNIST dataset. This model has 4 main hyperparameters including the number of layers (i.e., 2, 3, 4, 5, 6, 7, 8), the learning rate (i.e., 0.01, 0.04, 0.06, 0.08, 0.1), batch size (i.e., 8, 16, 32, 64), and epochs (i.e., 30, 40, 50). These hyperparameters can be combined into more than 420 different combinations ( $7 \times 5 \times 4 \times 3 = 420$ ), which is called search space. The MNIST dataset contains 70,000 grayscale handwritten digits in the size of  $28 \times 28$ . In order to find the optimal hyperparameter, we generated 14 models with different hyperparameter combinations from the above search space through an optimization algorithm (Async

Successive Halving Algorithm (ASHA)). Each model takes approximately 280 hours on a personal computer, the total time for 14 models is about 3920 hours ( $14 \times 280 = 3920$ ). However, we ran these models in parallel on the Northern Ireland High-Performance Computing (NI-HPC) facility and completed all models after about 294 hours. It means that the total time consumption of NI-HPC is about 13.3 times less than that of a personal computer. Therefore, NI-HPC can help us obtain optimal hyperparameters faster and significantly improve our research efficiency and quality.

## **Talk #2: Charles Gillan, Stephanie Beck**

**Centre for Data Science and Scalable Computing, Queen's University Belfast**

**Title: Application of neural networks and other AI methods to spice fraud detection.**

**Abstract:** Chemometrics is a key tool in analytical chemistry facilitating detection of adulteration in spices. AI methods are now used to complement multivariate methods, such as principal component analysis. However, in order to train models to achieve a level of predictive accuracy for commercial deployment, lengthy searches over spaces of hyper-parameters are needed. This paper reports on application of neural networks and other AI methods to spice fraud using The Kelvin system at the EPSRC funded Tier 2 NI HPC Centre. It explores the development, validation and performance of the various approaches in detecting adulteration in food fraud.

**Talk #3: Vahid Saranirad<sup>1</sup>, Martin McGinnity<sup>1</sup>, Damien Coyle<sup>1</sup>, Shirin Dora<sup>2</sup>**

**<sup>1</sup>Intelligent System Research Centre, Ulster University**

**<sup>2</sup>Department of Computer Science, Loughborough University**

**Title: Bioinspired Artificial Intelligence optimisation using High-Performance Computing.**

**Abstract:** Spiking neural networks (SNNs) mimic their biological counterparts more closely than their predecessors and are considered the third generation of artificial neural networks. This research introduces a new type of spiking neural network that draws inspiration and incorporates concepts from neuronal assemblies in the human brain. The proposed network, termed CDNA-SNN, assigns each neuron learnable values known as Class-Dependent Neuronal Activations (CDNAs), which indicate the neuron's average relative spiking activity in response to samples from different classes. A new learning algorithm is also presented that categorises the neurons into different class assemblies based on their CDNAs. These neuronal assemblies are trained via a novel training method based on Spike-Timing Dependent Plasticity (STDP) to have high activity for their class and low firing rate for other classes. The results on multiple benchmarks indicate that the proposed network can achieve higher performance with considerably fewer network parameters than other SNNs in comparison.

The proposed training algorithm includes parameters  $\alpha$  and  $\beta$  for all layers in the network. This research uses Nested Cross-Validation (n-CV) to perform hyperparameter optimisation. Nested cross-validation demands training CDNA-SNN with different combinations of hyperparameters which is very time-consuming. Table Below shows some examples of performing hyperparameter optimisation for CDNA-SNN.

| dataset               | #Samples | Parameters<br>#Trainable | Dimension<br>nCV | #Evaluations | #Evaluations<br>using HPC | Runtime of a<br>Single Eval. | Total<br>No HPC: | Runtime<br>HPC: Total |
|-----------------------|----------|--------------------------|------------------|--------------|---------------------------|------------------------------|------------------|-----------------------|
| Iris                  | 150      | 174                      | 5×5              | 260          | 3                         | 60 min                       | 11 day           | 3 hour                |
| Breast Cancer         | 683      | 185                      | 10×10            | 520          | 6                         | 110 min                      | 40 day           | 11 hour               |
| Pima Diabetes         | 768      | 438                      | 10×10            | 520          | 6                         | 130 min                      | 47 day           | 13 hour               |
| Liver disorders       | 345      | 400                      | 5×5              | 260          | 3                         | 90 min                       | 17 day           | 5 hour                |
| Ionosphere            | 351      | 955                      | 5×5              | 260          | 3                         | 75 min                       | 14 day           | 4 hour                |
| MNIST                 | 70000    | 45396                    | 5×5              | 260          | 3                         | 24 hour                      | 260 day          | 72 hour               |
| Fashion-<br>MNIST     | 70000    | 44188                    | 5×5              | 260          | 3                         | 20 hour                      | 217 day          | 60 hour               |
| Neuromorphic<br>MNIST | 70000    | 163074                   | 5×5              | 260          | 3                         | 30 hour                      | 325 day          | 90 hour               |

As can be seen without high performance computing resources, n-CV for all datasets takes a long time and it is almost impossible to perform. However, parallel computing employing the Tier 2 High-Performance Computing resources provided by the Northern Ireland

High-Performance Computing (NI-HPC) facility significantly reduced the whole training time of optimisation for different datasets. As an instance for the MNIST dataset, 5×5 n-CV is used with 25 different combinations of hyperparameters, leading to 260 evaluations. A single evaluation takes 24 hours, and without parallel computing, the total runtime is about 260 days, making this optimisation impossible with a personal computer. However, using NI-HPC and similar and assuming that we can run 100 jobs in parallel simultaneously, this n-CV takes about 72 hours. So far, our research has led to the following publications:

#### **Talk #4: Cathal Murray, Damien Coyle**

**Intelligent System Research Centre, Ulster University**

**Title: A comprehensive comparative analysis of state-of-the-art deep learning-based times series prediction algorithms accelerated using the NI-HPC.**

**Abstract:** Time series forecasting entails making predictions of time dependent variables. It has uses in many applications in industry, commerce, economics, education, and finance. State-of-the-art deep learning neural networks for time series forecasting often require extensive hyperparameter evaluation and optimisation and trainable parameter learning to generate an optimum forecasting model. Hyperparameter tuning is the adjustment or setting of an AI machines structural parameters to control the learning process with the aim of producing optimum accuracy. Trainable parameters are those updated during the learning process. For every hyperparameter setup to be optimised, the trainable parameters must be learned and depending on network architecture and datasets sizes and complexity of the task this problem can be intractable, especially without computational resources to parallelise training and evaluation. The most primitive and time consuming method of optimization is the gridsearch which attempts to build a model for every combination of parameter values influencing a models' structure during learning. The process of cycling through all combinations of hyperparameters is multiplicative and the number of models produced and tested with this process can be extremely

time consuming. To overcome this we apply Asynchronous successive halving algorithm (ASHA), an optimization method which inherently depends upon parallelization and can explore the hyperparameter space more efficiently. We accelerated trainable parameters learning using NVIDIA V100 GPU nodes on the NIHPC which are specifically suited to deep learning and neural networks. CPUs are more generally suited to serial processing of instructions. When a quantity of jobs are placed on a distributed HPC, the HPC provides parallelization by dispensing jobs from a job queue to different CPU units on the HPC cluster. However, in this analysis each job also requests a GPU to optimize a model using ASHA. A GPU is a processor dedicated to processing mathematical operations in parallel and consists of a number of Arithmetic Logic Unit (ALU) pipelines to parallelize instructions. ASHA uses this aspect of a GPU to train models in parallel with the aim of seeking the optimum hyperparameter combination.

Here we quantify the complexity of the training and optimisation problem:

1. A prediction horizon is the number of data points into the future a time series model is trained to predict. Models were trained to predict 1, 18, 25 and 50 steps in the future i.e., four different prediction horizons.
2. There were nine different model types under test, these included the Multilayer Perception (MLP), Long Short Term Memory (LSTM), Gated Recurrent Unit (GRU), BiDirectional LSTM, Convolutional LSTM (CNN-LSTM), Convolutional LSTM Autoencoder, Autoencoder, Attention mechanism and the Transformer.
3. Cross validation is used to ensure trustworthy AI through guaranteed generalisation performance estimation training and test data to be split into a number of folds and a unique model to be trained on each specific fold – here we employ 10-fold CV
4. To comprehensively compare the models capabilities we employ eight different datasets in the analysis. 2,880 models were to be created for this comprehensive analysis.

To evaluate the difference between parallel and non-parallel training time, a hyperparameter combination of 24 permutations was applied using parallel with ASHA vs non-parallel on an LSTM model. The training time was 2 hrs 44 minutes for GPU processed ASHA training as opposed to 6 hrs 48 minutes for CPU grid search model training. This training time varies depending on the model but using ASHA assisted by the GPU reduces the time taken considerably. As the HPC consists of many CPUs and GPUs to handle a large number of jobs this reduces the amount of time to complete model creation. For example, if 14 GPUs were running concurrently on 2880 models, and each job took 2 hrs 44 minutes then the time taken to complete training of all models would be under 24 days. For zero parallelization on one PC without a GPU, the time taken to produce the same 2880 models would be 816 days or approximately 2.24 years. The reduction of this time could only have been done using a high capacity HPC such as the Kelvin 2 cluster. Consequently, this analysis could not have been achieved without the use of the HPC cluster

### **Talk #5: Attila Korik, Damien Coyle**

**Intelligent System Research Centre, Ulster University**

**Title: Benefits of a high-performance computing cluster for calibrating brain-computer interface technology.**

**Background:** Brain-computer interface (BCI) technology provides an opportunity for the user to control an electronic device using information encoded in electroencephalography (EEG) without involving neuromuscular pathways. However, training AI to accurately decode voluntary modulations in EEG is a challenge and requires significant computational resources.

**Example study:** Currently, we are running 5 studies with similar datasets. In our 3-dimension (3D) arm movement decoding study, 10 subjects are trained over 10 sessions to achieve real-time control on a virtual arm directly through brain activity. Each session

involves 4 runs with actual movement and 4 runs with imagined movements, 2 blocks per run, and 4 trials per target are involved in each block. This scenario for each session results in 32 trials per target for both actual and imagined movements (256 trials per session in total, 2,560 trials for each subject).

**BCI framework:** The (artificial intelligence) AI includes extracting time-frequency patterns (MATLAB/Simulink) and translating these to control signals through a convolutional neural network (CNN) and a long-short term memory (LSTM) based kinematic data decoder module (Python). The feedback is presented with the Unity3D game engine.

**BCI calibration:** The subject-specific calibration of the BCI aims to find hyperparameters that dictate deep neural network architectures and learning ability to enable decoding the kinematic information with maximal accuracy from EEG. The parameters being optimized during the calibration phase can be separated into the following groups: (1) parameters of the preprocessing module for generating optimal features; (2) structural elements of the DL module; (3) Learning parameters of the DL algorithm; (+) Weights within the neural network (NN). Tuning these parameters is a time-consuming process which requires running multiple times the DL algorithm to find a combination of the hyperparameters which results in maximal decoding accuracy (DA).

**Calibration time:** The training of the DL module using a single instance of hyperparameter set with a high-end PC takes 20 minutes per CV fold for each session. As in our virtual arm study, all variations of the hyperparameter options within the above-described parameter groups count 41,472 options, the calculation time with a grid search would take approximately 13,824 hours (i.e., 576 days) per CV fold for each session. Human-supervised heuristic search can decrease the calculation time where a set of hyperparameter combinations are defined based on results obtained in the previous phase of the search. With this method, the number of calibration steps can be reduced significantly, but this strategy requires human supervision and may find only a local optimum in the full parameter space. Asynchronous successive halving algorithm (ASHA) and RayTune (a Scalable Hyperparameter Tuning) methods reduce the number of calibration steps by using an automatic method for inherently dependent parallelization.

**High-performance computing (HPC):** HPC facilities such as the Kelvin 2 cluster in the Northern Ireland HPC not only provide faster computation speed but also enable the use of multiple CPU and GPU (Tesla V100 and A100) to parallelize the analysis of the hyperparameter options. Running ASHA/RayTune optimization on the HPC cluster, the single session calibration time can be reduced significantly e.g., using 12 GPU units in parallel for 42 times repeated ASHA/RayTune hyperparameter optimization loop, the computation time from 576 days (full grid search) decreases to 14 hours.

**Conclusion:** NIHPC reduces the BCI calibration time to comprehensively assess hyperparameters between daily subject training sessions and enables us to perform high-quality research across multiple neurotechnology studies

### **Talk #6: Pal Schmitt**

**Marine research group, Queen's University Belfast**

**Title: Marine Engineering on NIHPC cluster.**

**Abstract:** Marine and coastal engineers make use of HPC applications to simulate structural and fluid dynamics problems. This presentation showcases methods and example applications developed in the Marine Research Group at QUB.

Key topics are:

- 1) Numerical Wave Tanks for the development and optimization of marine renewable energy devices (OpenFOAM);
- 2) Fluid-Structure-Interaction (FSI) for slender structures like turbine blades, foils and plant canopies (OpenFOAM);
- 3) Structural simulation tools for the development of Concrete Lattice Offshore Wind Towers (CodeAster).

**Talk #7: Evan P. Troendle<sup>1,†</sup>, Alan M. Rice<sup>1,†,‡</sup>, Stephen Bridgett<sup>1</sup>, Behnam F. Nejad<sup>2</sup>, Jennifer M. McKinley<sup>2</sup>, The COVID-19 Genomics UK consortium<sup>3</sup>, Timofey Skvortsov<sup>4,\*</sup>, David A. Simpson<sup>4</sup>**

**<sup>1</sup>Wellcome–Wolfson Institute for Experimental Medicine, Queen’s University Belfast**

**<sup>2</sup>School of Natural and Built Environment, Queen’s University Belfast**

**<sup>3</sup>The COVID-19 Genomics UK consortium**

**<sup>4</sup>School of Pharmacy, Queen’s University Belfast**

**\*Correspondence to [t.skvortsov@qub.ac.uk](mailto:t.skvortsov@qub.ac.uk) and [david.simpson@qub.ac.uk](mailto:david.simpson@qub.ac.uk)**

**† The authors consider these individuals to be Joint First Authors.**

**‡ Current address: Milner Centre for Evolution, Department of Life Sciences, University of Bath**

**Title: SARS-CoV-2 introductions to the island of Ireland.**

**Abstract:** As happened throughout the world, the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) virus had a devastating impact as waves of infection spread across the island of Ireland during the coronavirus disease 2019 (COVID-19) pandemic. A better understanding of how these successive introductions have spread could guide future public health responses. As SARS-CoV-2 replicates it slowly accumulates mutations and whole-genome sequencing (WGS) makes it possible to track these genetic changes. Phylogenetic analyses were performed using 7,603,547 GISAID WGS sequences available from across the globe to determine the viral importation events of major viral lineages (e.g., pangolin lineages B.1.177, B.1.1.7 ("Alpha"), B.1.617.2/AY ("Delta"), B.1.1.529/BA.1 ("Omicron BA.1"), and B.1.1.529.2/BA.2 ("Omicron BA.2")) to the island of Ireland as made possible by the Northern Ireland High Performance Computing (NI-HPC) Centre. Using geospatial timeseries analyses of the phylogenies and associated sample metadata across the pandemic, we describe in detail the importation and spread of multiple infection clusters to the island of Ireland for the first time. As expected, introductions spread primarily to adjacent districts, but with some longer distance hops, potentially associated with transport corridors. The methodological approach employed in this study can be applied broadly to

study the spatiotemporal kinetics of viral importations to any defined geographic region where sufficient WGS sequencing data is available. The high level of SARS-CoV-2 WGS, particularly in the UK, facilitated this analysis. Although our conclusions are nonetheless tempered by relatively sparse and variable sampling, this analysis demonstrates the additional value of WGS beyond surveillance of emerging variants. The opportunity provided by the continually decreasing cost of WGS should be used to adopt WGS more widely to monitor and geo-spatiotemporally investigate future infectious conditions.

**Talk #8: Jeffrey Johnston, Juliana Early, Dominic Chandar**

**School of Mechanical and Aerospace Engineering, Queen's University Belfast**

**Title: Large Eddy Simulation of Tilted Wind Turbines for Improving Power Production.**

**Abstract:** As the push for renewable sources of energy continues to accelerate, it is important to seek for optimum system performance, while minimising cost. In a wind farm, it is well established that there is a significant loss in power due to the wake effect – turbines extract energy from the wind, forming wakes which travel downstream and leave less energy available for the downstream turbines. The complex interaction of turbulence due to shear and buoyancy in the atmospheric boundary layer, together with the inherent turbulence of the wake, make it difficult to study these wake losses and the suggested methods of mitigating them through experimental studies alone. Large eddy simulation (LES) is an ever-more popular tool to help understand this complex fluid problem, as it can recreate the significant scales of turbulence with minimal flow-dependent empirical parameters and give a detailed and complete picture of the flow quantities and their statistics. In this study, the Kelvin HPC is used to run the Simulator for On/Offshore Wind Farms (SOWFA), version 2.4.x, developed by the National Renewable Energy Laboratory (NREL) in the US, and built on OpenFOAM 2.4.x. LES studies of one- and two- turbine arrays in neutral (no buoyancy) and unstable (convective) atmospheres are conducted. As it is not computationally feasible to directly model the geometry of the turbines in these cases, they are represented

instead by an actuator line model that calculates the force of the blades on the fluid using the solved velocity field. Specifically, the simulations are used to test a littleinvestigated method for mitigating the wake loss effect by tilting the upstream turbines, thus deflecting the wake. Each case is run for several weeks in parallel across 500 computing cores. It is found that the old kelvin compute nodes (using AMD Epyc processors) generally run these simulations significantly faster (in some cases more than twice as fast) than the new kelvin2 compute nodes (using Intel Xeon processors), even though the simulations must be split across many more nodes. This agrees with some unofficial benchmarks for OpenFOAM. Attempts were made to decompose some cases across many more cores (around 1000), but it was found that a drastic slow down occurred and eventually a complete crash of the simulation. Presently, the simulations have helped to confirm the viability of the tilt control method in a neutral atmosphere, but have also shown that it may be significantly less effective in an unstable atmosphere.

### **Talk #9: Fionntán Callan**

**Astrophysics Research Centre, Schools of Maths and Physics, Queen's University Belfast**

**Title: Modelling Stellar Explosions using Monte Carlo Radiative Transfer Simulations in the NIHPC cluster.**

**Abstract:** Type Ia supernovae are stellar explosions which play a number of key roles in astrophysics such as the production of heavy elements like iron and their role as cosmological distance indicators (which won the Nobel prize in physics in 2011). Modelling these explosions to better understand them is difficult due to the complicated dynamic structure and complex morphology exhibited by different explosions. In our modelling approach we carry out multi-dimensional, time-dependant, radiative transfer simulations using Monte Carlo methods. Solving the radiative transfer problem analytically is usually very challenging and quickly reaches its limit as the complexity of the problem increases. Therefore, it is often advantageous to solve the radiative transfer problem numerically, in our case using the Monte Carlo approach. In this method the radiative transfer is simulated by introducing a large number of test

particles (“photon packets”) which propagate and interact in the supernova in a similar way to photons. Monte Carlo simulations have many advantages including being relatively straightforward to generalise to multi-dimensional problems and being highly parallelisable. One disadvantage of the Monte Carlo approach is that due to the stochastic nature of the method large numbers of photon packets are required in order to reduce the statistical noise of the simulations. Due to this large number of photon packets and the complexity of the models simulated a large number of CPUs are required for our work to be feasible. HPC facilities are therefore essential for our simulations. Our research group uses various HPC facilities worldwide (including the Kelvin2 HPC here in NI) to carry out simulations of supernovae and other astrophysical explosions using our Monte Carlo Radiative Transfer code ARTIS (Sim 2007; Kromer and Sim 2009). In this talk I will discuss the computational methods used in our simulations with particular focus on the improved version of the code (Shingles et al. 2020) which makes ARTIS one of the leading codes in the field in terms of the detail with which the physical processes are simulated (Blondin et al. 2022). I will talk about the computational challenges associated with running the improved code version and our solutions to some of these challenges. For example, the use of OpenMP and MPI shared memory to reduce the memory requirements for individual CPUs during the simulations which can be limiting for our most detailed simulations. Finally, I will comment on the importance of continued support for HPC facilities here in NI not just for research areas new to High Performance Computing but also fields already well established in working with HPCs

**Talk #10: Dmitriy Makarov, Volodymyr Shentsov, Donatella Cirrone, Sergii Kashkarov, Vladimir Molkov**

**HySAFER Centre, Ulster University**

**Title: Computational Fluid Dynamics modelling to support the development of hydrogen safety strategies and design.**

**Abstract:** Hydrogen Safety Engineering and Research (HySAFER) Centre of Ulster University develops models and tools for assessment of hazards and associated risks following hydrogen incidents: hydrogen releases and dispersion, thermal and pressure

effects of jet fires, deflagrations, detonations, high-pressure storage tank rupture in a fire, etc. The thrust in HySAFER research is on numerical studies in close collaboration with centres of excellence in experimental research.

Safety of hydrogen traffic in underground infrastructure (tunnels, car parks, etc.) was the topic of a recent HyTunnel-CS project ([hytunnel.net](http://hytunnel.net)) coordinated by the Centre. The largest risks in case of hydrogen transport incidents are associated with high-pressure gaseous hydrogen storage (up to 700 bar) and its release via thermally activated pressure relief device (TPRD). The studies of hydrogen release and combustion hazards in the course of HyTunnel-CS project heavily relied on Computational Fluid Dynamics (CFD) modelling and massive numerical simulations. In particular, a parametric simulations of hydrogen unignited and ignited releases were performed for different TPRD diameters and release directions. Incident scenarios considered a light-duty vehicle at different locations within underground parking and various parking heights. The research concluded that the present day Regulations, Codes and Standards (RCS) for safety of tunnels and underground structures may be readily satisfied if the hydrogen-fuelled vehicle design follows safety strategies to minimise hydrogen hazards. Thus, length of hydrogen flammable cloud propagation and jet flame were found minimal when TPRD diameter was 0.5 mm or below and its release direction was oriented backward at 45° to the vertical. In such case the hazards to the driver, passengers, public and first responders are comparable to the hazards of an ordinary hydrocarbon fuel vehicle fire and first responders don't need to adopt a special tactics for fire intervention. The drawn research conclusions and safety strategies are included in one of the key project outputs "Recommendations for inherently safer use of hydrogen vehicles in underground traffic systems".

**Talk #11: Abdul-Akim Guseinov<sup>1,\*</sup>, Dmitry S. Karlov<sup>1</sup>, Stephen Garland<sup>2</sup>,  
Bianca Plouffe<sup>1</sup> Irina G. Tikhonova<sup>1</sup>**

**<sup>1</sup>School of Pharmacy and Wellcome-Wolfson Institute for Experimental Medicine, Queen's University  
Belfast**

**<sup>2</sup>NQuiX Ltd, Causeway House, Dane Street, Bishop's Stortford, CM23 3BT, UK**

**\*Corresponding author: a.guseinov@qub.ac.uk**

**Title: MD simulations to searching for membrane-facing allosteric sites of GPCRs: Neurotensin receptor.**

**Abstract:** Molecular dynamics (MD) simulations are extensively used for the search for druggable allosteric sites. Among them, probe MD simulations demonstrate rapid development. In these methods, the target is simulated in the presence of probe fragments, which are a prototype of drug molecules. Recently, our laboratory developed a probe-confined dynamic mapping protocol, which enhanced probe sampling within an area of interest. This protocol was successfully applied to map allosteric sites of GPCRs at various locations.

In this work, we apply this protocol to neurotensin receptor 1 (NTSR1) – a target for drug discovery against a wide range of cancers and several neurological diseases. No membrane-facing allosteric site has been reported for NTSR1, therefore it could be a good blind system to test our protocol.

First, identified 5 membrane-facing allosteric sites of NTSR1 using MDpocket – a geometric algorithm for cavity detection in MD trajectories. We next conducted probe-confined dynamic mapping focusing on these sites. The probe simulations were further explored by MD simulations of a docked probe molecule. The probe simulations allowed us to estimate the accessibility of selected allosteric sites and outline the key probe-receptor interactions.

To study the importance of these sites for receptor signalling, we mutated several residues forming each site. We monitored the signalling of the resulting NTSR1 mutants using bioluminescence resonance energy transfer (BRET)-based biosensors in HEK293 cells. We identified three non-biased and one biased allosteric sites.

From the druggability assessment in the simulations and the NTSR1 signalling pattern in the functional assays, we selected 2 allosteric sites showing promise as targets for ligand design.

## WHO WE ARE

NI-HPC is a scalable supercomputing cluster that supports researchers in QUB, UU, and EPSRC across the UK.

## RESOURCES

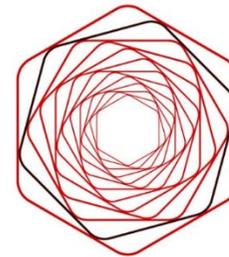
- 60 x 128 core Dell PowerEdge R6525 compute nodes with AMD EPYC 7702 dual 64-Core Processors (786GB RAM).
- 4 High memory nodes (2TB RAM).
- 32 x NVIDIA Tesla v100 GPUs.
- 2PB of lustre parallel file system for scratch storage.
- All compute nodes and storage are connected by EDR Infiniband fabric.
- Compute nodes run on Centos 7 operating system.

## EXEMPLAR RESEARCH AREAS

Neuroscience, Advanced Chemistry, Innovative Drug Delivery, Precision Medicine, Metabolomics, and Hydrogen Safety.

## CONTACT US

Website: <https://www.ni-hpc.ac.uk/>



**NI-HPC**  
NORTHERN IRELAND  
HIGH PERFORMANCE COMPUTING



The NI-HPC Tier 2 national supercomputing centre is the result of £2.1 million grant awarded by the Engineering and Physical Sciences Research Council (EPSRC) to a consortium of partners: Queen's University Belfast (QUB) and Ulster University (UU).