Viking Annual Report 2019/2020

HPC Management Committee

Summary

In January 2018, University Executive Board authorized £2.5M for the purchase and support of a new central high-performance computing facility for the University of York, supplied by Dell and supported by Alces Flight. This aimed to accelerate research in the university across all of the faculties and to provide a focus for developing interdisciplinary activities across the University. On February 14th 2019, Viking, as it was to be named, went live, providing 7000 cores, 42 TB of memory, 2.5 PB of storage, free at the point of delivery, to staff, postgraduates, and undergraduates across the university. This is a major university facility comparable to that provided by institutions much larger than York.

Since its launch, Viking has revolutionized the computational capability of the university. Supported by a team of 6 from IT Services and academic departments, the facility has supported 180 projects and 500 users. Despite the lag time between research and publication, 44 research outputs have been produced with support from Viking and this will rise substantially over the coming years.

The facility is currently well used, predominantly by the science departments but with some use across all faculties. An ongoing challenge will be engagement with the non-science faculties to ensure that they are able to take advantage of the transformative capacities of big data, machine learning, and computational methods in their disciplines. The recently created N8-CIR is helping here, with bespoke training in computational methods for Digital Humanities, Digital Health and other areas.

Although only halfway through the current lifecycle for Viking, attention must now turn to the future and the potential for the long term provision of high-performance computing at the University of York.

Statistics

- 500 registered users across 180 projects, spanning 17 departments
- 44 recorded research outputs, 86% of which would not have been possible without Viking
- 40 million CPU hours of work done over 7 million jobs
- Computing resources have been allocated to user jobs for 90% of the time in 2020
- 50% of Viking users are undergraduate/postgraduate taught students

Year Overview

Viking launched on February 14th 2019 with a successful event to mark its commission. Featuring academic and industrial speakers from a range of fields, the event marked a positive start to a transformative time for research at the University. In fact, prior to the official launch of the new computing facility, several University researchers had already enjoyed putting Viking through its paces, with the Biology department exploiting all 7000 of Viking's computational cores to great effect in reducing the time taken to complete genomic assemblies from several weeks on the University's previous computing facility, YARCC, to just over a day. Since becoming available to all University researchers, this kind of success has continued, with users from all faculties making use of the facility in their research endeavours.

Viking has seen steady growth in uptake over the course of its first year, with 180 projects and 500 users. Whilst computing facilities of this nature tend to be science-focused, Viking has managed to capture a user base across 17 academic departments in all three faculties.

Faculty of Sciences	Faculty of Social Sciences	Faculty of Arts and Humanities
Biology	**Economics and Related Studies**	**Archaeology**
Chemistry	**Education**	**Music**
Computer Science	**The York Management School**	**Philosophy**
Electronic Engineering	**Centre for Health Economics**	
Environment and Geography		
Health Sciences		
Hull York Medical School		
Mathematics		
Physics		
Psychology		

Figure 1: Departments with registered Viking projects

It is encouraging to see uptake from the Faculty of Arts and Humanities! There is scope for increased engagement with departments in both the Social Sciences and Arts and Humanities, which a growing research software engineering capacity in IT Services will help with.

Viking has seen a large volume of computational work done with over 40 million CPU hours of work completed across more than 7 million jobs. Viking resources have been in use with user jobs 89% of the time since the beginning of 2020.

Research Outputs

Through search of research output databases and user surveys, we have recorded a total of 44 research outputs (papers, articles, datasets, software etc.) that have been produced using Viking.

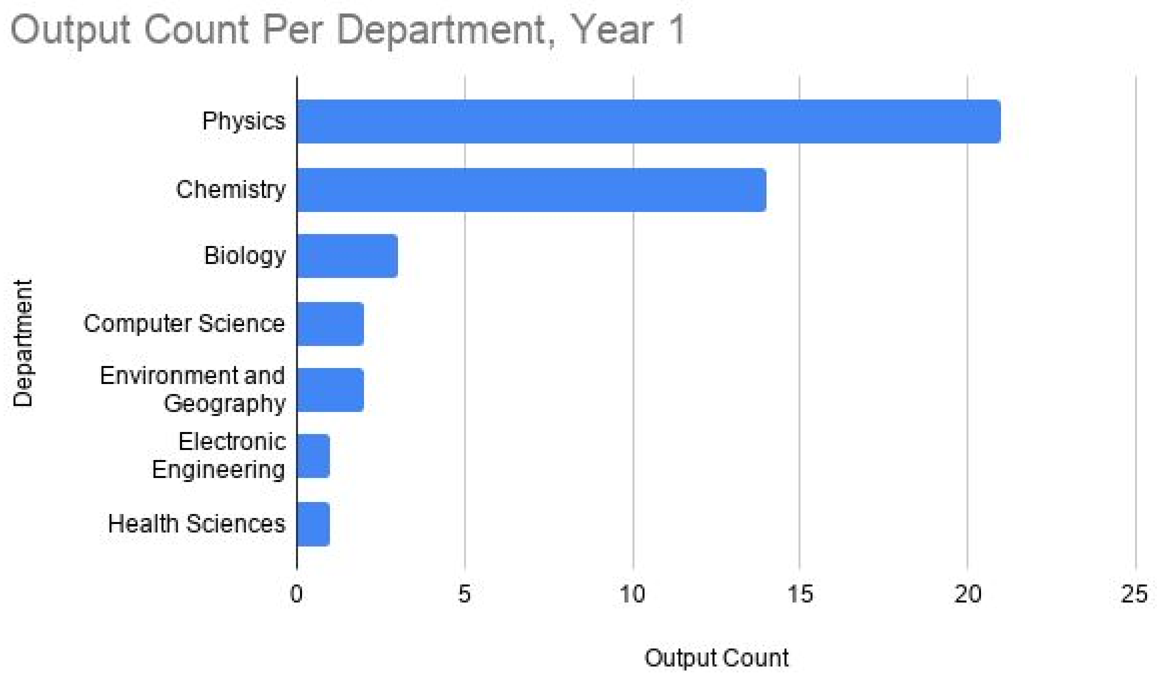

Figure 2: Research outputs which made use of Viking, counts by department

The above figure reflects the split of researchers on Viking, and the greater experience with HPC in the natural sciences. As we increase engagement with other Faculties and existing social science / arts and humanities projects complete, we expect this picture to change.

Additionally, 86% of survey respondents stated that without Viking, the research output would not have been possible to produce on University of York facilities.

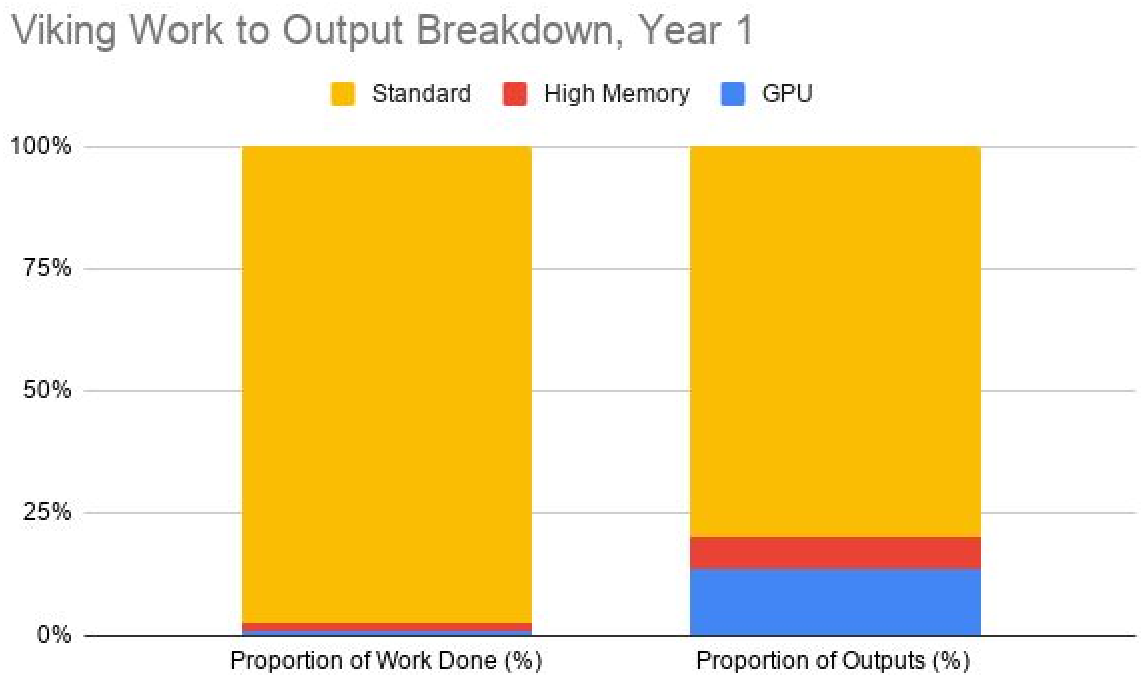

Figure 3: Proportion of work done by different types of servers in Viking and proportion of research outputs produced using each of these different types of server.

The standard nodes in Viking were responsible for 80% of the Viking research outputs. The specialist components (GPU, high memory) were responsible for a disproportionate amount of the recorded research outputs as these are facilities that are not available to researchers anywhere else on campus. All of these resources would typically incur cost on regional/national facilities or in the cloud but were provided free at the point of use to York researchers.

Case Studies

We have selected some highlights from work done on Viking throughout year 1 for case studies. Additionally, a full reference list for Viking research outputs can be found in 'Research Output References'.

The Unusual Skulls from the River Thames

Eleanor Joan Green, Ph.D. student in the Department of Archaeology, University of York and the Department of Life Sciences, Natural History Museum, London. This project is supported by the Arts & Humanities Research Council (AHRC) Collaborative Doctoral Award scheme (AH/N005015/1).

Figure 1: Three of the skulls recovered from the River Thames held in the River Thames collection at the Natural History Museum, London.

Over the past two centuries, archaeological finds have been, often inadvertently, recovered from various locations along the course of the River Thames, southeast England. Amongst these finds are ornate metalwork, preserved wood, and human remains, commonly skulls. The Natural History Museum, London, holds a collection of approximately 300 skulls recovered from the Thames in the past 200 years. The small number of skulls which have been radiocarbon dated thus far span a wide age range, from the Neolithic (from ~6000 years ago) to the Post-Medieval period.

As part of my PhD research, I have extracted the surviving DNA from 30 of the skulls from the Natural History Museum's collection. After death DNA is no longer maintained by the body, and begins to decay. As a result of this, when we investigate ancient DNA we are looking for small DNA fragments which have been subject to base changes. The DNA recovered from the skulls is being analysed using Viking, the University of York's new high performance computing cluster. Whilst working with ancient DNA, we retrieve both host and non-host DNA sequences. In the case of the Thames skulls, the host DNA is that from the individual human, the non-host DNA is mainly bacterial and derives from the depositional environment of the skulls. The human DNA will be used to create a genetic profile of each ancient individual (chromosonal sex and genetic affinities) by comparing the Thames skulls data to previously published genomic data derived from modern and ancient European populations. To identify the source of non-host DNA we use Viking to map against large databases of known DNA sequences. Identification of the non-host DNA will enable us to build up a more thorough picture of the history of this unusual collection.

Computational Magnetism Group

Sarah Jenkins, Mara Strungaru, Sergiu Ruta, Luke E Elliott, Daniel Meilak, Roy W Chantrell, and Richard F L Evans

Over the past year Viking has enabled the computational magnetism group in the Department of Physics at the University of York to conduct a wide variety of projects on fundamental and applied magnetic materials, with applications in novel cancer therapies, new magnetic memory technology and 2D magnetic materials. The new capability enabling large-scale simulations on hundreds of computer cores has enabled us to significantly increase the types of systems we can simulate and the quality of data that can be collected.

Some results of particular interest are those of two-dimensional magnetic materials such as CrI_3, essentially a magnetic equivalent of graphene. Until recently it was thought that such materials should not possess long-range magnetic order due to the Mermin-Wagner theorem, but our simulations have revealed rich magnetic behaviours not previously explored.

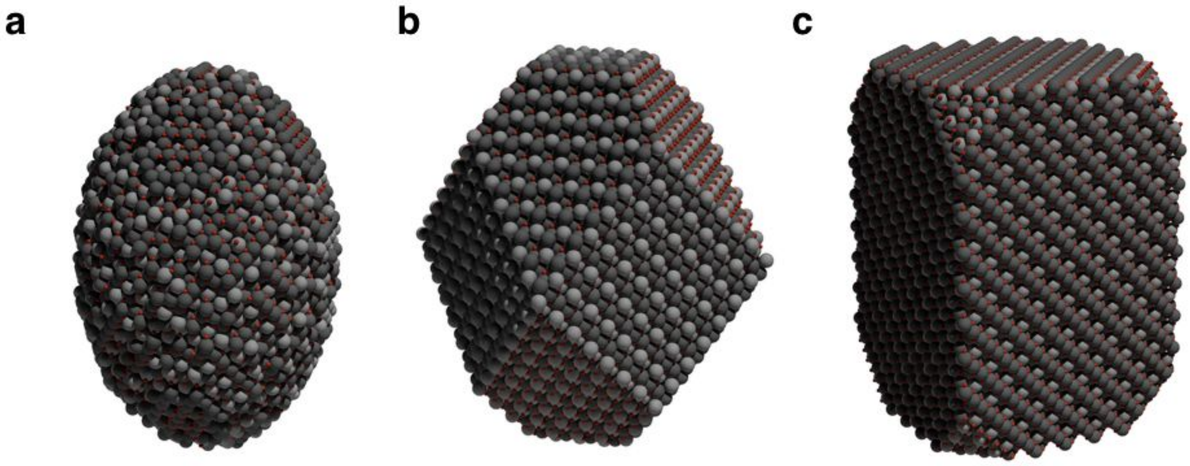

We have also performed simulations on Magnetic Random Access Memory (MRAM) devices, showing the complex size evolution of the magnetic stability as sub-50 nm length scales, as well as the origins of stray magnetic fields and their ability to form complex magnetic spin structures. This research has significant implications for the development of future devices with reduced dimensions, and could enable an industry-wide shift towards non-volatile memory technology.

Visualisation of different magnetic nanoparticles used for magnetic hyperthermia

Viking has also been used to explore magnetic shape anisotropy in small magnetite nanoparticles with applications in novel cancer therapy through magnetic hyperthermia. The shape anisotropy is often thought to be the dominant contribution to the heating which kills cancerous cells, but until now the detailed origins of the effect in nanoparticle-sized systems have proved illusive. Through state-of-the-art atomistic simulations performed on Viking we have been able to determine the contributions of shape and faceting to the magnetic shape anisotropy, which will help to accelerate the development of this promising cancer treatment.

We have used the large capacity of Viking to enable previously impossible simulations of the fundamental properties of IrMn, an industrially important antiferromagetic material with applications in magnetic data storage, spintronics, and neuromorphic computing. Using the power of thousands of computing cores, we were able to determine the complex cubic-like form and temperature dependence of the magnetic anisotropy for the first time. We were also able to simulate the excitations of IrMn coupled to a ferromagnetic material, showing a strong enhancement of the effective anisotropy due to the presence of interfacial spins. In a related paper, we determined the fundamental nature of exchange anisotropy, finding an unusual scaling with magnetization that has broad implications for the interpretation of experimental measurements. Another large scale study of the dynamics of ferromagnetic FePt used for heat assisted magnetic recording (HAMR) has explained complex experimental measurements that previously disagreed with theoretical estimations. Our findings will contribute to the development of next-generation magnetic recording technology enabling large increases in data storage density.

The provision of the Viking service has been revolutionary for our research, and has already enabled studies that would have been impossible. The work we are now able to produce is of the highest international standard, and will significantly enhance the profile of researchers from the University of York in the fields of spintronics and magnetism.

Modelling Surface Ozone Concentrations in Beijing

Peter Ivatt, Wolfson Atmospheric Chemistry Laboratories

Ozone is an important pollutant which degrades human health, inhibits plant growth and contributes to climate change. However, it is not directly emitted into the atmosphere but forms from complex chemistry occurring within the atmosphere. Computer simulations which attempt to replicate the chemistry, meteorology, emissions, and other processes in the atmosphere are used to try to understand this chemistry and to evaluate policy decisions to reduce ozone's concentration.

Due to rapid economic growth, recent attention has focused on the sources of ozone in China and the interplay of different compounds emitted into the atmosphere. We have focused on data collected in Beijing in the summer of 2017. Using the open-source GEOS-Chem atmospheric chemistry transport model (www.geos-chem.org) we have simulated the observations collected during this campaign.

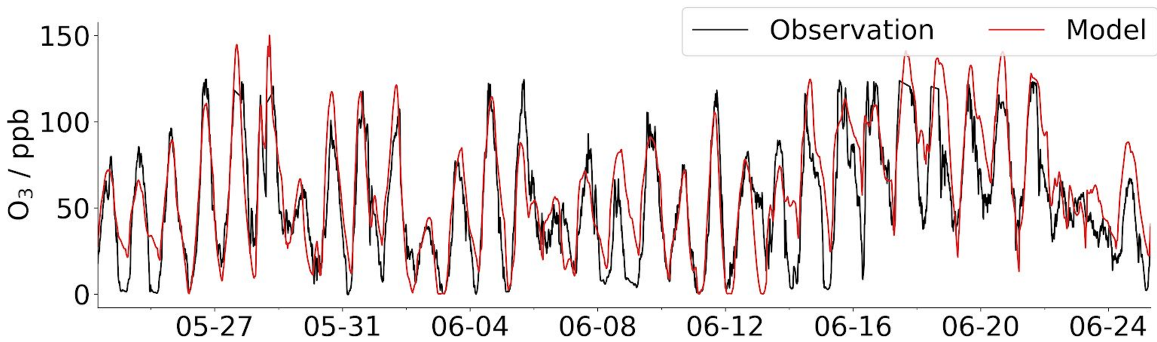

Figure 1. Ozone concentration measured in Beijing in May and June 2017. Black line indicates observed concentrations and red line indicates that simulated concentrations above 50 ppbv for 8 hours would break WHO standards.

Conventionally, once a 'reasonable' base model simulation has been developed a relatively small number of studies are conducted to explore the model sensitivity. This is often done in an ad-hoc manner. With the increase in computational capabilities offered by Viking, we have been able to take a different approach.

With Viking, we have been able to run hundreds of almost identical simulations, systematically varying one parameter in turn. With our previous cluster, it took a week to run 4 simulations. With Viking, this has become a week to run around 200 simulations. This increase in computational throughput allows us to better explore the complex chemistry and emissions that underlie a problem such as air quality in China and so provide increased insights into this important societal problem.

CryoEM structure of the Nipah virus nucleocapsid assembly

Huw Jenkins, York Structural Biology Laboratory.

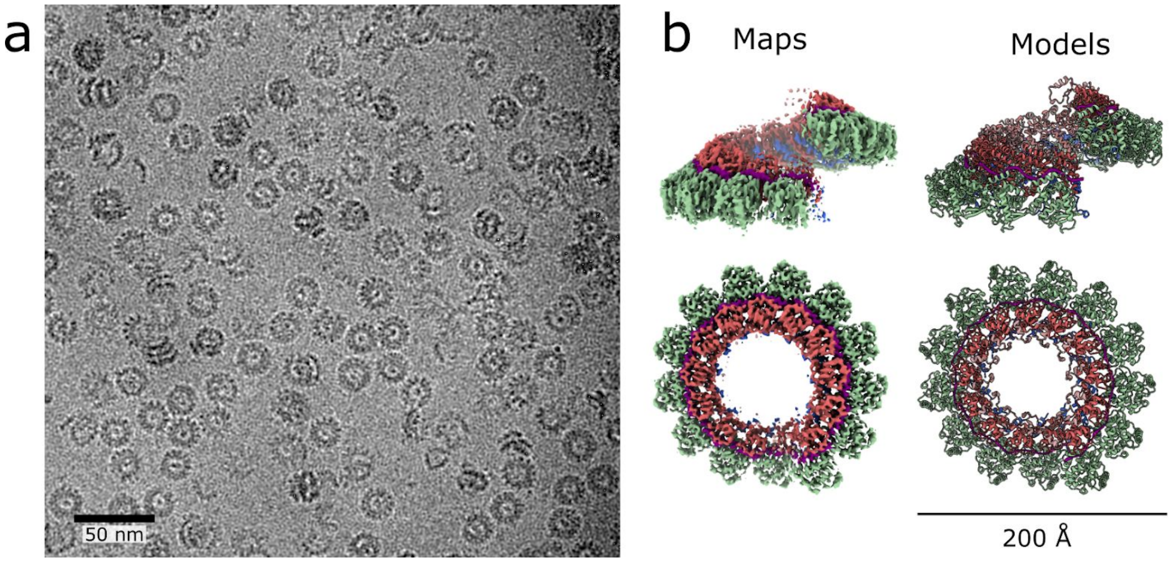

Figure 1: Micrograph and final reconstruction. **a)** representative micrograph scale bar is 50 nm). **b)** 3.6 Å reconstruction (left) and atomic model (right).

Nipah virus (NiV) is an emerging, highly pathogenic virus causing acute respiratory illness and fatal encephalitis. There is no vaccine or approved treatment available. Since the initial NiV outbreak in Malaysia in 1998, there have been sporadic outbreaks in Bangladesh and India almost every year with case-fatality rates of around 75% and confirmed human to human transmission. The World Health Organisation has classed NiV as a priority disease alongside Ebola, Zika and SARS due to the pandemic potential of this virus.

The RNA genome of NiV is wrapped by the nucleocapsid protein to form a helical nucleocapsid. This not only protects the viral genetic material from degradation by nucleases in the host cell but also regulates viral replication. Therefore understanding of the mechanism by which the nucleocapsid protein binds to RNA could play a role in development of antiviral therapeutics. The NiV nucleocapsid protein can be recombinantly produced in bacteria where it binds to the bacterial RNA.

In order to determine the structural basis of nucleocapsid protein:RNA interaction, purified ribonuclear protein complexes were imaged by electron cryomicroscopy (cryoEM) at the UK national electron bio-imaging centre (eBIC). An example micrograph is shown in Figure 1a. A variety of spiral shaped assemblies are present in the sample resulting from the variation in length of RNA bound. In order to produce a high resolution cryoEM structure of the complex many thousands of images of the protein:RNA complex with the same conformation must be aligned and averaged, therefore, this heterogeneity in the sample presents a major challenge.

Fortunately, recent developments in cryoEM reconstruction enable “computational purification” where images of complexes in different conformations can be separated. The aim is to identify a subset of images where a single conformation of the complex has been imaged in a variety of orientations. As the number of conformations present in the sample is unknown, multiple attempts must be made and so the process is computationally very demanding. It was, therefore, essential for this project to be able to make extensive use of the resources available on Viking. In the RELION software the computation is parallelized down to the level of individual particles and so the work can efficiently make use of the resources on Viking. Even so for a classification run using ~150 CPUs the typical wall-clock time required was 24 hours. Over 350 classification runs were required to optimise parameterization and then characterise the heterogeneity in the sample (Fig 2), to finally enable identification of a subset of particles that could be used to generate a sub-4Å resolution reconstruction (Fig. 1b). This resolution was sufficient to build an atomic model of the nucleocapsid protein bound to RNA (Fig. 1b). The model enabled both the structural basis for RNA binding by the nucleocapsid protein to be determined and the protein:protein interactions involved in helical assembly of the nucleocapsid to be defined. A pre-print describing this work has been submitted (CryoEM structure of the Nipah virus nucleocapsid assembly - doi: https://doi.org/10.1101/2020.01.20.912261)

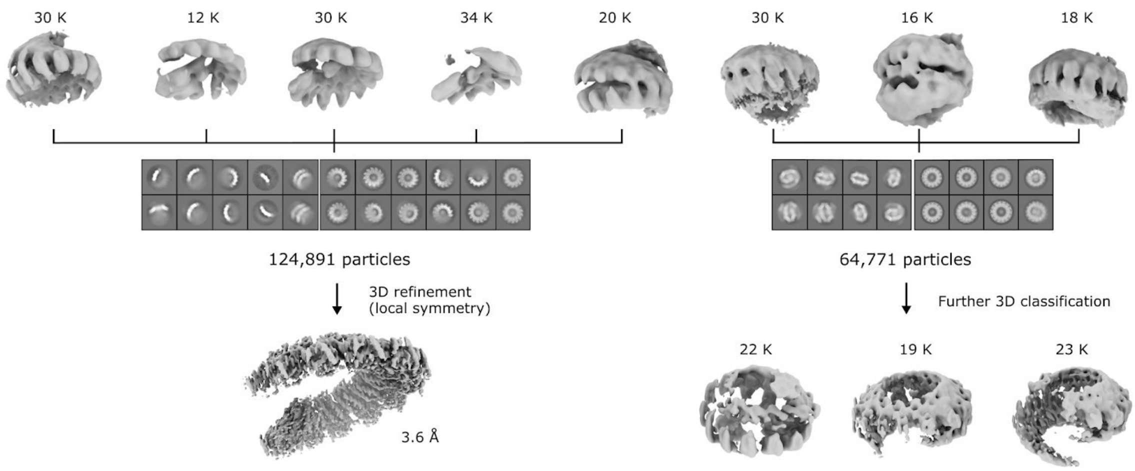

Figure 2: 3D classification workflow.

Targeted sequencing in DLBCL, molecular subtypes, and outcomes: a Haematological Malignancy Research Network report

Stuart Lacy, Health Sciences.

Diffuse large B-cell lymphoma (DLBCL) is an aggressive cancer and the most common form of non-Hodgkin lymphoma. Recently, research attention has focused on classifying genetic variants of this cancer in order to identify potential targets for novel treatments; two recent whole exome sequencing studies proposed partially overlapping classification systems.

The Haematological Malignancy Research Network (HMRN) provides a registry of all new haematological malignancy diagnoses (leukemias, lymphomas and myelomas, and related blood disorders) from a catchment population of ~4 million. Using a representative population-based cohort of 928 DLBCL patients from HMRN, tumour DNA was sequenced using a targeted gene panel. A computationally-intensive machine learning algorithm was applied, confirming the existence of six molecular subtypes of DLBCL, providing evidence that genomic tests have prognostic significance in non-selected DLBCL patients.

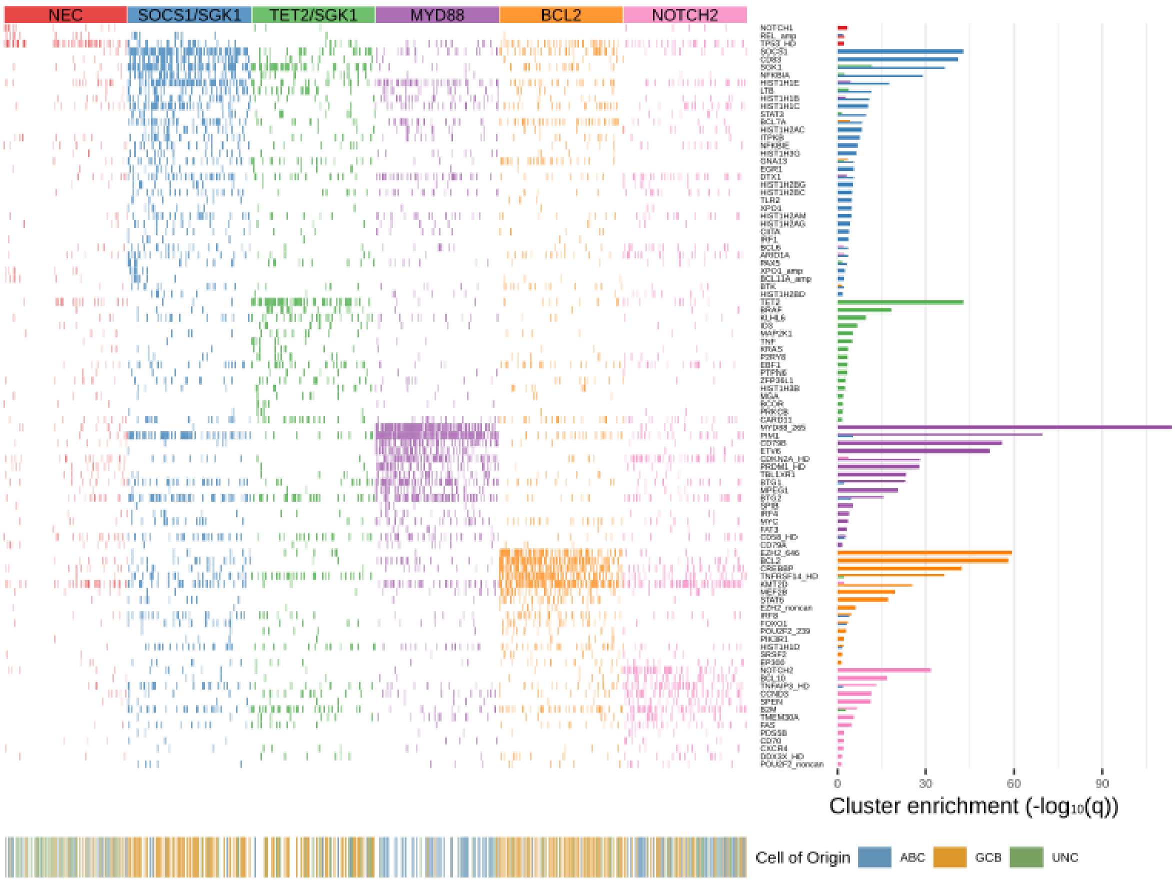

A consensus-clustering method was employed to quantify the strength of these six subtypes. The original clustering algorithm was refitted on a large number (1,000) of bootstrapped datasets (1,000) and the proportion of pairs of samples that were assigned to the same cluster was calculated; the more robust the cluster, the more frequently its members would be grouped together. Fitting the mixture-model takes around 20 minutes, meaning the whole process would take 33 hours, or over 18 days, if run sequentially. Fortunately, the Viking HPC facility at York enabled this to be run in parallel, meaning results could be obtained in a matter of hours. The plot below summarises the results, showing that the MYD88 group was the most robustly defined, and NOTCH2 the least.

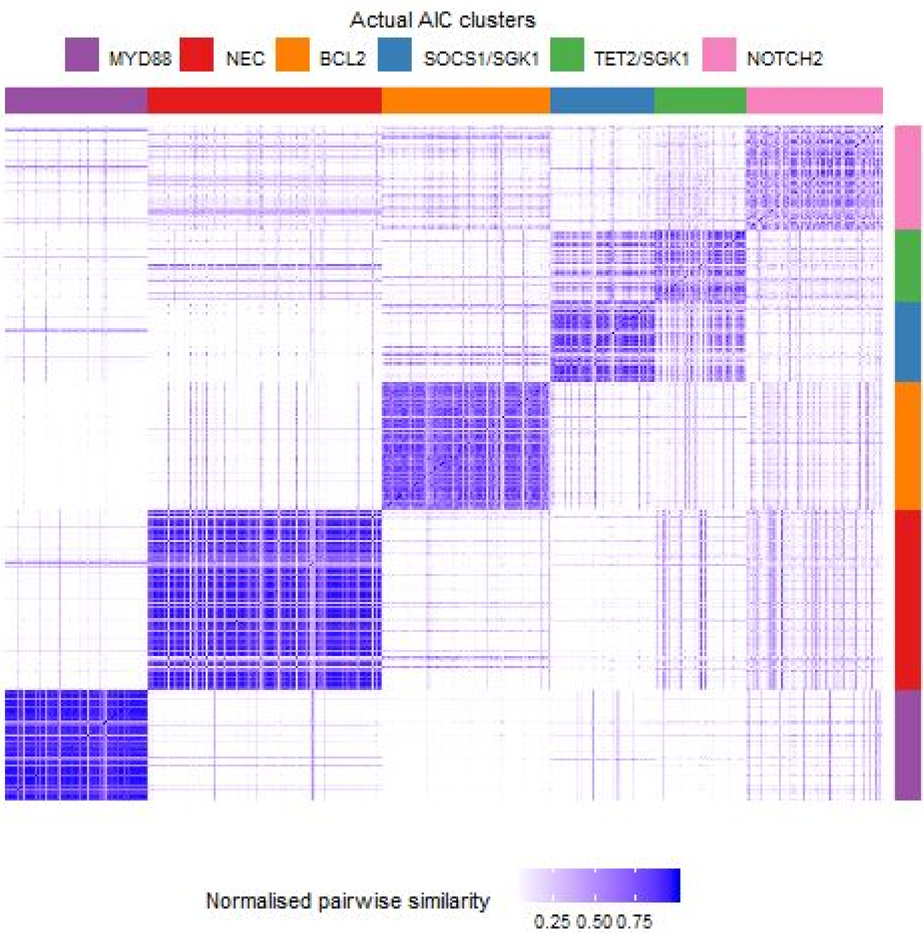

Events, Support, and Training

Events

Viking's commission began with a successful launch event, filling the Bowland Auditorium and drawing speakers from a range of institutions. Professor Sophia Ananiadou (Director of The National Centre for Text Mining), Professor Carla Molteni (Physics, King's College London), and Conrad Wolfram (CEO, Wolfram Research Europe) presented on text mining, computational simulations of materials, and new developments in mathematics education through computing, presentations which were well-received by attendees and prompted engaging discussion.

Since the launch, the Viking team has had a presence at several events, including YorkTalks and subject specific conferences, and has run a successful showcase event, bringing together researchers from across faculties to present their computational work on Viking. This event offered researchers from many departments the opportunity to share knowledge and experience on computational research topics, and has led to valuable connections being made between the Viking support team and the University research community.

Support

Over the past year, the Research and High Performance Computing Support Team has grown from its original size of 4 staff working on rotation to provide generic support, to 6 staff providing both dedicated first and second-line support to all Viking users and specialised research computing support to several departments in the Faculty of Sciences. This growth in staff has been funded using a mixed model, with funding from IT Services, academic departments, and research projects.

Growth in support capability has resulted in an improved user experience on Viking, allowing minor issues and queries to be resolved swiftly and leaving room for deep-dive support on the kinds of issues that facilitate new research or more quickly delivering on research outputs. The 2019 IT Survey turned up several positive items of feedback for Viking, and the additional capacity in the team has also allowed us to work on longer term plans based on feedback for improving the service.

IT Services have also been able to support research activities on cloud computing infrastructure - as part of the Viking procurement, the University received $10,000 of Amazon Web Services credits, which have been allocated to nine research projects. These credits allow for the exploration of technologies that are unsupported on Viking, and provide insight into technologies to consider for research computing procurement.

Training

IT Services have a number of relevant training courses for Viking in development. We are currently collaborating with Teaching & Learning to deliver an Introduction to the Linux Shell, foundational knowledge for productive working on Viking. Following the introduction to Linux shell, a generally-targeted Introduction to Viking will run, covering the tools and best practices for getting value from the facility, and where requested, we are happy to develop versions of this training that are customised for the needs of departments across the faculties. IT Services also maintains a curated list of online training resources for researchers, which was recently referenced by UCL as a ‘go to’ list of resources for online research computing learning.

In addition to group training, individual consultations with a member of research computing support have been made good use of - University researchers have been able to talk through their computational ideas and problems with an experienced member of staff, resulting in new ways of working and improvements on existing computational workflows.

Finally, IT Services have been working with the University’s Coding Club, run by members of staff from the Physics department and IT Services, to provide informal support for computational research. Many of these interactions have resulted in an improvement of user experience on Viking!

Regional and National Links

This is an exciting time for regional and national research computing at the University. The N8 Centre of Excellence in Computationally Intensive Research (N8CIR) has recently been awarded £3.1M from EPSRC, supplemented by £5.3M from N8 member institutions, to establish a new research computing facility in the North of England. The **N**orthern **I**ntensive **C**omputing **E**nvironment, NICE, has been designed to match the architecture of some of the fastest supercomputers in the world, and will facilitate research from across the University faculties - particularly that which can benefit from GPU computing. You can find more information about this new facility on the N8 CIR website. IT Services will be able to support the use of NICE when it becomes available for use, and look forward to the potential new discoveries/improvements it could facilitate.

Not only is there a new regional facility on the way, the national supercomputing facility ARCHER is currently undergoing a teardown and rebuild into a completely new facility (ARCHER2), having received £79M of government funding in early 2019. With potential for impressive performance improvements (nearly 10x in many cases), as demonstrated by the system benchmarking results, several groups at the University who have made use of the current ARCHER facility will be looking forward to a significantly improved system. More information about the new ARCHER2 system can be found on the ARCHER website.

We see the establishment of these two new facilities as a great opportunity to support University researchers in developing their computational work to incorporate new technologies and run at a larger scale than Viking can reasonably facilitate. Viking is therefore an important stepping stone for researchers looking to make use of these facilities, and we will be looking for ways in which we can support the jump up.

Plans for the Future

Viking is now stable, with the initial teething troubles dealt with. Most users are able to exploit the system without support. This affords the Viking support team an opportunity to develop the service further. Projects here include:

- Redesigning user documentation and developing specialised support for individual departments to encourage more uptake within the University
- Developing a status monitoring system to provide Viking users with information about the current usage of Viking
- Developing our support for reproducible research practices through tools which allow work to be straightforwardly transferred between computing facilities, allowing groups to move from local to regional and ultimately national computing systems
- Work closely with other elements in IT Services to provide on-demand access to cloud computing resources, giving researchers the ability to flexibly move beyond the limitations of Viking

Research Output References

1. R. J. Hall et al., ‘Simulating the evolutionary trajectories of metabolic pathways for insect symbionts in the Sodalis genus’. Cold Spring Harbor Laboratory, Oct. 25, 2019 [Online]. Available https://dx.doi.org/10.1101/819946.

2. V. Daza-Cajigal et al., ‘Loss of Janus Associated Kinase 1 Alters Urothelial Cell Function and Facilitates the Development of Bladder Cancer’, Frontiers in Immunology, vol. 10, Sep. 2019 [Online]. Available https://dx.doi.org/10.3389/fimmu.2019.02065.

3. R. McNeill et al., ‘Specificity of the Metallothionein-1 Response by Cadmium-Exposed Normal Human Urothelial Cells’, International Journal of Molecular Sciences, vol. 20, no. 6, p. 1344, Mar. 2019 [Online]. Available https://dx.doi.org/10.3390/ijms20061344.

4. H. E. Ho et al., ‘Visible-light-induced intramolecular charge transfer in the radical spirocyclisation of indole-tethered ynones’, Chemical Science, vol. 11, no. 5, pp. 1353–1360, 2020 [Online]. Available https://dx.doi.org/10.1039/C9SC05311E.

5. D.-S. Ker et al., ‘CryoEM structure of the Nipah virus nucleocapsid assembly’. Cold Spring Harbor Laboratory, Jan. 20, 2020 [Online]. Available http://dx.doi.org/10.1101/2020.01.20.912261.

6. R. J. Pound et al., ‘Influences of oceanic ozone deposition on tropospheric photochemistry’. Copernicus GmbH, Nov. 20, 2019 [Online]. Available http://dx.doi.org/10.5194/acp-2019-1043.

7. N. G. S. McGregor et al., ‘Rational Design of Mechanism-Based Inhibitors and Activity-Based Probes for the Identification of Retaining α-l-Arabinofuranosidases’, Journal of the American Chemical Society, vol. 142, no. 10, pp. 4648–4662, Feb. 2020 [Online]. Available https://dx.doi.org/10.1021/jacs.9b11351.

8. N. G. K. Wong et al., ‘Direct Observation of Photochemical Free Radical Production from the Sunscreen 2-Phenylbenzimidazole-5-Sulfonic Acid via Laser-Interfaced Mass Spectrometry’, ChemPhotoChem, vol. 3, no. 12, pp. 1231–1237, Aug. 2019 [Online]. Available https://dx.doi.org/10.1002/cptc.201900149.

9. R. J. Chance et al., ‘Global sea-surface iodide observations, 1967–2018’, Scientific Data, vol. 6, no. 1, Nov. 2019 [Online]. Available https://dx.doi.org/10.1038/s41597-019-0288-y.

10. T. Sherwen et al., ‘A machine learning based global sea-surface iodide distribution’. Copernicus GmbH, Mar. 26, 2019 [Online]. Available https://dx.doi.org/10.5194/essd-2019-40.

11. A. Badia et al., ‘Importance of reactive halogens in the tropical marine atmosphere: a regional modelling study using WRF-Chem’, Atmospheric Chemistry and Physics, vol. 19, no. 5, pp. 3161–3189, Mar. 2019 [Online]. Available https://dx.doi.org/10.5194/acp-19-3161-2019.

12. S. Inamdar et al., ‘Estimation of Reactive Inorganic Iodine Fluxes in the Indian and Southern Ocean Marine Boundary Layer’. Copernicus GmbH, Feb. 04, 2020 [Online]. Available http://dx.doi.org/10.5194/acp-2019-1052.

13. L. Zhu et al., ‘Effect of sea salt aerosol on tropospheric bromine chemistry’, Atmospheric Chemistry and Physics, vol. 19, no. 9, pp. 6497–6507, May 2019 [Online]. Available https://dx.doi.org/10.5194/acp-19-6497-2019.

14. X. Wang et al., 'The role of chlorine in global tropospheric chemistry', Atmospheric Chemistry and Physics, vol. 19, no. 6, pp. 3981–4003, Mar. 2019 [Online]. Available https://dx.doi.org/10.5194/acp-19-3981-2019.

15. T. I. G.-C. U. Community, 'geoschem/geos-chem: GEOS-Chem 12.6.0'. Zenodo, Oct. 18, 2019 [Online]. Availablehttps://zenodo.org/record/3507501.

16. T. I. G.-C. U. Community, 'geoschem/geos-chem: GEOS-Chem 12.7.0'. Zenodo, Feb. 03, 2020 [Online]. Available https://zenodo.org/record/3634864.

17. T. Sherwen, 'tsherwen/sparse2spatial: sparse2spatial v.0.1.1 - Predictions for iodide, CH2Br2 and CHBr3'. Zenodo, Aug. 15, 2019 [Online]. Available https://zenodo.org/record/3369212.

18. E. Alharbi et al., 'Comparison of automated crystallographic model-building pipelines', Acta Crystallographica Section D Structural Biology, vol. 75, no. 12, pp. 1119–1128, Nov. 2019 [Online]. Available https://dx.doi.org/10.1107/s2059798319014918.

19. B. Muller and W. Smith, 'A Hierarchical Loss for Semantic Segmentation', in Proceedings of the 15th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, 2020 [Online]. Available https://dx.doi.org/10.5220/0008946002600267.

20. H. Xie et al., 'Numerical and Analytical Analysis of Stochastic Electromagnetic Fields Coupling to a Printed Circuit Board Trace', IEEE Transactions on Electromagnetic Compatibility, pp. 1–8, 2019 [Online]. Available https://dx.doi.org/10.1109/TEMC.2019.2954303.

21. A. Baker et al., 'Modelling the ecological impacts of tidal energy barrages'. Center for Open Science, Jan. 09, 2020 [Online]. Available https://dx.doi.org/10.31223/osf.io/vapmu.

22. J. Hill, 'Resolving the tsunami wave: interpreting palaeotsunami deposits by integrating numerical modelling and sedimentology'. figshare, 2019 [Online]. Available https://figshare.com/articles/Resolving_the_tsunami_wave_interpreting_palaeotsunami_deposits_by_integrating_numerical_modelling_and_sedimentology/9771917/3.

23. S. E. Lacy et al., 'Targeted sequencing in DLBCL, molecular subtypes, and outcomes: a Haematological Malignancy Research Network report', Blood, Mar. 2020 [Online]. Available https://dx.doi.org/10.1182/blood.2019003535.

24. S. Jenkins et al., 'Magnetic stray fields in nanoscale magnetic tunnel junctions', Journal of Physics D: Applied Physics, vol. 53, no. 4, p. 44001, Nov. 2019 [Online]. Available https://dx.doi.org/10.1088/1361-6463/ab4fbf.

25. S. Jenkins et al., 'Magnetic anisotropy of the noncollinear antiferromagnet IrMn3', Physical Review B, vol. 100, no. 22, Dec. 2019 [Online]. Available https://dx.doi.org/10.1103/PhysRevB.100.220405.

26. R. Moreno et al., 'The role of faceting and elongation on the magnetic anisotropy of magnetite Fe3O4 nanocrystals', Scientific Reports, vol. 10, no. 1, Feb. 2020 [Online]. Available https://dx.doi.org/10.1038/s41598-020-58976-7.

27. B. Parks et al., 'Magnetoresistance Dynamics in Superparamagnetic Co−Fe−B Nanodots', Physical Review Applied, vol. 13, no. 1, Jan. 2020 [Online]. Available https://dx.doi.org/10.1103/PhysRevApplied.13.014063.

28. S. Jenkins and R. F. L. Evans, ‘Scalable space and time hierarchical dipole- dipole interactions in the VAMPIRE code’, Zenodo, 2020 [Online]. Available https://zenodo.org/record/3669966.

29. S. M. João et al., ‘KITE: high-performance accurate modelling of electronic structure and response functions of large molecules, disordered crystals and heterostructures’, Royal Society Open Science, vol. 7, no. 2, p. 191809, Feb. 2020 [Online]. Available https://dx.doi.org/10.1098/rsos.191809.

30. S. M. João et al., ‘KITE: high-performance quantum transport software’. Zenodo, Jun. 13, 2019 [Online]. Available https://zenodo.org/record/3245011.

31. J. W. Shepherd et al., ‘The emergence of sequence-dependent structural motifs in stretched, torsionally constrained DNA’, Nucleic Acids Research, vol. 48, no. 4, pp. 1748–1763, Jan. 2020 [Online]. Available https://dx.doi.org/10.1093/nar/gkz1227.

32. E. Iacocca et al., ‘Spin-current-mediated rapid magnon localisation and coalescence after ultrafast optical pumping of ferrimagnetic alloys’, Nature Communications, vol. 10, no. 1, Apr. 2019 [Online]. Available https://dx.doi.org/10.1038/s41467-019-09577-0.

33. A. Ceballos et al., ‘Role of Element-Specific Damping on the Ultrafast, Helicity-Independent All-Optical Switching Dynamics in Amorphous (Gd,TB)Co Thin Films, preprint, Nov. 2019 [Online]. Available https://arxiv.org/abs/1911.09803 .

34. M. Strungaru et al., ‘Model of Damping and Anisotropy at Elevated Temperatures: Application to Granular FePt Films’, preprint, Feb. 2020 [Online]. Available https://arxiv.org/abs/2002.02865.

35. J. A. Quirk et al., ‘Electronic Properties of {112} and {110} Twin Boundaries in Anatase TiO 2’, Advanced Theory and Simulations, vol. 2, no. 12, p. 1900157, Oct. 2019 [Online]. Available https://dx.doi.org/10.1002/adts.201900157.

36. S. Petzold et al., ‘Forming-Free Grain Boundary Engineered Hafnium Oxide Resistive Random Access Memory Devices’, Advanced Electronic Materials, vol. 5, no. 10, p. 1900484, Aug. 2019 [Online]. Available https://dx.doi.org/10.1002/aelm.201900484.

37. K. P. McKenna, ‘Forming-Free Grain Boundary Engineered Hafnium Oxide Resistive Random Access Memory Devices’. University of York, 2019 [Dataset, Online]. Available https://pure.york.ac.uk/portal/en/datasets/formingfree-grain-boundary-engineered-hafnium-oxide-resistive-random-access-memory-devices(ddd777f1-0e57-41dc-b843-449d7268e617).html.

38. S.-H. Hung and K. McKenna, ‘First-Principles Investigation of the Structure and Properties of Au Nanoparticles Supported on ZnO’, The Journal of Physical Chemistry C, vol. 123, no. 34, pp. 21185–21194, Jul. 2019 [Online]. Available https://dx.doi.org/10.1021/acs.jpcc.9b02639.

39. S.-H. Hung, ‘First principles investigation of the structure and properties of Au nanoparticles supported on ZnO’. University of York, 2019 [Dataset, Online]. Available https://pure.york.ac.uk/portal/en/datasets/first-principles-investigation-of-the-structure-and-properties-of-au-nanoparticles-supported-on-zno(e4abc92f-2363-4b8e-b5ca-f77c0e9e5967).html.

40. A. C. M. Padilha and K. P. McKenna, ‘First principles investigation of Y2O3-doped HfO2’, Journal of Applied Physics, vol. 126, no. 8, p. 84105, Aug. 2019 [Online]. Available https://dx.doi.org/10.1063/1.5110669.

41. A. C. C. Padilha, ‘First principles investigation of Y2O3-doped HfO2’. University of York, 2019 [Dataset, Online]. Available

https://pure.york.ac.uk/portal/en/datasets/first-principles-investigation-of-y2o3doped-hfo2(6af02367-39f8-4577-bcd9-4bd7bd0f9e77).html.

42. J. J. Carey and K. P. McKenna, ‘Does Polaronic Self-Trapping Occur at Anatase TiO2 Surfaces?’, The Journal of Physical Chemistry C, vol. 122, no. 48, pp. 27540–27553, Nov. 2018 [Online]. Available https://dx.doi.org/10.1021/acs.jpcc.8b09437.

43. J. J. Carey and K. P. McKenna, ‘Screening Doping Strategies To Mitigate Electron Trapping at Anatase TiO2 Surfaces’, The Journal of Physical Chemistry C, vol. 123, no. 36, pp. 22358–22367, Aug. 2019 [Online]. Available https://dx.doi.org/10.1021/acs.jpcc.9b05840.

44. C.-J. Tong and K. P. McKenna, ‘Passivating Grain Boundaries in Polycrystalline CdTe’, The Journal of Physical Chemistry C, vol. 123, no. 39, pp. 23882–23889, Sep. 2019 [Online]. Available https://dx.doi.org/10.1021/acs.jpcc.9b08373.