



Proceedings

Eighteenth ACC Cyfronet AGH HPC Users' Conference



Editors: Marek Magryś
Marian Bubak
Robert Pająk
Andrzej Zemła

Zakopane, 13-15 April 2026

Proceedings

Eighteenth
ACC Cyfronet AGH
HPC Users'
Conference

Zakopane
13-15 April 2026

Editors: Marek Magryś
Marian Bubak
Robert Pająk
Andrzej Zemła

Published in April 2026

by Academic Computer Centre Cyfronet AGH
Nawojki 11, 30-950 Kraków, P.O.Box 6, Poland

© The Authors mentioned in the Table of Contents

All rights reserved. This book or part thereof, may not be reproduced in any form or by any means electronic or mechanical, including photocopying, recording or any information storage and retrieval system now known or to be invented, without written permission of the Authors and Publisher.

ISBN 978-83-61433-51-4

Cover design and book typesetting by Joanna Kasina.

Organization

KU KDM 2026 was organized by the Academic Computer Centre Cyfronet AGH, Nawojki 11, 30-950 Kraków, Poland.

Organizing Committee

Marek Magryś, Joanna Kasina, Magdalena Maryańska, Kamil Mucha, Robert Pająk

KU KDM'26 Program Committee

Andrzej Zemła, PhD (chairman)	Academic Computer Centre Cyfronet AGH
Marian Bubak, PhD	Sano Centre for Computational Medicine / Academic Computer Centre Cyfronet AGH
Izabela Czekaj, Prof.	Cracow University of Technology
Andrzej Eilmes, Prof.	Jagiellonian University
Marek Gorgoń, Prof.	AGH University of Krakow
Agnieszka Janiuk, Prof.	Center for Theoretical Physics PAS
Jacek Kitowski, Prof.	AGH University of Krakow / Academic Computer Centre Cyfronet AGH
Jacek Korchowicz, Prof.	Jagiellonian University
Paweł Łabaj, PhD	Jagiellonian University
Marek Magryś	Academic Computer Centre Cyfronet AGH
Jacek Niemiec, Prof.	Henryk Niewodniczański Institute of Nuclear Physics PAS
Irena Roterman-Konieczna, Prof.	Jagiellonian University Medical College
Bartłomiej Szafran, Prof.	AGH University of Krakow
Maciej Szaleniec, Prof.	Jerzy Haber Institute of Catalysis and Surface Chemistry PAS
Michał Tomza, PhD	University of Warsaw
Tadeusz Uhl, Prof.	AGH University of Krakow
Kazimierz Wiatr, Prof.	Academic Computer Centre Cyfronet AGH

Table of Contents

Invited Talks	
<hr/>	
The Stubborn Old Protein Folding Problem in the Era of ExaFlops, Qbits and AI	9
<i>O. Zimmermann</i>	
Towards European Infrastructure for In Silico Medicine	10
<i>M. Bubak</i>	
Current Activities and Further Development of the National Supercomputing Infrastructure of the Czech Republic	11
<i>V. Vondrák</i>	
<hr/>	
Contributed Papers	
<hr/>	
HARM-COMBO: Development and Performance Tests of the Hybrid (MPI+OpenMP) GRMHD Code	13
<i>E. Klepczarek, A. Janiuk, I. Janiuk, P. Płonka</i>	
AI-Based Job Runtime Prediction for the ALICE Computing Grid at CERN	15
<i>T. Lelek, B. Baliś, M. Kurdziel, M. Zasada</i>	
Benchmarking a Heterogeneous Metagenomics Pipeline on Ares and Helios Supercomputers	17
<i>P. Kica, K. Zielińska, M. Orzechowski, K. Zajęc, T. Kościółek, M. Malawski</i>	
HPC in Global Quantification and Personalization of Regulated Cardiovascular Model - a Cohort Study	19
<i>K. Tlalka, I. Halliday, A. Narracott, M. Malawski</i>	
Automating CFD Workflows on HPC with PyFluent: from Batch Simulations to Large-Scale Study	21
<i>K. Zajęc, M. Otta, M. Malawski, A. Narracott</i>	
PubDB: an Institutional Publication Database for KPI Monitoring and Visibility	23
<i>T. Zhyhulin, A. Nowak, J. Meizner, P. Nowakowski</i>	
Evaluating the Stability of Deep Neural Networks for Brain Microstructural Parameter Estimation under Varying Diffusion Gradient Schemes	25
<i>D. Ciupek, M. Malawski, T. Pięciak</i>	
High-Throughput Statistical Shape Modelling of Venous Anatomy Using Containerised Deformetrica on Ares Supercomputer	27
<i>M. Otta, P. Kica, M. Malawski</i>	

Using High Performance Computing for the Application of PINNs in Cardiovascular Models <i>E. Gómez del Pozo, K. Tlalka, K. Zajac, M. Malawski</i>	29
CR39 Track Detection and Segmentation <i>M. Rejdak, S. Rybski, L. Grzanka, J. B. Christensen</i>	31
Radiative GRMHD Simulations of Puffy Accretion Discs: Numerical Modelling of Sub-Eddington Black Hole Accretion <i>D. Lančová, M. Wielgus, M. Abramowicz, A. Różańska, G. Török, W. Kluźniak</i>	33
Sovereign AI Chat Platform with RAG on PLGrid: Data Privacy, Polish Language Support, and Open-Source HPC Deployment <i>M. Makuch, T. Balawajder, M. Cholewa, P. Dąbrowska-Wierzbicka, Ł. Flis, J. Kocot, M. Leśniak, S. Mazurek, A. Mytnik, H. Siejkowski, J. Sto, B. Wend</i>	35
Smart Data Management and ML-based Workflow Prediction for Cognitive Compute Continuum in SPICE <i>M. Orzechowski, P. Kica, M. Zasada, Ł. Opiola, B. Kryza, J. Liput, A. Raczek, J. Karczewski, W. Szmelich, D. Nikolow, Ł. Dutka, B. Baliś, R. G. Słota, J. Kitowski</i>	37
Selecting Knowledge Representation Tools for a Cloud-Native Context <i>S. Głomski, P. Zajdel, J. Kosińska</i>	39
Towards Event-Based In-Cabin Sensing: a Synchronized Dual-Camera Dataset Pipeline <i>M. Kowalczyk, M. Długosz, K. Jeziorek, M. Komorkiewicz, T. Kryjak, D. Marchewka, M. Szelest, P. Wzorek, P. Skruch</i>	41
Simulation of Texture Evolution in Rolled Cu with Neural Network Support <i>B. Sulkowski</i>	43
Sampling Spin-Glass Systems with Transformers <i>A. Stefański, D. Zapolski, P. Białas, P. Korcyl, T. Stebel</i>	45
Transformer as an Efficient Sampler for the Ising Model <i>D. Zapolski, P. Białas, P. Korcyl, T. Stebel</i>	47
A More Noise Resilient Search Space Method for the Quantum Approximate Optimization Algorithm <i>A. Wolk, K. Capala, K. Rycerz</i>	49
Kinetic Simulations of Nonresonant Instability in the Precursor of Young Supernova Remnant Shock Containing Positrons <i>O. Kobzar</i>	51
Towards Fast, Machine-Learning-Based Calorimeter Simulations in the ALICE Experiment at CERN <i>E. Majerz, Ł. Dubiel, J. Otwinowski, W. Dzwiniel, J. Kitowski</i>	53
General-relativistic Magneto-hydro-dynamical Simulations of Relativistic Jets from Accreting Spinning Black Holes <i>K. Nalewajko</i>	55

VERONA: a GPU-Accelerated Special Relativistic Hydrodynamics (SRHD) Code for Astrophysical Applications <i>P. Plonka, M. Kapusta, A. Janiuk</i>	57
Finite Element Ion Recombination Simulations for Dosimetry and FLASH on PLGrid <i>R. Benedykciński, J. Kaliński, L. Grzanka, J. B. Christensen</i>	59
Super-Eddington Accretion onto Neutron Stars, Ultra-compact Stars, and Black Holes: Radiative General Relativistic Magnetohydrodynamics Simulations <i>F. Kayanikhoo, M. Wielgus, D. Lančová, L. Zdunik, M. Urbanec, W. Kluźniak</i>	61
Anthropogenic Hazards Research with EPOS TCS AH EPISODES Platform <i>J. Kocot, T. Balawajder, M. Leśniak, M. Makuch, H. Stejkowski, J. Sto, B. Wenda, S. Lasocki, A. Leśnodorska, A. Mtupa-Ndiaye, B. Orlecka-Sikora</i>	63
Molecular Dynamics Study of Tetraglyme Solutions of Two Lithium Salts with Isomeric Anions: LiTFSI and LiPF ₆ <i>A. Eilmes, P. Kubisiak, C. Nicotri</i>	65
Applications of Inter-reactant Interaction Surfaces in the Study of Membrane Degradation in Alkaline Fuel Cells <i>M. Różga, A. Michalak</i>	67
Theoretical Study of Alkaline Degradation Pathways of Carbazolium Cations in Asymmetric Anion-Exchange Membranes <i>O. Żurowska, M. Kukulka, A. Michalak</i>	69
Computational Investigation of Temperature-Dependent Luminescence in Transition Metal Complexes <i>P. Bonarek, R. Jankowski, J. Wang, S. Chorąży</i>	71
Palmitoylation as a Structural Switch for the CFTR Ion Channel <i>J. Mróz, W. Kopeć, M. Sikora</i>	73
Mechanistic Plasticity of the Diels - Alder Reaction: the Role of Substituent and Solvent Effects <i>A. Łapczuk</i>	75
Abduction-Guided Experimental Design for Biomimetic CaP Synthesis via MD Screening of Ca-P Ionic Interactions <i>K. Stafin, P. Śliwa, M. Piątkowski</i>	77
Theoretical Modeling of Two D-sp ³ -A Type Donor-Acceptor Systems with Extreme Solvatochromism <i>W. Radzik, M. Andrzejak</i>	79
Development of AMBER Force Field Parameters for W-cofactors and Application to Study W-enzymes' Catalytic Properties <i>M. Szaleniec, V. Baerle, T. Attucci, C. Andreini, A. Raczyńska</i>	81

Mathematical Model Explaining the Structural Transformation of Proteins Causing Neurodegenerative Diseases	83
<i>I. Roterman-Konieczna, D. Dulak, K. Stapor, L. Konieczny</i>	
Spatial Model of the Presynaptic Bouton for Neurotransmitter Transport Simulation by Using Cellular Automaton	85
<i>M. Gierdziewicz, R.-M. Stama</i>	
Mechanistic Insights into the Cycloaddition of (2E,4E)-2,5-Dinitro-2,4-hexadiene with Diazomethane from Density Functional Theory	87
<i>K. Kula, R. Jasiński</i>	
Language Driven Therapy Design in Predictive Oncology	89
<i>F. Ręka</i>	
DFT Calculations of Flavin Derivatives in Their Anionic Forms	91
<i>I. Gulaczyk, D. Prukala, E. Zubova, R. Cibulka, M. Sikorski</i>	
Domain Bias in Deep Learning for Cytology	93
<i>J. Krupiński, E. Jamro, M. Wielgosz, A. Dąbrowska-Boruch</i>	
Computer Vision Pipelines for High-Content Microscopy Reveal Heritable Single-Cell Immune States	95
<i>T. Co Nguyen, J. A. Redondo, M. Kochończyk, P. Paszek</i>	
Cross-Polarized Skin Image Synthesis from Non-Polarized Inputs Using Deep Generative Models	97
<i>P. R. Popielski, S. Wilczyński, J. Żmudzki, K. Kwieciński</i>	
Author Index	99

The Stubborn Old Protein Folding Problem in the Era of ExaFlops, Qbits and AI

Olav Zimmermann

Simulation & Data Laboratory “Biology”
Jülich Supercomputing Centre (JSC)
Forschungszentrum Jülich GmbH, Germany
olav.zimmermann@fz-juelich.de

When in 2024 the Nobel Prize in Chemistry was awarded to Demis Hassabis and John Jumper from Google DeepMind (developers of AlphaFold2), many, including many scientists claimed that the 50 year old protein folding problem had been solved.

This talk will point out the difference between the structure prediction problem that AlphaFold attempts to solve and the protein folding problem and explain why the latter is much harder and, apart from few successes in the last decades, still essentially unsolved.

After presenting several semi-successful, pseudo-successful, and failed attempts of the community (including a large part of our own work using Monte Carlo simulations) the talk will try to dissect this old and stubborn problem and investigate the question how new capabilities available at Supercomputer centers such as exascale computing, quantum computing, and AI could be harnessed to finally solve it. Along the way I will also try to analyse why a single group at DeepMind could outcompete the entire academic community in the structure prediction competition CASP within a few years, what lessons can be learned from this, and why the protein folding problem might be well suited to assess the current capabilities and competitiveness of academic research.

Towards European Infrastructure for In Silico Medicine

Marian Bubak^{1,2}

¹ Sano Centre for Computational Medicine, Czarnowiejska 36, 30-054 Kraków, Poland

² ACC Cyfronet AGH, ul. Nawojki 11, 30-950 Kraków, Poland

<https://sano.science/>; m.bubak@sanoscience.org

Keywords: in silico medicine, digital twin for healthcare, virtual human twin, EU platform

Investigations in area in silico (computational) medicine [1], accelerated the Physiome Project [2], are conducted by many researchers supported by the VPH Institute (currently the Society for In Silico Medicine), and outcomes are presented at the VPH Conferences [3].

The goal of this lecture is to present the emerging European infrastructure supporting research in computational medicine and the practical use of their results.

In 2022-2025, a systematic approach to computational medicine was pursued by the consortium of the EU Project EDITH, aiming to facilitate implementation of the concept of digital twins in healthcare (DTH). It resulted in an idea of virtual human twin (VHT) as an integrated environment enabling multi-scale, multi-temporal and multi-disciplinary representation of quantitative human physiology and pathology [4]. The EDITH Roadmap [5] presents recommendations to enable efficient implementation of VHT related to assessment of creators and consumers, basic building blocks of VHT technologies and infrastructures, ELSI, standards, regulatory aspects, and sustainability conditions. With VHT one may: access every digital twin developed so far; search catalogue by data type, anatomical location, age of the patient; match any digital twin with any available dataset valid as input for that model; run models on every digital dataset available; orchestrate multiple digital twins to build multiscale or multisystem models; script it, and save scripts for automation or reuse. The Cyfronet and Sano Centre team elaborated the demonstrator to execute DTHs on HPC resources for scale-out studies; it enables reproducibility of computer simulations.

The European Commission decided to support practical implementation of the EDITH Roadmap as the European VHT initiative [6]. The first step is elaboration and operation of the VHT platform by the VHT NET Consortium (CINECA, Sano Centre, ENG, NTT Data, In Silico Trials). The demonstrator will be exploited in development of the VHT platform

Acknowledgements. I acknowledge the support of EU grants Teaming Sano N857533, EDITH 101083771, GEMINI 101136438 (HORIZON). I'm very grateful to Marek Kasztelnik, Maciej Malawski, Jan Meizner, Piotr Nowakowski, and Piotr Połec for discussions.

References

1. Peter Coveney, Roger Highfield: Virtual You. How Building Your Digital Twin Will Revolutionize Medicine and Change Your Life, Princeton University Press, 2023.
2. Physiome Project - <https://www.auckland.ac.nz/en/abi/our-research/research-groups-themes/physiome-project.html>
3. VPH Society for In Silico Medicine - <https://vph-society.org/>
4. M. Viceconti, M. De Vos, S. Mellone, L Geris: Position paper: from the digital twins in healthcare to the virtual human twin: a moon-shot project for digital health research, IEEE Journal of Biomedical and Health Informatics 28 (1), 491-501, 2023.
5. EDITH European Virtual Human Twin <https://www.edith-csa.eu/>
6. European VHT initiative <https://www.virtualhumantwins.eu/>

Current Activities and Further Development of the National Supercomputing Infrastructure of the Czech Republic

Vít Vondrák

IT4Innovations national supercomputing center, VSB-Technical University of Ostrava, 17. listopadu
2172/15, 70800 Ostrava-Poruba, the Czech Republic

`vit.vondrak@vsb.cz`

IT4Innovations National Supercomputing Center at the VSB – Technical University of Ostrava operates and provides the most powerful computing capacities for the research environment of the Czech Republic. Under the leadership of CESNET, the operator of the national research and education network, and together with CERIT-SC at Masaryk University, we form the strategic national research e-infrastructure, e-INFRA CZ. In addition to providing services of the e-INFRA CZ supercomputing infrastructure, IT4Innovations also conducts research activities in the areas of high-performance computing (HPC), large-scale data analytics, artificial intelligence, and quantum computing.

Currently, IT4Innovations operates the EuroHPC supercomputer Karolina, which is being integrated with the VLQ quantum computer acquired by the LUMI-Q consortium and EuroHPC. This integration is based on the LEXIS software platform developed at IT4Innovations, which enables simple, seamless, and secure access to diverse computing and data resources, including quantum computers. In addition, it enables the execution of complex computational workflows across different computing and data resources, including hybrid HPC and quantum computations. The LEXIS platform will also become a component of the European federated platform integrating all EuroHPC computing resources.

The presentation will introduce the current and planned capacities of IT4Innovations, the policies for accessing them, and their use by user communities. It will also present the current services and in-house developed software platforms that extend e-INFRA CZ capabilities. Finally, the newly established Czech AI Factory, particularly its planned services and potential topics for collaboration with other AI factories will also be introduced.

HARM-COMBO: Development and Performance Tests of the Hybrid (MPI+OpenMP) GRMHD Code

Lukasz Klepczarek¹, Agnieszka Janiuk¹, Ireneusz Janiuk², Piotr Płonka¹

¹ Center for Theoretical Physics, Al. Lotników 32/46, 02-668 Warsaw, Poland

² Self-employed

{lklepczarek, agnes, pplonka}@cft.edu.pl

Keywords: computational astrophysics, black holes, hybrid model, scalability, magnetohydrodynamics

1. Introduction

Continuous development of the computer hardware allows for pushing the boundaries in numerical astrophysics. At the same time, rising energy costs require careful allocation of resources. To address this issue, we report recent updates to HARM-COMBO, a C-based GRMHD code using hybrid MPI+OpenMP parallelism. We characterize its scaling behavior across various configurations to facilitate optimal usage of computational power for simulations.

2. Description of the problem

Physical motivation for this work is to enable the evolution of spin and mass of the compact object, as well as account for the self-gravity force acting on accreting material together with the calculations of microphysics in the astrophysical plasmas described by a three-parameter (density, temperature, electron fraction) equation of state (EOS). Such a scenario can represent a system with more massive surrounding material than the central object.

Technical motivation, on the other hand, is to reduce fragmentation of the physics, improve maintainability of the code, unify data structures, find efficient use of resources for a given configuration, and to modernize a code building system for better portability.

3. Related work

In contrast to general-purpose GRMHD frameworks (such as Athena++ [4], BHAC [5], Cosmos++ [6]), the specific focus of our HARM-COMBO development is to incorporate the self-gravity of the accreting material in a way that allows the central compact object's mass and spin to evolve during the simulation.

4. Solution of the problem

To reduce fragmentation and improve maintainability, we merged the previously separate HARM-SELFG and HARM-EOS developments into a single main branch, HARM-COOL, while targeting consistent outputs and similar runtime compared to the pre-merge states (physical description can be found in [1], [2], and [3], respectively). Validation was enforced through bash regression tests in GitHub Actions, comparing outputs against reference data.

Legacy build system using makefiles was replaced with a CMake + configuration file workflow, in which a global defaults file plus a problem init file select the physics options and feature toggles. Each out-of-source build directory stores the resolved options, enabling multiple, coexisting, reproducible configurations without modifying the source tree.

Performance was evaluated in strong scaling on the Ares supercomputers: ~300k CPU-hours. We tested MPI/OpenMP performance for the 2D grid (1536x768) and for 3D grid (288x256x128)

on the Ares with 48 cores per node within the range of 8-256 nodes. In the first case, the best result is for 1 MPI x 48 threads, with 2x24 close behind, while MPI-heavy layouts scale poorly. In the latter case, the optimal configuration shifts to a moderate hybrid balance with 4x12 best at the largest scale, then 2x24 and 8x6 close. Both extremes, pure OpenMP and pure MPI, underperform.

5. Conclusions

This work consolidates and modernizes the hybrid MPI+OpenMP GRMHD code HARM-COMBO to support simulations combining self-gravity with a three-parameter EOS. Merging former branches and adding automated regression tests establishes a workflow that preserves correctness during ongoing development. Migration to CMake improves portability across HPC systems. Strong scaling tests on the Ares supercomputer show that the optimal MPI/OpenMP balance depends on the dimensionality: for the 2D grid (1536/768), the best results are obtained with OpenMP-heavy setups (1-48/2-24 MPI/OpenMP ratio), but for the 3D grid (288x256x128) the optimal result is for a hybrid configuration (4-12 MPI/OpenMP ratio).

Acknowledgements. NCN 2023/50/A/ST9/00527; PLGrid/ACC Cyfronet AGH, PLG/2025/018186.

References

1. Janiuk et al., *A&A* 677 (2023).
2. Janiuk et al., *A&A*, in press, arXiv:2511/05473 (2026).
3. Janiuk, *Acta Phys. Pol. B Proc. Suppl.* 16 (2023).
4. White et al., *ApJS* 225 (2016).
5. Porth et al., *Comput. Astrophys.* 4 (2017).
6. Anninos et al., *ApJ* 635 (2005).

AI-Based Job Runtime Prediction for the ALICE Computing Grid at CERN

Tomasz Lelek¹, Bartosz Baliś^{1,2}, Marcin Kurdziel¹, Mikołaj Zasada¹

¹ Faculty of Computer Science, AGH University of Krakow, Poland

² ACC Cyfronet AGH, Kraków, Poland

t.lelek@agh.edu.pl

Keywords: ALICE, CERN, runtime prediction, deep learning, meta-learning, job scheduling

1. Introduction

ALICE (A Large Ion Collider Experiment) is one of the major experiments at the Large Hadron Collider at CERN, studying heavy-ion collisions to investigate the quark-gluon plasma. Its computing demands are handled by the ALICE Grid, a worldwide infrastructure of about 60 clusters that processes up to 500,000 tasks per day. We are developing an AI-based system for data-driven prediction of job runtimes, integrated with the ALICE Job Broker (Fig. 1).

2. Description of the problem

Job scheduling in the ALICE Grid relies on a time-to-live (TTL) parameter - the declared maximum execution time. A job is assigned to a node only if the node's remaining availability window exceeds the declared TTL. Operators assign TTL values manually, typically defaulting to 20 hours regardless of actual runtime [1]. Since nodes lease resources in 24-hour windows, tail-end slots shorter than the inflated TTL go unused, wasting capacity across the Grid. Accurate prediction of job runtimes can mitigate this effect.

3. Related work

Our prior work [3] applied classical machine learning to job runtime prediction in the ALICE Grid. We now extend it with a deep learning pipeline, GPU-accelerated inference [1], and Reptile-based meta-learning for adaptation to non-stationary workloads [2].

4. Solution of the problem

The ALICE scheduling pipeline consists of several services (Fig. 1). Users submit jobs described in the Job Description Language (JDL) to a Task Queue. When a compute node becomes available, its Site Job Agent sends a match request to the central Job Broker, which selects a suitable job. In the proposed architecture, a Prediction Service is hosted at ACC Cyfronet AGH alongside the training infrastructure. When the Broker cannot find a job whose operator-assigned TTL fits the node's remaining time, it queries the Prediction Service for a tighter runtime estimate and uses it as the effective TTL. This fallback path widens the pool of eligible jobs without risking premature termination.

On the data-collection side, runtime monitoring logs and JDL metadata are continuously gathered and stored in a data lake at Cyfronet. From this raw data we extract 20 categorical and numerical features (host identity, CPU model, collision system type, production campaign, etc.) that feed into model training on an HPC cluster. Trained models are deployed directly to the co-located Prediction Service. The full cycle - data ingestion, training, deployment - runs periodically so that models stay current with evolving workloads.

A key challenge is that the Prediction Service must respond within roughly 40 ms per request to keep up with the Job Broker’s matching rate. We address this with GPU-accelerated inference using INT8 weight quantization and 2:4 semi-structured sparsity on Nvidia A100 hardware. INT8 quantization achieves a $1.8\times$ single-sample speedup (0.39 ms vs. 0.71 ms baseline), while sparsity delivers up to $4\times$ throughput gain for large batches - all without degrading prediction accuracy [1].

For model accuracy, the core predictor is a feedforward neural network with learned embeddings for high-cardinality categorical features. On a dataset of over one million ALICE Grid jobs it achieves an average RMSE of 1.9 hours on held-out test data, compared with 14.23 hours for manual TTL assignment [1]. Because ALICE workloads are non-stationary - new production campaigns and infrastructure changes cause data and concept drift - we further apply Reptile-based meta-learning [2] that adapts the model to unseen job types using as few as 40 samples, delivering a 15–17 % reduction in MAE and SMAPE over non-adaptive baselines.




Fig. 1. Architecture of the ALICE computing framework with the proposed AI-based prediction service. Green elements are the components developed in this work.

5. Conclusions

We presented a complete pipeline for AI-driven job runtime prediction in the ALICE Grid, from data collection and feature engineering through adaptive deep learning models to a low-latency prediction service ready for integration with the ALICE Job Broker. The system reduces the TTL estimation error from over 14 hours (manual) to under 2 hours, while meeting the sub-40 ms latency budget required by the scheduler. Ongoing work focuses on production deployment of the Prediction Service and evaluation of its impact on overall Grid resource utilization.

Acknowledgements. This work is co-financed by the Polish Ministry of Science and Higher Education under Agreement No. 2023/WK/07 and through the PMW program. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid and ACC Cyfronet AGH for providing computer facilities and support within computational grant no. PLG/2024/017775 and PLG/2024/017612.

References

1. Lelek, T., Mazurek, S., Wielgosz, M., Balis, B. (2026). Fast prediction of job execution times in the ALICE Grid through GPU-Based Inference with Quantization and Sparsity Techniques. Submitted to Journal of Computational Science.
2. Lelek, T., Zasada, M., Balis, B., Kurdziel, M. (2025). Adapting Runtime Prediction Models for the ALICE CERN Grid Using Meta-Learning. Submitted to Journal of Grid Computing.
3. Balis, B., Lelek, T., Bodera, J., Grabowski, M., Grigoras, C. (2024). Improving prediction of computational job execution times with machine learning. *Concurrency and Computation: Practice and Experience*, 36(2), 7905.

Benchmarking a Heterogeneous Metagenomics Pipeline on Ares and Helios Supercomputers

Piotr Kica^{1,2}, Kinga Zielińska⁴, Michał Orzechowski^{1,2,3},
Karol Zajac¹, Tomasz Kościółek^{1,2}, Maciej Malawski^{1,2}

¹Sano Centre for Computational Medicine, Kraków, Poland

²Faculty of Computer Science, AGH University of Krakow

³ACC Cyfronet AGH

⁴Małopolska Centre of Biotechnology, Jagiellonian University, Kraków, Poland

{p.kica, m.orzechowski, k.zajac, t.kosciolek, m.malawski}@sanoscience.org,
kinga.zielinska@uj.edu.pl

Keywords: HPC, benchmarking, Nextflow, metagenomics, SLURM

1. Introduction

Metagenomics is the study of microorganism genetic material recovered directly from environmental or clinical samples, e.g. bacteria, viruses, fungi, without culturing individual organisms. A typical analysis involves quality control of raw sequencing reads, removal of (host) contamination, performing taxonomy and functional potential-focused sequence classification, as well as recovery of metagenome-assembled genomes (MAGs).

Nextflow [1] is an established bioinformatics tool that natively handles scheduling heterogeneity. Each process in the pipeline declares its own CPU count, memory, and Nextflow submits individual jobs accordingly.

With the introduction of the Helios supercomputer for the general public, we decided to migrate the workflow for higher efficiency of computations.

This paper compares the pipeline's performance on 100 metagenomic samples across higher efficiency computers: **Ares** and **Helios**.

2. Description of the problem

The pipeline consists of 13 processing steps whose resource demands vary by orders of magnitude. MEGAHIT assembly and Bakta gene annotation are the two heaviest steps, together consuming about 67% of total CPU-hours in our benchmark. Bakta and CheckM2 are memory-intensive steps, each requiring over 150 gigabytes of RAM. At the other end, steps such as Sylph taxonomy and TPM calculation finish in minutes per sample. This heterogeneity makes uniform resource allocation impractical: requesting the same hardware for all steps wastes either time on heavy processes or allocation on lightweight ones. Several steps also depend on large reference databases (GTDB r226, Bakta v6, CheckM2) that must be accessible at runtime. Moreover, tools such as bowtie2, MEGAHIT and Bakta are very I/O intensive.

3. Related work

Authors of [4] benchmarked Nextflow executor backends - HPC, Kubernetes, and local on a genomics variant-calling pipeline, showing that performance varies significantly across infrastructure types and dataset sizes. Our work extends this approach to metagenomics workloads by comparing two HPC systems with different hardware specifications.

4. Solution of the problem

We built a Nextflow pipeline that automates the analysis of scheduling heterogeneity via three parallel branches after shared preprocessing: **Branch 1** performs taxonomic profiling with Sylph [2] against the GTDB r226 database. **Branch 2** constructs a per-sample gene catalog using Bakta [3] annotation and CD-HIT clustering, then quantifies gene abundance as TPM. **Branch 3** recovers MAGs through MetaBAT2 binning, assesses their quality with CheckM2, and classifies high-quality bins with GTDB-Tk.

All bioinformatics tools are installed through Conda environments, which are cached on each cluster and reused across runs. The pipeline uses `errorStrategy = „ignore”` so that a single sample failure does not halt the rest of the batch. Nextflow saves execution traces which are used for the comparison presented below.

We processed 100 metagenomic samples on both Ares and Helios with the same configuration and compared 1163 Nextflow process executions. Table 1 shows results for processing the samples on the supercomputers.

Tab. 1. Aggregated experiment results.

	Helios	Ares
Total CPU usage	1235 h	1457 h
Experiment wall-time	147 min	162 min
Weighted mean efficiency	61.2 %	65.3 %

In total we processed 145 GB of FASTQ.GZ data twice and Helios was also faster on average across all workflow steps by about 15%. Interestingly Ares reports higher efficiency - most likely due to I/O intensive nature of the pipeline which reduces CPU efficiency when accessing storage.

5. Conclusions

The main takeaway is that MEGAHIT and Bakta dominate the compute budget. Any effort to reduce total pipeline runtime should focus on these two steps first, e.g. with better resource tuning. The main difference in processing speed can be explained by Helios having *AMD EPYC 9654 96-Core @ 2.4 GHz* processors which are faster than Ares' *Intel(R) Xeon(R) Platinum 8268 CPU @ 2.90 GHz*.

Acknowledgements. Polish Ministry of Science and Higher Education: Grant MEiN/2023/DIR/3796; European Commission: H2020 Grant 85753. This research was supported by PLGrid Infrastructure. Computations have been performed on the Ares and Helios supercomputers at ACC Cyfronet AGH (grant no. PLG/2025/019015 and PLG/2025/018235).

References

1. Di Tommaso, Paolo, et al. "Nextflow enables reproducible computational workflows." *Nature biotechnology* 35.4 (2017): 316-319.
2. Shaw, Jim, and Yun William Yu. "Rapid species-level metagenome profiling and containment estimation with sylph." *Nature Biotechnology* 43.8 (2025): 1348-1359.
3. Schwengers, Oliver, et al. "Bakta: rapid and standardized annotation of bacterial genomes via alignment-free sequence identification." *Microbial genomics* 7.11 (2021): 000685.
4. Spišáková, Viktória, et al. "Nextflow in bioinformatics: Executors performance comparison using genomics data." *Future Generation Computer Systems* 142 (2023): 328-339.

HPC in Global Quantification and Personalisation of Regulated Cardiovascular Model - a Cohort Study

Karolina Tlałka^{1,2}, Ian Halliday^{2,3}, Andrew Narracott^{2,3}, Maciej Malawski^{1,4}

¹ Sano Centre for Computational Medicine, Nawojki 11, 30-072 Kraków, Poland

² Division of Clinical Medicine, School of Medicine and Population Health, University of Sheffield, Sheffield, United Kingdom

³ Insigneo Institute for in silico Medicine, Sheffield, United Kingdom

⁴ Faculty of Computer Science, AGH University of Kraków, Kraków, Poland

{k.tlalka,m.malawski}@sanoscience.org,
{i.halliday,a.j.narracott}@sheffield.ac.uk

Keywords: cardiovascular models, personalization, global optimization, modelling, simulation

1. Introduction

Mechanical cardiovascular (CV) models utilise the principles of fluid mechanics to predict pressures, volumes and flows in vascular locations [1]. Supplied with patient-specific data, they can be used to assess personal CV health. CV parameters change constantly in response to innervation, driven by a changing environment (temperature, body position) and it is therefore necessary to introduce a baroreflex submodel, which accounts for short term neural feedback control of the mechanical CV system [1,2].

2. Description of the problem

Personalisation is perhaps the crucial feature of the critical path between patient-specific CV models and the clinic [1]. Adding the necessary innervation to a CV model results in increased model non-linearity and a complication of the personalisation task, by introducing the need of computationally expensive global methods for personalisation [1]. This drawback can be offset by using High Performance Computing (HPC). Here, we present the method by which we used Ares (PLGrid, Cyfronet AGH) to personalise model outputs of stroke volume, mean arterial pressure and heart rate, in the setting of induced central hypovolemia, in 10 subjects [3], using global optimisation.

3. Related work

To justify less expensive local methods of sensitivity analysis and personalization, most prior art assumes that interactions between model input parameters are not influential, or that true solutions lie close to initial parameter estimates. Our research has exposed the limitations of this assumption [1,4]; indeed, it argues that global methods are essential.

4. Solution of the problem

We used a 4-chamber zero-dimensional CV model published before [1] with baroreflex regulation model proposed by Ursino [2]. Personalisation to subject data was performed by solving **global optimization task and minimizing a bespoke loss function** [1]. We used Differential Evolution method [5] (120 agents, maximum 35k iterations, Julia implementation [6]) and repeated the optimization 100 times for each subject to ensure convergence and to avoid landing on a false solution. Overall, we performed 1000 simulations (10 array jobs) with average efficiency exceeding 80%. 998 simulations were successful. Peak memory usage was approxima-

tely 11 GB per single optimiser run. Average time of one simulation was 30.6 hours with standard deviation of 14.3 hours. Cohort personalization was parallelized. After 61 hours, we achieved personalized responses for all subjects. There was a variability between execution time for subjects (Fig. 1). Estimated time for 1000 simulations conducted on local machine, considering RAM limitations (maximum 12 GB), which enforce to calculate in series, is approximately 1200 days.

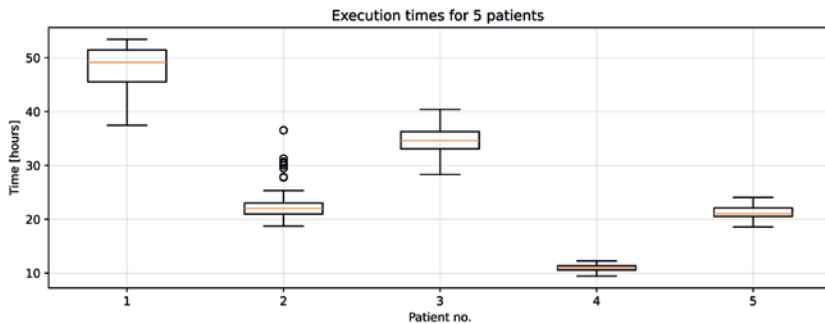


Fig. 1. Execution time for subjects 1-5.

The quality of fit was significantly improved. Mean errors of stroke volume, mean arterial pressure and heart rate prediction do not exceed 6%. Full convergence was achieved for 4 subjects.

5. Conclusions

We achieved fully convergent solutions for some subjects. The discrepancies can be the effect of variability between specific subject responses and mechanisms other than baroreflex which influence cardiovascular balance (such as emotional stress). Further research involves testing different loss function and different subsets of parameters.

Acknowledgements. The publication was created within the project of the Ministry of Science and Higher Education “Support for the activity of Centers of Excellence established in Poland under Horizon 2020” on the basis of the contract number MEiN/2023/DIR/3796. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 857533. This publication is supported by Sano project carried out within the International Research Agendas programme of the Foundation for Polish Science, co-financed by the European Union under the European Regional Development Fund. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grants no. PLG/2024/017108 and PLG/2025/018168.

References

1. Tlalka et al., Personalization of a fully closed-loop lumped-parameter cardiovascular model with baroreflex in response to orthostatic stress. Submitted to PLOS One.
2. Ursino, doi: 10.1152/ajpheart.1998.275.5.H1733. PMID: 9815081
3. Høiseth et al., doi: 10.1097/CCM.0000000000000766. PMID: 25513787
4. Tlalka et al., PLOS Comp Biol, doi: 10.1371/journal.pcbi.1012377
5. Price et al., doi: 10.1007/3-540-31306-0
6. Dixit, Rackauckas, doi: 10.5281/zenodo.7738525

Automating CFD Workflows on HPC with PyFluent: from Batch Simulations to Large-Scale Study

Karol Zajac¹, Magdalena Otta^{1,3}, Maciej Malawski^{1,2}, Andrew Narracott³

¹Sano Centre for Computational Medicine, Kraków, Poland

²Faculty of Computer Science, AGH University of Kraków

³School of Medicine and Population Health, University of Sheffield, UK

{k.zajac, m.otta, m.malawski}@sanoscience.org,
a.j.narracott@sheffield.ac.uk

Keywords: HPC, Large-scale, CFD Simulation, Workflow Automation, PyFluent

1. Introduction

High-fidelity Computational Fluid Dynamics (CFD) studies are increasingly performed on High Performance Computing (HPC) systems, but manual and interactive setup of simulations in tools like ANSYS Fluent limits scalability and reproducibility. We explore automation of CFD workflows using PyFluent [1] on HPC, enabling batch execution, parameter studies, and integration with modern workflow engines for large-scale simulation campaigns.

2. Description of the problem

Typical Fluent workflows often rely on GUI-driven preprocessing in ANSYS SpaceClaim, manual solver setup, and interactive post-processing. This slows experimentation, prevents reproducibility, and complicates running large parameter sweeps, especially on clusters. Additionally, PyFluent is not yet fully stabilised in terms of API consistency or feature coverage, which poses challenges for building reliable end-to-end CFD automation.

3. Related work

CFD batch execution has been explored through scripting, transcribing and journaling in ANSYS Fluent and open-source frameworks such as OpenFOAM. Workflow management tools like EasyVVUQ and Dask enable scalable sensitivity analysis and uncertainty quantification studies, but full CFD workflow automation remains limited [2].

4. Solution of the problem

We designed a semi-automated CFD workflow (Fig. 1) that executes Fluent simulations in batch mode on HPC using PyFluent scripts and SLURM array jobs. Geometry preparation, meshing, solver configuration, and execution are orchestrated through Python pipelines that reduce manual effort and improve reproducibility.

Experiments can be customised via user-defined pre- and post-processing tasks, enabling centralised management within a Python ecosystem. Tools such as Dask and EasyVVUQ can replace the array job batch execution approach, enabling large-scale sensitivity analysis and uncertainty quantification studies. While meshing and solver stages are efficiently executed on HPC resources, ongoing work focuses on automating geometry preparation for more complex case studies (e.g. marking named selections), as well as methods to extract additional simulation metrics that are not readily available in standard Fluent output formats. Critically, HPC resources enabled not only single-case simulations but also systematic VVUQ studies with extensive boundary-condition and geometry perturbations, making nearly 1000 simulations feasible across several studies [3] - a scale unattainable on local machines.

Applied automation not only provides faster execution in comparison to local machines but also reduced experiment preparation time. This demonstrates how workflow scripting can replace traditional GUI-based operations, paving the way toward fully automated Fluent CFD pipelines on HPC systems.

This work complements our ongoing research on statistical shape modelling [4] and anatomical variability using Deformetrica, where shape-based uncertainty is quantified within venous and multi-scale modelling workflows [5].

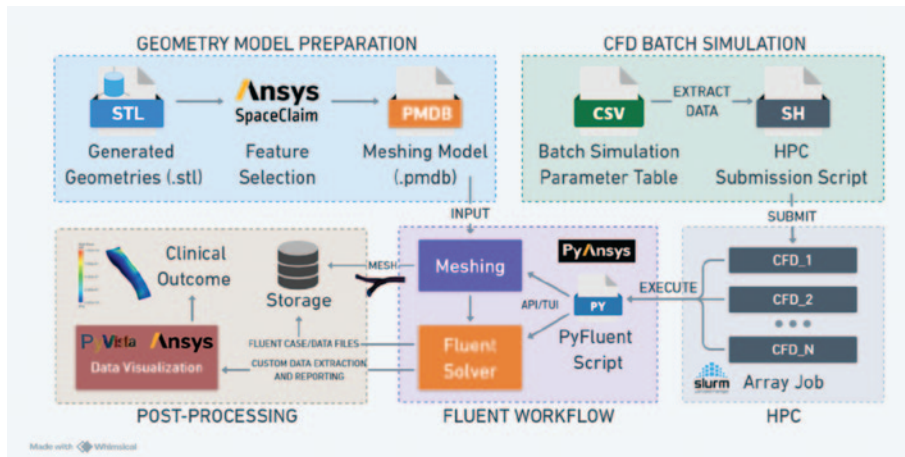


Fig. 1. Basic batch execution of CFD workflow with PyFluent on HPC.

5. Conclusions

Automating Fluent workflows with PyFluent on HPC significantly accelerates CFD experimentation and enables scalable research studies. Despite current API limitations, the approach shows strong potential for reproducible, large-scale CFD research as required manual intervention in geometry preparation and advanced simulation result extraction is progressively reduced.

Acknowledgements. This publication is part of a project that has received funding from the European Union’s Horizon Europe research and innovation programme under grant agreement No 101136438. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018168.

References

1. PyFluent Documentation. PyAnsys Project web site: <https://fluent.docs.pyansys.com/version/stable/index.html>
2. Morris, P.D. et al. Computational fluid dynamics modelling in cardiovascular medicine. *Heart* 2 (102), 18–28 (2016). doi:10.1136/heartjnl-2015-308044
3. Otta, M. et al. Towards Sensitivity Analysis: 3D Venous Modelling in the Lower Limb. ICCS 2025. doi:10.1007/978-3-031-97557-8
4. Otta, M. et al. High-throughput statistical shape modelling of venous anatomy using containerised Deformetrica on Ares supercomputer; submitted to KUKDM 2026.
5. Otta, M. et al. bioRxiv preprint, under review, 2026. doi:10.64898/2026.02.17.706277

PubDB: an Institutional Publication Database for KPI Monitoring and Visibility

Taras Zhyhulin¹, Adam Nowak¹, Jan Meizner^{1,2}, Piotr Nowakowski^{1,2}

¹ Sano Centre for Computational Medicine, Czarnowiejska 36/C5 st., 30-054 Kraków, Poland

² ACC Cyfronet AGH, Nawojki 11 st., 30-950 Kraków, Poland

{t.zhyhulin, a.nowak, j.meizner, p.nowakowski}@sanoscience.org

Keywords: institutional publication database, research output tracking, KPI monitoring, bibliographic metadata, Ruby on Rails

1. Introduction

With the growth of research institutions, the number of publications produced by their researchers is steadily increasing. Scientific output plays an important role in evaluating research performance, securing research grants, and building institutional visibility. Therefore, institutions need reliable and structured ways to manage and present their publication records. These circumstances motivate the development of digital tools that support systematic collection, curation, and presentation of institutional scientific output for different groups of users.

2. Description of the problem

The proposed solution stems from a need for a centralised application that meets the needs of multiple stakeholders. Researchers should be able to register their publications, provide detailed metadata, attach identifiers and links, associate them with related conferences and journals, and search for works authored by other members of the institution using advanced filtering options. Managers and administrators require tools to extract publication statistics, monitor productivity, calculate KPIs, and support project reporting and evaluation. External users should have transparent access to publication records, citation information, and aggregated institutional statistics.

To address these needs, the proposed system aims to implement a centralised publication database with automated integration of external academic services (e.g., Google Scholar APIs and Dataverse widgets), enabling efficient data retrieval, structured storage, and statistical analysis.

3. Related work

PBN [1] is a national Polish system used to register and aggregate publications of institutions, authors and journals. It provides advanced filtering and relational linking between entities; however, it does not offer built-in tools for flexible statistical reporting or dashboard creation. Baza Publikacji AGH [2] serves as an institutional registry of researchers' publications and supports statistical data collection, although its analytical capabilities remain limited.

4. Solution of the problem

The PubDB application [3] is a full-stack solution, implemented in Ruby on Rails with a PostgreSQL database based on a normalised domain model (Fig. 1).

Authentication and role-based access control are integrated with Microsoft Entra ID, with organisational roles mapped to application roles (user and moderator). Moderators manage research groups and advanced publication metadata (e.g., KPI indicators, open access charges),

while users register publications and maintain metadata and additional materials (e.g., identifiers and links).

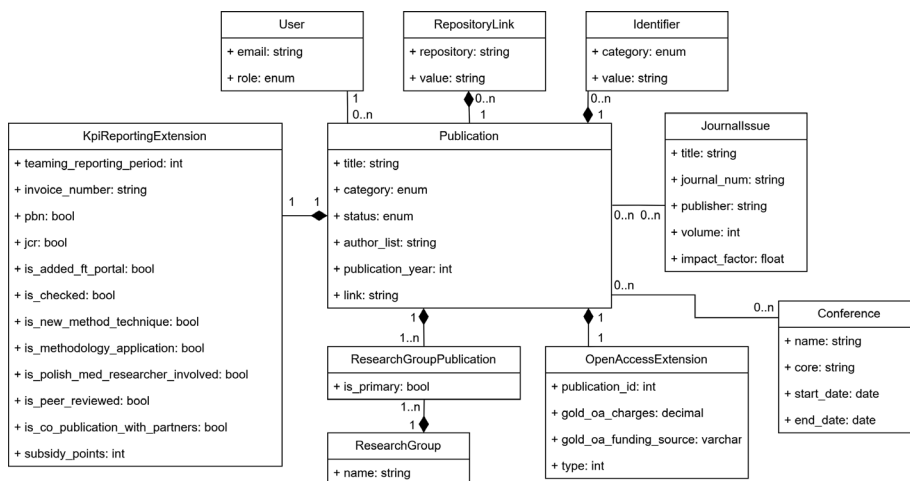


Fig. 1. UML class diagram of the PubDB system reflecting the underlying relational database schema.

The application provides advanced multi-criteria search and exposes a REST API with aggregated statistical indicators, visualised in a React-based KPI dashboard [4] that can be embedded into the institution’s website to present and promote scientific output.

The system is at an early stage of development; planned work includes customisable reporting features enabling detailed researcher- or team-specific report generation, exporting citations export, automated publication synchronisation with external indexing services (e.g., Google Scholar, Scopus, Web of Science), and Dataverse integration to detect and display associated datasets with citation and download options.

5. Conclusions

The designed system, together with the integrated services, enables institutional publication tracking, monitoring, and statistical analysis, while promoting researchers’ scientific output and facilitating publication retrieval. It assists researchers in managing their publication records, and supports managers in statistical analysis and settlement of research grants.

Acknowledgements. This publication is partly supported by the EU H2020 grant Sano (857533) and by the Minister of Science and Higher Education “Support for the activity of Centers of Excellence established in Poland under Horizon 2020” number MEiN/2023/DIR/3796. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACC Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018168.

References

1. Polska Bibliografia Naukowa. Publication registration portal link: <https://pbn.nauka.gov.pl>
2. Baza Danych o Autorach i Publikacjach AGH (BaDAP). Institutional publication database portal link: <https://badap.agh.edu.pl/>
3. PubDB. Sano’s institutional publication database link: <https://kpi.sano.science/>
4. Sano Publications Statistics. Sano website: <https://sano.science/sano-statistics/>

Evaluating the Stability of Deep Neural Networks for Brain Microstructural Parameter Estimation under Varying Diffusion Gradient Schemes

Dominika Ciupek¹, Maciej Malawski^{1,2}, Tomasz Pięciak³

¹ Sano Centre for Computational Medicine, Kraków, 30-054 Poland

² AGH University of Kraków, Kraków, 30-059 Poland

³ ETSI Telecomunicación, Universidad de Valladolid, Valladolid, 47011 Spain

{d.ciupek, m.malawski}@sanoscience.org, tpieciak@tel.uva.es

Keywords: bipolar disorder, magnetic resonance imaging, diffusion tensor imaging, deep learning

1. Introduction

Bipolar disorder (BD) is a chronic condition characterized by profound mood and energy fluctuations. Correct diagnosis remains challenging, with a mean delay of 5-10 years from symptom onset to clinical confirmation [1]. Diffusion-weighted magnetic resonance imaging offers potential for the identification of more objective biomarkers in psychiatric evaluation. This study analyzes the potential of using deep learning (DL) models to correctly estimate diffusion tensor imaging (DTI) [2] microstructural parameters, particularly radial diffusivity (RD), while preserving differences between patients and healthy controls (HC).

2. Description of the problem

DTI requires tailored protocols for reproducible results. Standard 6-direction protocols often lack the stability required for the voxel-level precision needed to distinguish BD from HC [3]. While DL methods can estimate parameters from limited data, their evaluation is still inadequate. Current studies often treat microstructural parameters as natural images, using metrics like MSE or MSSIM to assess performance. However, these metrics prioritize general image similarity over clinical accuracy. Consequently, they do not ensure that the estimated parameters preserve the biological differences necessary for objective psychiatric diagnosis.

3. Related work

Existing DL research for estimating microstructural parameters [4,5] primarily focuses on healthy populations, creating a gap in its clinical application for pathological conditions. Additionally, current literature lacks a systematic evaluation of the stability during the training phase. The impact of various factors, such as data partitioning and weight initialization, on the clinical reliability of these networks remains unexplored.

4. Results and conclusions

The study used the UCLA dataset [6], including HC and BP. A U-Net model was employed to estimate RD across three different numbers of gradient directions: 6, 32, and 64. To verify the stability of the estimates, 100 training runs were conducted for each setup, utilizing different weight initializations. While training a single model requires ~121 hours on a local CPU, training on the Athena cluster reduces the time per model to ~35 minutes using a single GPU. By running up to 40 independent training jobs simultaneously, the full set of 300 models was completed in ~5 hours of wall-clock time. This corresponds to a reduction from over four years of sequential CPU computation to just a few hours using parallel GPU resources.

Performance was assessed using standard metrics alongside Tract-Based Spatial Statistics (TBSS) to localize RD differences between HC and BD. Although all models performed well in terms of image similarity, TBSS analysis of the splenium of the corpus callosum (Tab. 1 and Fig. 1) revealed a systematic inflation of the number of statistically significant voxels in models trained on fewer gradient directions. These findings confirm that reliance on natural image metrics can provide a misleading sense of accuracy. This highlights the critical need for evaluation methods that accurately reflect clinical quality rather than voxel-wise similarity.

Tab. 1. Averaged results in the SCC region on BP data for models trained in three different configurations.

		Number of gradient directions		
		6	32	64
Evaluation metrics	MSE [10^{-6}]	127.7 ± 16.8	70.4 ± 14.3	6.4 ± 1.4
	MSSIM	0.9880 ± 0.0013	0.9930 ± 0.0012	0.9994 ± 0.0001
TBSS analysis	TPR [%]	91.9 ± 11.0	94.2 ± 12.3	62.0 ± 44.6
	FPR [%]	315.9 ± 148.8	305.8 ± 144.2	91.3 ± 123.9



Fig. 1. Visualization of true positive and false positive voxels in the SCC region for models trained in three different configurations.

Acknowledgements. The numerical experiment was possible through computing allocation on the Ares and Athena systems at ACC Cyfronet AGH under the grants PLG/2025/018009 and PLG/2026/019157. This work is supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 857533 and the International Research Agendas Programme of the Foundation for Polish Science No MAB PLUS/2019/13. This publication was created within the project of the Minister of Science and Higher Education “Support for the activity of Centers of Excellence established in Poland under Horizon 2020” on the basis of the contract number MEiN/2023/DIR/3796. Tomasz Pieciak acknowledges the NAWA for grant PPN/BEK/2019/1/00421.

References

1. M. L. Phillips, et al.: Bipolar disorder diagnosis: challenges and future directions. *The Lancet*, 381(9878), 2013, pp.1663-1671.
2. P. J. Basser, et al.: MR diffusion tensor spectroscopy and imaging. *Biophysical journal*, 66(1), 1994, pp.259-267.
3. H. Ni, et al.: Effects of number of diffusion gradient directions on derived diffusion tensor imaging indices in human brain. *American Journal of Neuroradiology*, 27(8), 2006, pp.1776-1781.
4. E. K. Gibbons, et al.: Simultaneous NODDI and GFA parameter map generation from subsampled q-space imaging using deep learning. *Magnetic resonance in medicine*, 81(4), 2019, pp.2399-2411.
5. Q. Tian, et al.: Deepdti: High-fidelity six-direction diffusion tensor imaging using deep learning. *NeuroImage*, 219, 2020, pp.117017.
6. K. J. Gorgolewski, et al.: Preprocessed consortium for neuropsychiatric phenomics dataset. *F1000Research*, 6, 2017.

High-Throughput Statistical Shape Modelling of Venous Anatomy Using Containerised Deformetrica on Ares Supercomputer

Magdalena Otta^{1,2}, Piotr Kica^{1,3}, Maciej Malawski^{1,3}

¹ Sano Centre for Computational Medicine, Kraków, Poland

² University of Sheffield, Sheffield, United Kingdom

³ Faculty of Computer Science, AGH University of Krakow

{m.otta, p.kica, m.malawski}@sanoscience.org

Keywords: HPC, Deformetrica, Statistical Shape Modelling, Docker, Venous Anatomy, Apptainer, Slurm

1. Introduction

Statistical Shape Modelling (SSM) captures and analyses shape variability across a population of anatomical geometries. SSM of venous anatomy performed in Deformetrica [1] requires computationally demanding pipelines, particularly for deterministic atlas construction and subsequent Principal Geodesic Analysis (PGA). To address these requirements and enable reproducible large-scale experimentation, the workflow was migrated to the Ares supercomputer and executed in a containerised environment.

2. Description of the problem

Meaningful PGA outputs depend on a well-converged atlas template, which in practice requires extensive parameter exploration due to sensitivity to kernel widths, smoothness settings, and control-point configurations. Deformetrica is an older library with fragile dependencies, however there is no valid alternative, making containerisation not merely convenient but necessary for reliable execution in HPC environment. HPC also enables access to its GPU implementation, although this was not used here.

3. Related work

Most SSM studies focus on methodology or haemodynamics, with HPC-based atlas construction rarely reported. Containerisation remains uncommon in shape-modelling, especially for vascular anatomy, and Apptainer - despite its maturity in HPC bioinformatics - has seen limited use in atlas-based shape work. Likewise, systematic parameter exploration via HPC array jobs - conceptually similar to distributed hyperparameter optimisation in machine learning - is rarely applied to atlas construction.

4. Solution of the problem

To obtain meaningful PGA modes, a well-converged atlas is needed. We packaged Deformetrica 4.3 in Docker and executed on Ares via Apptainer. Atlas parameters were swept across 50-100 configs using Slurm array-jobs which are well-suited to embarrassingly parallel parameter sweeps; each task is independent, enabling backfill scheduling and high cluster utilisation even during peak load. For a set of 12 venous meshes (~1 mm edges), a run typically used 15 CPUs and 20–50 GB RAM (60–80% efficiency, 10–90 min execution time), scaling with control points (200–800). This pipeline (Fig.1) underpins our recent SSM applications [2,3].

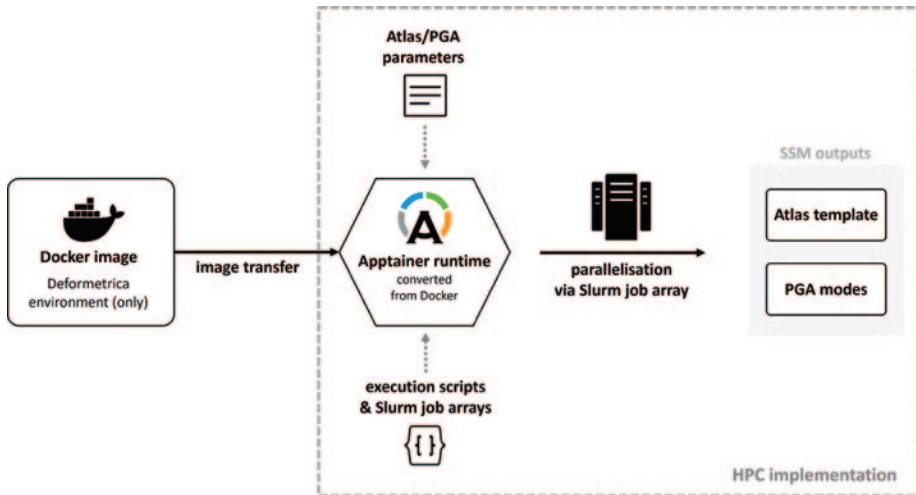


Fig. 1. A workflow for Deformetrica implementation on Ares HPC.

5. Conclusions

Containerisation and HPC improved throughput and reproducibility of SSM modelling. The Slurm array-based sweep provides systematic, reproducible coverage of the atlas parameter space - transforming informal tuning into a documented experimental protocol. The Docker image encapsulates the complete Deformetrica 4.3 dependency graph, ensuring long-term computational reproducibility critical for clinical shape studies. The workflow integrates with our downstream CFD analysis using PyFluent [4].

Acknowledgements. This work was supported by the National Science Centre (NCN) under Preludium'22 grant agreement UMO-2023/49/N/ST6/04252. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 857533 and was created within the project of the Minister of Science and Higher Education "Support for the activity of Centers of Excellence established in Poland under Horizon 2020" on the basis of the contract number MEiN/2023/DIR/3796 and is supported by Sano project carried out within the International Research Agendas programme of the Foundation for Polish Science, co-financed by the European Union under the European Regional Development Fund. The numerical experiments were possible through computing allocation on the Ares system at ACC Cyfronet AGH under the PLGrid allocation PLG/2025/018168.

References

1. Deformetrica 4.3.0 [Computer software] Aramis Lab. <http://www.deformetrica.org>
2. Otta M, et al. *Shape vs Flow: A 2D Statistical Shape Analysis of the Projection of Common Iliac Veins in Patients with Deep Vein Thrombosis*. ShapeMI 2025, LNCS 16171 Springer 2026.
3. Otta M, et al. *Iliac vein morphology and wall shear stress: a statistical shape modelling and CFD analysis of patient-specific geometries* bioRxiv preprint, under review, 2026. doi:10.64898/2026.02.17.706277
4. Zajac K, et al. *Automating CFD Workflows on HPC with PyFluent: From Batch Simulations to Large-Scale Study*; submitted to KUKDM 2026.

Using High Performance Computing for the Application of PINNs in Cardiovascular Models

Elena Gómez del Pozo¹, Karolina Tlalka^{1,3}, Karol Zając¹, Maciej Malawski^{1,2}

¹ Sano Centre for Computational Medicine, Nawojki 11, 30-072 Kraków, Poland

² Faculty of Computer Science, AGH University of Krakow, Kraków, Poland

³ Division of Clinical Medicine, School of Medicine and Population Health, University of Sheffield, Sheffield, United Kingdom

{e.gomez, k.tlalka, m.malawski}@sanoscience.org

Keywords: Physics-Informed Neural Networks (PINNs), Cardiovascular Modeling, Hybrid Data-Driven Physics-Based Modeling, Model Personalization

1. Introduction

Physics-Informed Neural Networks (PINNs) provide a promising framework to combine known physical constraints with experimental data for human model personalization [1]. They allow for results with a significantly reduced amount of training data by infusing knowledge through a physics loss term. In this work, we apply PINNs to cardiovascular lumped-parameter models, where the human cardiovascular system is reduced to an electrical circuit analog.

2. Computational challenge

Traditional calibration and sensitivity analysis of cardiovascular models are computationally expensive due to their nonlinear, time-dependent dynamics, making personalization long and often yielding non-unique solutions [2]. Simpler models are easier to calibrate but fail to capture critical physiological features, such as sharp pressure peaks during cardiac phase transitions – which may be clinically relevant [2].

3. State of the art

There has been significant recent progress in the application of Physics-Informed Neural Networks (PINNs) to forecasting, fluid dynamics, and other PDE systems [1,3]. The application of these approaches to ODE systems like the cardiovascular zero-dimensional models has been investigated [4], although certain limitations remain to be addressed.

4. Proposed solution

The proposed approach employs a hybrid ODE system trained to enhance a two-chamber cardiovascular system model by learning corrective terms added to its derivatives, enabling it to approximate the dynamics of a more complex four-chamber model. This process can be interpreted as infusing physics-informed neural networks into the differential equation solver, such that the baseline mechanistic model is enriched without losing its original structure.

The implementation is primarily based on the PINN_Infuser library [5], which has been adapted to better accommodate the specific case of cardiovascular modeling applications. The four-chamber reference system, used as ground truth during training and evaluation, is derived from a ready-to-use CellML model [6].

Several experimental scenarios were tested assigning 1 core to each. Total execution time on Ares HPC was 1.8 hours with 2.9 GB of memory usage.




Fig. 1. Time evolution of left ventricular volume (Vlv) and systemic venous flow (Qsv) for three models (PINN enhanced 2 chamber model: PINN, 2 chamber model: ODE, 4 chamber model: DATA).

5. Conclusions

The resulting augmented model presents greater detail and fidelity to real physiology. However, the neural network contributions introduce small instabilities that accumulate over a few cardiac cycles and disrupt periodicity. Possible future research includes experimenting with periodicity loss terms or a post-training correction to achieve a zero-net contribution to the derivatives over a cycle.

Acknowledgements. This research was supported by the Polish Minister of Science and Higher Education under the program “Support for the activity of Centers of Excellence established in Poland under Horizon 2020” (MEiN/2023/DIR/3796) as part of Sano Scholarship for Empowering Women in STEM. The project also received funding from the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 857533). We also gratefully acknowledge Polish high-performance computing infrastructure PLGrid, PLG/2025/018168 (HPC Center: ACK Cyfronet AGH), for providing computer facilities and support within computational grant.

References

1. M. Raissi, P. Perdikaris, and G. E. Karniadakis, doi:10.1016/j.jcp.2018.10.045
2. H. Saxton, *Uncertainty quantification and personalisation of lumped parameter models of the cardiovascular system*, Ph.D. dissertation, Sheffield Hallam Univ., Sheffield, U.K., 2025.
3. Ahmad et al., <https://doi.org/10.1016/j.array.2026.100688>
4. W. Perczak and M. Dydek, *A Library in Julia Programming Language for Modelling Cardiovascular System Based on Physics Informed Neural Networks (PINN)*, Diploma project, Wydział Informatyki, Akademia Górniczo-Hutnicza im. Stanisława Staszica w Krakowie, Kraków, Poland, 2026.
5. W. Perczak and M. Dydek. PINNFuser.jl Accessed 01.02.2026. Available at: <https://github.com/mdydek/PINNFuser.jl>
6. Y. Shi, P. Lawford, and R. Hose, doi: 10.1186/1475-925X-10-33

CR39 Track Detection and Segmentation

Michał Rejdak¹, Szymon Rybski¹, Leszek Grzanka¹, Jeppe Brage Christensen²

¹AGH University of Krakow, al. Mickiewicza 30, 30-059 Kraków, Poland

²Paul Scherrer Institute, Forschungsstrasse 111, 5232 Villigen PSI, Switzerland

{mrejdak, szymonrybski}@student.agh.edu.pl

Keywords: CR39, detection, tracks, dosimetry, image processing, segmentation

1. Introduction

CR39 plastic nuclear track detectors are widely used in neutron dosimetry and radiation measurements. Accurate identification and counting of etched tracks is essential for reliable dose estimation. With increasing amounts of data, automated analysis methods are often replacing the time-consuming manual analysis of irradiated detectors, reducing the human bias at the same time.

2. Description of the problem

Automatization of this process remains challenging - aside from tracks, detectors often contain dirt, manufacturing defects and mechanical damage that could be mistaken for tracks. Moreover, high radiation doses often result in overlapping tracks which are especially problematic to detect and count reliably.

3. Related work

Recent studies apply deep learning methods to automate particle track detection and dose estimation in CR39 detectors. Approaches based on Faster R-CNN [1] treat the task as an object detection problem by predicting bounding boxes for individual tracks, improving discrimination between tracks and artifacts. RecoilNet [2] extends the Mask R-CNN framework by performing instance segmentation, enabling pixel-level delineation of track shapes. YOLO-based detectors [3] have also been explored for fast localization and counting of tracks. However, these methods primarily focus on detecting and counting tracks and often perform poorly in regions with overlapping or densely clustered tracks, where multiple tracks may be merged into a single detection or incorrectly segmented. In addition, most existing approaches do not explicitly address the identification of tracks produced by different particle types within the same detector image.

4. Solution of the problem

Current work focuses on exploring several approaches to the problem. As a baseline, we implement classical computer vision methods, including morphological processing with adaptive thresholding. These techniques exploit the geometric nature of particle tracks and provide a computationally lightweight solution. However, simple geometric filters do not reliably detect all particle types, especially when track shapes vary from the expected patterns. Despite this limitation, they are used to generate preliminary annotations that serve as training data for more advanced models.

In parallel, we investigate the use of a YOLO model for track segmentation. We also consider lightweight convolutional architectures that incorporate information about the real physical dimensions of tracks, addressing a limitation of typical object detection models. All approaches are first tested on small batches of detector images to assess their feasibility before further development and integration into a pipeline comparing detected tracks with modeled radiation-dose relationships.

5. Conclusions

Morphological segmentation can partially separate overlapping tracks by exploiting the geometric properties of etched pits. However, these methods are sensitive to false positives caused by dirt, detector defects, and other track-like artifacts, and they may introduce over-segmentation in densely populated regions. On the other hand, deep learning approaches rely on large object detection or instance segmentation models designed for general visual recognition tasks. These architectures are typically optimized for scale invariance and diverse datasets, which does not fully match the constrained geometry and scale of CR39 tracks. Therefore, lighter and geometry-aware methods may be more suitable for this problem. Such approaches should better exploit the physical properties of etched tracks while reducing computational complexity and enabling more accurate analysis of radiation detector images.

Acknowledgements. This work has been supported by the Paul Scherrer Institute. It was on the Ares system at ACC Cyfronet AGH under the grants plgccbmc14.

References

1. Yuji Miao et al., Application of deep learning-based track detection technique in neutron personal dose monitoring research, *Radiation Measurements*, Vol. 189, 2025, 107526.
2. <https://github.com/CR39-analysis/RecoilNet>
3. Xu, Hongbo & Xiao, Feng & Yang, Xinyue & Zu, Chenxi & Mao, Xianfa & Luo, Shicheng & Liu, Jia & Liu, Yi & Luo, Cheng & You, Haoyu & You, Hao & Yuan, Hongzhi & Tan, Yanliang. (2024). Automatic counting method for alpha tracks on CR-39 detectors based on the YOLOv8 image recognition algorithm. 10.13140/RG.2.2.34051.82724

Radiative GRMHD Simulations of Puffy Accretion Discs: Numerical Modelling of Sub-Eddington Black Hole Accretion

Deborá Lančová^{1,2*}, Maciej Wielgus³, Marek Abramowicz^{1,2,4}, Agata Różańska¹, Gabriel Török², and Włodzimierz Kluźniak¹

¹ Nicolaus Copernicus Astronomical Centre, Polish Academy of Sciences, Warsaw, Poland

² Silesian University in Opava, Czech Republic

³ IAA-CSIC, Granada, Spain

⁴ Department of Physics, Göteborg University, Göteborg, Sweden

*debor.lancova@gmail.com

Keywords: accretion, accretion discs -- black hole physics -- MHD -- X-rays: binaries

1. Introduction

X-ray binaries containing black holes (BHs) are driven by matter under extreme conditions. The luminous spectral state is commonly interpreted using geometrically thin disc [1]. Despite observational success, these models are unstable in the radiation-pressure-dominated regime and neglect magnetic fields, which are expected to stabilise the disc. We simulated a BH accretion disc to study its behaviour and observable signatures and compare the results with predictions of analytical models.

2. Description of the problem

We study accretion flow onto a stellar mass BH using 3D general relativistic radiative magnetohydrodynamics (GRRMHD). The flow is modelled self-consistently, capturing magnetic turbulence and radiation-matter coupling. Reaching a quasi-steady state from a distant initial torus requires extremely long run times, making this problem tractable only with HPC resources. To obtain a quasi-stable state we average the simulation data over $\sim 20000 t_g$, where the dynamical timescale $t_g = GM/c^3$.

3. Related work

Analytical thin [1] and slim disc [2] models form the standard framework. Magnetic stabilization of radiation-dominated discs was demonstrated numerically by [3]. The puffy disc structure was first reported in [4] and the properties and observational signature analysed in [5] and [6].

4. Solution of the problem

Simulations were performed with the KORAL code [7], solving GRMHD and radiation transport equations using the M1 closure scheme on a static grid. Four simulations covering mildly sub-Eddington accretion rates were performed. KORAL solves a set of conservation equations for the gas coupled with radiation and the magnetic field, together with equation of state. The code is parallelised using MPI [8] and a few millions of core hours is required for one simulation.

5. Conclusions

The simulations produce magnetically stabilized puffy discs, geometrically and optically thick all the way to the BH horizon. The disc inner edge lies well inside the plunging region,

effective viscosity rises steeply in the innermost region, and the vertical structure resembles a warm corona. These properties differ significantly from analytical models and introduce systematic uncertainties in BH spin estimation from X-ray spectral fitting.




Fig. 1. The structure of puffy disc, with photosphere with black full line and the magnetic field-lines with white contours.

Acknowledgements. This work has been supported by GAČR grant 25-16928O (DL), NCN grants 2021/41/B/ST9/04110 (AR), and Ramón y Cajal grant RYC2023-042988-I (MW). The numerical experiment was possible through computing allocation on the Prometheus system at ACC Cyfronet AGH under the grants PLG/2025/018611.

References

1. Shakura & Sunyaev 1973, A&A, 500, 33.
2. Abramowicz et al. 1988, ApJ, 332, 646.
3. Sądowski 2016, MNRAS 459, 4397.
4. Lančová et al. 2019, ApJ 884, L37.
5. Wielgus et al. 2022, MNRAS 514, 780.
6. Lančová et al. in preparation.
7. Sądowski et al. 2013, MNRAS 429, 3533.
8. MPI Forum, 2021, MPI: A Message-Passing Interface Standard, Version 4.0, <https://www.mpi-forum.org/docs/>

Sovereign AI Chat Platform with RAG on PLGrid: Data Privacy, Polish Language Support, and Open-Source HPC Deployment

Mieszko Makuch, Tomasz Balawajder, Mieszko Cholewa, Patrycja Dąbrowska-Wierzbička,
Łukasz Flis, Joanna Kocot, Maciej Leśniak, Szymon Mazurek, Adam Mytnik,
Hubert Siejkowski, Jakub Sto, Bartłomiej Wenda

Academic Computer Centre Cyfronet AGH, ul. Nawojki 11, 30-950 Kraków, Poland

```
{mieszko.makuch, tomasz.balawajder, mieszko.cholewa,  
patrycja.dabrowska, lukasz.flis, joanna.kocot, maciej.lesniak,  
szymon.mazurek, adam.mytnik, hubert.siejkowski, jakub.sto,  
bartlomiej.wenda}@cyfronet.pl
```

Keywords: high-performance computing, large language models, data sovereignty, retrieval-augmented generation, Polish NLP, open-source, PLGrid

1. Introduction

Commercial AI chat platforms such as ChatGPT process user data on infrastructure outside the European Economic Area (EEA), where providers may retain or process conversations and uploaded documents to train future models. For small and medium enterprises (SMEs) handling proprietary or sensitive data, this poses unacceptable privacy and regulatory risks. We present a sovereign conversational AI platform with RAG, deployed entirely within the EEA on PLGrid HPC infrastructure and built on open-source components.

2. Description of the problem

Small and medium enterprises need a secure, ChatGPT-like environment where they can build custom knowledge bases from proprietary documents and query them using natural language. All system components: from LLM inference and embedding generation, through vector and document databases, file storage, to speech-to-text; must run on infrastructure within the EEA to ensure that no corporate data leaves a controlled, secure environment.

3. Related work

European sovereign AI initiatives such as OpenEuroLLM, Bielik [1] and EuroLLM [2] focus on training multilingual models on EEA infrastructure. However, a gap remains in providing organizations with a complete, privacy-preserving chat-and-RAG deployment stack built on these models.

4. Solution of the problem

Algorithm. The platform implements a two-phase RAG pipeline [3]: documents are split into chunks with configurable size and overlap, stored in PGVector. At query time, relevant fragments are retrieved via cosine similarity and appended to the LLM prompt.

Software. The platform is built on open-source components: MongoDB for persistence, MeiliSearch for search, and PGVector for embeddings. LLM Lab [5] exposes inference via an OpenAI-compatible API using vLLM [4], serving models including Bielik [1] for Polish.

Resources. All services are hosted on the infrastructure of ACK Cyfronet AGH, with LLM inference performed on NVIDIA GH200 superchips of the Helios supercomputer. We plan to expand to LUMI supercomputer resources for inference of state-of-the-art open-source LLMs.

5. Conclusions

All data and computation - conversations, documents, vector embeddings, embedding generation, and LLM inference - remain entirely within the EEA. No data is shared with third parties or used for model training, giving SMEs a safe, privacy-preserving AI platform with Polish language support.

Acknowledgements We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018674. Jointly funded by EuroHPC JU (GA 101234208) and Polish Ministry of Science and Higher Education (agreement MNiSW/2025/DIR/853).

References

1. SpeakLeash: Bielik Polish language model: <https://huggingface.co/speakleash>
2. OpenEuroLLM Consortium: EuroLLM: <https://openeurollm.eu>
3. P. Lewis et al.: Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In Proc. NeurIPS, 2020.
4. W. Kwon et al.: Efficient Memory Management for Large Language Model Serving with PagedAttention. In Proc. SOSP, 2023.
5. LLM Lab: PLGrid LLM inference platform: <https://llmlab.plgrid.pl>

Smart Data Management and ML-based Workflow Prediction for Cognitive Compute Continuum in SPICE

Michał Orzechowski^{1,2}, Piotr Kica¹, Mikołaj Zasada¹, Łukasz Opiola², Bartosz Kryza¹, Jakub Liput¹, Agnieszka Raczek¹, Jakub Karczewski¹, Wojciech Szmelich¹, Darin Nikolow¹, Łukasz Dutka², Bartosz Baliś¹, Renata G. Słota¹, Jacek Kitowski^{1,2}

¹AGH University of Krakow, Faculty of Computer Science, al. Mickiewicza 30, Kraków, Poland

²Academic Computer Centre Cyfronet AGH, Nawojki 11, Kraków, Poland

email: {morzech, rena, kito}@agh.edu.pl, lukasz.dutka@cyfronet.pl

Keywords: data management, metadata management, workflow processing, digitisation, compute continuum

1. Introduction

The growing scale of cloud computing has driven the emergence of edge and IoT computing, forming a „compute continuum” spanning from edge sensors to data centers and HPC facilities. Industry 4.0, smart energy, and medical imaging increasingly rely on this continuum for time-critical, AI-assisted processing. In manufacturing, laser cutting systems generate sensor data requiring edge-to-cloud pipelines with data anonymization; in energy, smart grids process large volumes of geospatial data across distributed infrastructure. SPICE (Smart Data Pipelines for the Cognitive Compute Continuum) [1] is a Horizon Europe collaboration of 10 partners from 4 EU countries, creating a platform for SMART data pipelines: Secure, Meaningful, Adaptable, Reliable, and Trustworthy. This paper presents AGH’s contributions: combining Onedata [2] with ML-based prediction for intelligent pipeline execution across the edge-cloud continuum.

2. Description of the problem

A critical challenge is efficiently managing distributed, multistage data pipelines across heterogeneous infrastructure with varying latency, bandwidth, and compute capabilities. These pipelines must handle massive data volumes, often hundreds of terabytes annually, while optimizing for conflicting objectives: throughput, cost, and data locality. Industrial contexts add further challenges: data sovereignty, cross-organizational sharing, and privacy compliance.

3. Related work

Onedata [2], developed by Cyfronet through EU-funded and national projects, provides a decentralized data layer for unified access to geographically distributed resources. It processes dataset fragments without full replication and implements trust-driven access control for secure cross-organizational sharing. Widely adopted in European research infrastructure, Onedata powers EGI DataHub, integrates with EOSC [3] services, and serves as the data backbone in EOSC Node Poland. LSTM networks [4], effective at time-series forecasting, are well-suited for predicting resource consumption in dynamic pipeline environments.

4. Solution of the problem

AGH contributes two complementary components: a smart data management layer based on Onedata and ML-based techniques for runtime optimization. A key innovation is smart data movement, using ML to anticipate data needs and position data along optimal transfer paths. AGH also develops an LLM-based discovery service providing semantic search, categorization,

and contextual analysis, enabling insight extraction from raw data regardless of expertise. For runtime optimization, AGH develops an ML-based pipeline predictor for dynamic workloads and resource uncertainties. It employs LSTM models to forecast workloads and resource consumption from historical and recent monitoring data. These predictions feed a cognitive mapper combining design-time global optimization with lightweight runtime heuristics. The mapper balances throughput and cost using strategies like Shortest Transmission Time First and Shortest Queue Length First. Continuous monitoring of transfer times, latency, and task invocations enables proactive SLO violation detection and schedule adaptation.

Both components facilitate the Cyfronet infrastructure, including the Athena and Ares supercomputers. Use cases include Industry 4.0 processing TB-scale sensor data, energy grids handling ~200TB/year of geospatial data, and medical imaging with whole-slide images exceeding 30GB each.

5. Conclusions

SPICE addresses managing data pipelines across the compute continuum. AGH contributes to Onedata for decentralized data management and ML-based prediction for runtime optimization. Together, these enable efficient large-scale distributed processing while maintaining security and European data space compliance. Future work focuses on full EOSC integration, broadening the platform's applicability across EU research infrastructure.

Acknowledgements. This work is co-founded within the European Union NextGenerationEU, Recovery and Resilience Plan under the SPICE project (09I02-03-V01-00012). MO, RS, and JK were partially supported by the Ministry of Education and Science funds assigned to AGH University. ŁO and ŁD are grateful for the support of the Polish Ministry of Education and Science under the program International Co-financed Projects (5399/DIGITAL/2023/2).

References

1. SPICE Project (09I02-03-V01-00012) - Smart Data Pipelines for the Cognitive Compute Continuum. <https://spice-project.eu/>
2. Orzechowski, M., Wrzeszcz, M., Kryza, B., Dutka, Ł., Słota, R.G., & Kitowski, J. (2023). Global data access to distributed data in a multi-cloud environment. *Future Generation Computer Systems*, 148, 150-159.
3. *EOSC Portal - European Open Science Cloud*. <https://eosc.pl/>
4. Hochreiter, S., & Schmidhuber, J. (1997). Long Short-Term Memory. *Neural Computation*, 9(8), 1735-1780.

Selecting Knowledge Representation Tools for a Cloud-Native Context

Szymon Głomski, Patryk Zajdel, Joanna Kosińska

AGH University of Krakow, al. Mickiewicza 30, 30-059 Kraków, Poland

{sglomski, zajdelp}@student.agh.edu.pl, kosinska@agh.edu.pl

Keywords: Cloud-native, semantics, Knowledge Representation, evaluation, tools

1. Introduction

Managing the scale and dynamism of cloud-native systems requires more than imperative scripts; it requires explicit, machine-readable knowledge representations (KRs). In this context, “knowledge” encompasses the structure of services and resources, the behavioral constraints and security policies, API interface semantics, and the temporal runtime context. Without a structured approach to KR, teams face fragmented visibility, inconsistent policy enforcement, and operational silos. This work explores KR tools families available for cloud-native systems.

2. Description of the problem

KR in cloud-native architectures is pragmatically defined as artifacts that enable automation, deployment, and reasoning. The core domains of cloud-native knowledge include:

- **Structure:** Definitions of microservices, infrastructure resources, dependencies, and domain boundaries.
- **Behavior and Constraints:** Security rules, admission policies, and resource placement regulations.
- **Interfaces:** Formalized communication agreements between microservices. This includes payload schemas, endpoint definitions, versioning strategies, semantics for synchronous (e.g. gRPC) and asynchronous interactions. Explicitly representing interfaces is critical for service decoupling, automated testing, and independent deployment.
- **State and Context:** Runtime interactions, telemetry data, and service affinities.

3. Related work

Existing literature on managing infrastructure complexity often points to Configuration Management Databases, which struggle to keep pace with the ephemeral and dynamic nature of microservices. While some approaches explore Semantic Web technologies [3] ontologies are often too heavy for agile platform engineering pipelines. Recent GitOps and Infrastructure-as-Code trends [2] focus mostly on deployment mechanics rather than formalized KR.

4. Solution of the problem

We categorize the solution space into four primary tool families, each addressing a specific domain of the knowledge problem. In Table 1 we map primary knowledge needs to corresponding tools. Furthermore, utilizing the infrastructure of Cyfronet, we subject these tools to rigorous benchmarking to assess their performance and usability under intensive loads.

Graph Databases: Neo4j and JanusGraph model service topologies, mapping dependencies for failure-impact and root-cause analysis by linking static models with runtime signals.

Policy-as-Code: Open Policy Agent (OPA) unifies policies across the stack using Rego, while Kyverno uses standard CRDs [1] to validate and mutate Kubernetes resources.

Application modeling: Crossplane and KubeVela bridge the gap between low-level infrastructure and high-level semantics. They allow teams to represent external cloud resources and application workflows as declarative artifacts, hiding underlying complexity.

Metadata ecosystems: Backstage serves as a central metadata hub, linking ownership, software entities, and API contracts into a searchable domain model. Configuration languages like CUE unify schema validation to prevent configuration drift.

Tab. 1. Tool selection based on primary knowledge need.

Primary knowledge need	Typical questions	Suitable tools
Topology	What depends on what? What breaks if X fails?	Neo4j, JanusGraph
Constraints & Policies	Is this deployment compliant? Who has access?	OPA/Rego, Kyverno
Application models	How is this app composed across environments?	Crossplane, KubeVela (OAM)
Configuration schemas	How can manifests be validated and reused?	CUE, Jsonnet, CRDs
Organizational metadata	Which team owns this service? Where is the API?	Backstage Catalog
API Contracts	What does this API do?	OpenAPI, AsyncAPI

5. Conclusions

Modern cloud-native systems require shifting from implicit knowledge to explicit, machine-readable representations. Our analysis demonstrates that no single tool can comprehensively cover all dimensions of cloud-native KR. Instead, platform engineering teams must adopt a composite, purpose-built toolchain, combining graph databases for topological visibility, policy engines for automated compliance, and metadata catalogs for organizational context. Ultimately, treating infrastructure, policies, and API contracts as formal knowledge artifacts is a foundational step toward building robust, autonomous, and self-healing internal developer platforms.

Acknowledgements. This research was funded in whole by National Science Centre, Poland [Grant number:2025/57/B/ST6/01885]. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission. The experiment was possible through computing allocation at ACC Cyfronet AGH under the grants [ACC Cyfronet for testing].

References

1. Kubernetes - custom resources. <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>. Accessed: 2026-02-18.
2. Abiola Ajayi, Oriyomi Badmus, Lucky Ehizojie, and Shokenu Segun. Comparative analysis of gitops tools and frameworks. *Path of Science*, 11:2009–2022, 05 2025.
3. J. Kosińska, G. Brotoń, and M. Tobiasz: Knowledge representation of the state of a cloud-native application. *International Journal on Software Tools for Technology Transfer*, 26(1):21–32, Feb 2024.

Towards Event-Based In-Cabin Sensing: a Synchronized Dual-Camera Dataset Pipeline

Marcin Kowalczyk, Marek Długosz, Kamil Jeziorek, Mateusz Komorkiewicz, Tomasz Kryjak, Dariusz Marchewka, Marcin Szelest, Piotr Wzorek, Paweł Skruch

Department of Automatic Control and Robotics, Faculty of Electrical Engineering, Automatics, Computer Science and Biomedical Engineering, AGH University of Krakow

{kowalczyk, mdlugosz, kjeziorek, komorkiewicz, kryjak, dmar, mszelest, pwzorek, pawel.skruch}@agh.edu.pl

Keywords: event camera, image processing, DVS, DMS, dataset

1. Introduction

Event cameras represent a paradigm shift in visual sensing; instead of capturing absolute brightness frame-by-frame, individual pixels operate independently, reacting only to localized changes in illumination. This unique mechanism unlocks microsecond precision, ultra-low latency, and an exceptional dynamic range. These qualities are highly beneficial for Driver Monitoring Systems (DMS) operating under unpredictable in-cabin lighting. However, because these sensors output a continuous, sparse spatio-temporal point cloud, conventional frame-based algorithms and existing automotive databases are rendered obsolete. To facilitate the development of robust models for driver state analysis, constructing a dedicated, event-driven dataset is a fundamental first step.




Fig. 1. Comparison of a frame from a classic camera and a corresponding frame created based on events. The detected and projected faces are marked.

2. Description of the problem

Developing robust DMS using event-based vision relies heavily on neural networks to analyze driver behavior, as deep learning allows to achieve high accuracy on continuous and sparse event data. The critical bottleneck in advancing these algorithms is the lack of extensive training materials. Designing and evaluating effective machine learning solutions requires processing large-scale datasets that accurately reflect the automotive environment. The main problem this work tackles is the preparation and structuring of a dedicated event-camera dataset, providing the necessary foundation to train and test future DMS algorithms.

3. Related work

Although event cameras offer significant advantages for Driver Monitoring Systems (DMS), their practical application is still emerging. Recent algorithmic advancements include yawn and seatbelt detection using synthetic data [1], spiking neural networks for drowsiness [2], and DMS-specific camera bias optimization [3]. Despite this progress, real-world datasets remain extremely scarce. Pioneering efforts like the EDDD [4] (drowsiness detection) and multi-modal NeuroIV [5] (gaze and gesture recognition) datasets demonstrate the viability of in-cabin event sensors. However, their major limitation is accessibility: EDDD is untraceable in open repositories, and NeuroIV is restricted by regional hosting. This lack of reproducibility highlights a critical need for a fully open event-based dataset to advance DMS research.

4. Solution of the problem

To address the lack of accessible event-based DMS data, we utilized a synchronized and calibrated dual-sensor setup comprising an event camera and a standard frame camera. Because annotating sparse event streams directly is inherently difficult, we adopted an automated cross-modal labeling strategy. Specifically, the InsightFace framework detects faces and facial landmarks on the standard video frames. Using calibration parameters, these precise annotations are then projected into the event camera's spatial domain. This pipeline relies heavily on High-Performance Computing (HPC). The Athena supercomputer executes the deep learning inference, while the Ares supercomputer handles data preprocessing and coordinate conversions, with all datasets maintained on dedicated HPC storage.

5. Conclusions

The research carried out has made it possible to design a comprehensive pipeline for acquiring and annotating a large-scale event-based dataset for Driver Monitoring Systems. In addition, the challenging and time-consuming process of event data labeling was successfully automated by utilizing cross-modal annotations from the synchronized frame camera. The preparation of this dataset recording pipeline would not have been possible without the use of the computing resources of the ACC Cyfronet AGH centre.

Acknowledgements. This work was supported by the research subsidy (funds from the Ministry of Science and Higher Education). We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACC Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2026/019101.

References

1. KIELTY, Paul, et al. "Neuromorphic driver monitoring systems: A proof-of-concept for yawn detection and seatbelt state detection using an event camera." *IEEE Access* 11 (2023): 96363-96373.
2. SHARIFF, Waseem, et al. "Spiking-DD: Neuromorphic event camera based driver distraction detection with spiking neural network." *IET Conference Proceedings CP887*. Vol. 2024. No. 10. Stevenage, UK: The Institution of Engineering and Technology, 2024.
3. DILMAGHANI, Mehdi Sefidgar, et al. "Optimization of event camera bias settings for a neuromorphic driver monitoring system." *IEEE Access* 12 (2024): 32959-32970.
4. CHEN, Guang, et al. "EDDD: Event-based drowsiness driving detection through facial motion analysis with neuromorphic vision sensor." *IEEE Sensors Journal* 20.11 (2020): 6170-6181.
5. CHEN, Guang, et al. "NeuroIV: Neuromorphic vision meets intelligent vehicle towards safe driving with a new database and baseline evaluations." *IEEE Transactions on Intelligent Transportation Systems* 23.2 (2020): 1171-1183.

Simulation of Texture Evolution in Rolled Cu with Neural Network Support

Bartosz Sułkowski^{1,2}

¹ Department of Materials Science and Non-Ferrous Metals Engineering, Faculty of Non-Ferrous Metals
AGH University of Krakow, Kraków, Poland

² ACC Cyfronet AGH, ul. Nawojki 11, 30-950 Kraków

bartosz.sulkowski@agh.edu.pl

Keywords: texture simulations, viscoplastic self-consistent model, neural networks

1. Introduction

Texture is one of the most important metallic materials feature which has effect on many properties [1,2]. For example it affects the Yield Strength (YS), Ultimate Tensile Strength (UTS) and plasticity of many rolled materials (Cu, Al, Fe) [3,4]. This is even more important for deformed hexagonal metals (Mg, Zn) with limited independent slip systems [5]. However, the texture measurements are time consuming and expensive having effect on the design of the final material. It would be very convenient to have a tool for texture prediction especially for Finite Element Method (FEM) simulations to design the best processing method in advance.

2. Description of the problem

Nowadays, the most important and well-known model for the texture simulations is the viscoplastic self-consistent model (vpssc) [1,2]. It gives very good results comparable with the experimental measurements. In the vpssc model the problem of texture evolution is based on solving the non-linear problem of finding the stress state in a deformed grain under known deformation gradient. However, to solve the problem the minimization mathematical equations have to be applied. This kind of approach has many downsides such as finding local minima, dependency of the calculation time on the distance between final solution and the first guess and poor scalability. This is impractical in the simulations with the FEM methods and prevent the use of the Graphics Processing Unit (GPU).

3. Related work

Neural networks are promising tools for improving calculations and finding solutions of difficult problems in many fields [6,7]. They were applied for texture simulations for prediction of cold rolling texture of steels with very good results [8]. However, that approach is base on training the network having known input parameters such as temperature, slip systems, deformation method etc. The problem is that the solution is very prone to the quality and number of the texture measurements. In the present study there is another approach, the data for the training NN will be generated using vpssc model. This will help to improve the quality and number of output data.

4. Solution of the problem

In the present study the neural network model was applied to reduce the calculation time and improve the simulations scalability of the texture evolution in rolled Cu. The neural network was composed with one input layer with 9 weights, 2 hidden layers and output layer with 9 weights. For all layers the sigmoid activation function was applied.

The aim of the present study was to model the rolling texture of Cu using NN model. The data such as deformation gradient and the stress state of 100 grains in 30 steps of simulations using vpssc model were generated to train and validate the outputs of the NN model. The deformation was equal to $\epsilon = 0,9$ at room temperature.




Fig. 1. (111) pole figures of rolled Cu (a), calculated with the use of the vpssc model (b) and calculated with the use of the NN model (c).

As it can be observed in Fig. 1, the results obtained from the simulations with the use of the NN model (Fig. 1c) is in agreement with the experimental texture of rolled Cu (Fig. 1a) and the simulations from the vpssc model. The main components of the Cu rolled texture measured by the experiments are $\{112\}\langle 111 \rangle$ (Copper – component) and $(110)\langle -112 \rangle$ (Brass component) with weak $(123)\langle 63-4 \rangle$ (S component) [9]. The vpssc model predicted correctly the main strong Copper component and the weak S. The NN model predicted correctly the main components (Copper and S) similar to vpssc model. However, there are some shifts of the main components from their ideal positions and the symmetry is somehow worse. Nevertheless, the results obtained from the NN simulations are in very good agreement with the experiments and vpssc simulations while the calculations time was strongly reduced and the scalability was improved.

5. Conclusions

In the present study the NN model was applied with success to predict the main texture components of the Cu rolled texture. The NN model correctly predicted Copper and S texture components, however there are shifts of the components from their ideal positions and the symmetry is worse. Calculations time was strongly reduced and the scalability was improved.

Acknowledgements. The numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grants plgcysoftware2026-cpu.

References

1. R.A. Lebensohn, C.N. Tomé, A self-consistent anisotropic approach for the simulation of plastic deformation and texture development of polycrystals: application to zirconium alloys, *Acta Metall. Mater.* 41 (1993) 2611–2624.
2. P. Van Houtte, A comprehensive mathematical formulation of an extended Taylor-Bishop-Hill model featuring relaxed constraints, the Renouard-Wintenberger theory and a strain rate sensitivity model, *Textures Microstruct.* 8 & 9 (1988) 313–350.
3. H.R. Piehler, *Crystal-Plasticity Fundamentals*, ASM Handbook Volume 22A: Fundamentals of Modeling for Metals Processing (#05215G), 2009.
4. Chapuis, Q. Liu, Modeling strain rate sensitivity and high temperature deformation of mg-3Al-1Zn alloy, *J. Magnes. Alloys* 7 (2019) 433–443, <https://doi.org/10.1016/j.jma.2019.04.004>
5. S.R. Agnew, M.H. Yoo, C.N. Tome, Application of texture simulation to understanding mechanical behavior of mg and solid solution alloys containing Li or Y, *Acta Mater.* 49 (2001) 4277–4289.
6. C.M. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, 1995-11-01.
7. R. Tadeusiewicz, *Sieci neuronowe*, Akademicka Oficyna Wydawnicza, 1993.
8. A. Brahme, M. Wining, D. Raabe, Prediction of cold rolling texture of steels using an Artificial Neural Network, *Computational Materials Science* Volume 46, Issue 4, October 2009, Pages 800-804.
9. H. Tian, H.L. Suo, O.V. Mishin, Annealing behaviour of a nanostructured Cu–45 at.%Ni alloy, *J Mater Sci* 48, 4183–4190 (2013). <https://doi.org/10.1007/s10853-013-7231-y>

Sampling Spin-Glass Systems with Transformers

Adam Stefański, Dawid Zapolski, Piotr Białas, Piotr Korcyl, Tomasz Stebel

Jagiellonian University, Kraków, 31-007 Poland

adam.stefanski@student.uj.edu.pl, dawid.zapolski@doctoral.uj.edu.pl,
piotr.bialas@uj.edu.pl, piotr.korcyl@uj.edu.pl, tomasz.stebel@uj.edu.pl

Keywords: Machine Learning, Variational Autoregressive Neural Networks, Transformers, Statistical Physics, Spin Systems, Markov Chain Monte Carlo

1. Introduction

We propose a way to adapt the transformer network architecture to sample spin systems. Our model expands the idea of Variational Autoregressive Networks (VAN) proposed in [1], enhancing their effectiveness with the usage of self-attention mechanism. The algorithm is called tVAN (transformer VAN). It allows sampling large systems of 64×64 spins.

In the spin-glass model, the interaction between neighboring spins $s^i = \pm 1$ on an $L \times L$ grid adds the term $-J_{i,j}s^i s^j$ to the energy of the system, with $J_{i,j} = \pm 1$.

2. Description of the problem

The aim of tVAN is sampling systems of spins from a given probability distribution and computing the free energy of such systems - a task hard to achieve using traditional Markov Chain Monte Carlo methods.

For spin-glass, the probability of a lattice configuration is given by the Boltzmann distribution:

$$p(\mathbf{s}) = \frac{1}{Z} \exp(\beta E(\mathbf{s})),$$

where β is the inverse temperature and Z is the partition function: $Z = \sum_{\mathbf{s}} \exp(-\beta E(\mathbf{s}))$ with the sum performed over all 2^{L^2} configurations. The free energy is defined as

$$F = -\frac{1}{\beta} \log Z.$$

3. Related work

Simulating spin systems with dense networks was proposed in [1], where the conditional probabilities of consecutive spins are modelled autoregressively, minimizing the variational free energy of the system. Another model proposed in [2] makes use of two body interactions between spins.

4. Solution of the problem

The algorithm is a physics-informed version of the well-known transformer architecture [3]. It consists of a sequence of connected blocks, each containing self-attention mechanism and feed-forward network.

Transformer generates spin lattice samples with *ansatz* joint probability $q_\theta(\mathbf{s})$, adjusting variational parameters θ by minimizing variational free energy

$$F_q = \frac{1}{\beta} \sum_{\mathbf{s}} q_\theta(\mathbf{s}) [\beta E(\mathbf{s}) + \ln q_\theta(\mathbf{s})]$$

to bring $q_\theta(\mathbf{s})$ close to the actual distribution $p(\mathbf{s})$ [1].

The algorithm is implemented in the Python language using PyTorch module and CUDA as the framework for tensor operations on Nvidia A-100 GPUs.

5. Conclusions

Our tVAN algorithm can effectively simulate spin-glass systems with 64×64 spins. This makes it a state of the art sampler of spin systems.

Acknowledgements. This research has been supported by the Polish National Science Centre (NCN) Grant No. 2021/43/D/ST2/03375. The numerical experiment was possible through computing allocation on the Athena system at ACC Cyfronet AGH under the grant no. PLG/2025/018811.

References

1. D. Wu, L. Wang, P. Zhang, Phys. Rev. Lett. 122 (2019) 080602.
2. I. Biazzo, Commun. Phys. 6 (2023).
3. A. Vaswani et al., in: Adv. Neural Inf. Process. Syst., volume 30, Curran Associates, Inc., (2017).

Transformer as an Efficient Sampler for the Ising Model

Dawid Zapolski^{1,3}, Piotr Białas¹, Piotr Korczyk², Tomasz Stebel²

¹ Institute of Applied Computer Science,

² Institute of Theoretical Physics,

³ Doctoral School of Exact and Natural Sciences,

Jagiellonian University, ul. Łojasiewicza 11, 30-348 Kraków, Poland

dawid.zapolski@doctoral.uj.edu.pl

Keywords: machine learning, statistical physics, Markov chain Monte Carlo, transformer, spin systems

1. Introduction

Generative Transformers have revolutionized scientific computing, yet validating their output fidelity at scale remains challenging. Unlike natural language processing (NLP) where perplexity proxies quality, physical simulations require rigorous guarantees. We present a framework using the 2D Ising model as a “white-box” stress test for autoregressive generation, providing a quantitative fidelity metric via the Effective Sample Size (ESS).

2. Description of the problem

The 2D Ising model is a fundamental system in statistical physics consisting of a lattice of binary variables, or “spins”. Generating independent spin configurations with Markov chain Monte Carlo (MCMC) techniques near phase transitions suffers from critical slowing down. Variational Autoregressive Networks (VANs) overcome this, but introduce severe inference latency as one must sequentially generate L^2 tokens for an $L \times L$ lattice.

3. Related work

Reference [1] introduced the VAN method for statistical physics, introducing a self-supervised training scheme that requires no dataset, optimizing directly against the target Boltzmann distribution. Other works such as [2], [3] build on that including physics-inspired improvements.

4. Solution of the problem

We propose the single- or double-layer transformer (tVAN) working with $(a \times b)$ patches (up to 3×4) instead of the single spins. This architectural changes together with KV cache speed up generation and training time by a few orders of magnitude comparing to the single spins, and giving results better than [2] and [3].

Tab. 1. ESS for our model as well as [2], [3] at critical temperature for $L=128$. $ESS \sim 1$ indicates the learned distribution closely matches the target.

Model name	1-layer tVAN	2-layer tVAN	HAN [2]	RiGCS [3]
ESS	0.4	0.7	2×10^{-5}	0.02

We train our models, written with PyTorch library, using up to 4 AMD MI250x GPU modules via data parallelism, as the self-supervised training requires a batch size too large to fit a single GPU.

5. Conclusions

We demonstrate that the patch-based transformer can work as an efficient sampler for the Ising model up to 128×128 system size. By providing rigorous convergence metrics absent in NLP benchmarks, this framework offers a verifiable stress test for long-context inference optimizations.

Acknowledgements. We acknowledge Polish high-performance computing infrastructure PLGrid for awarding this project access to the LUMI supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CSC (Finland) and the LUMI consortium through PLL/2025/08/018112. T.S. and D.Z. acknowledge the support of the Polish National Science Center (NCN) Grant No. 2021/43/D/ST2/03375. P.K. acknowledges the support of the Polish National Science Center (NCN) grant No. 2022/46/E/ST2/00346.

References

1. Wu, Dian, Lei Wang, and Pan Zhang. "Solving statistical mechanics using variational autoregressive networks." *Physical review letters* 122.8 (2019): 080602.
2. Białas, Piotr, Piotr Korcyl, and Tomasz Stebel. „Hierarchical autoregressive neural networks for statistical systems." *Computer Physics Communications* 281 (2022): 108502.
3. Singha, Ankur, et al. "Multilevel Generative Samplers for Investigating Critical Phenomena." *International Conference on Learning Representations*, edited by Y. Yue et al., vol. 2025, (2025), pp. 25450–25483.

A More Noise Resilient Search Space Method for the Quantum Approximate Optimization Algorithm

Arkadiusz Wołk^{1,2}, Karol Capała^{1,2}, Katarzyna Rycerz^{1,2}

¹ AGH University of Krakow, Faculty of Computer Science, Mickiewicza 30, Kraków, 30-059, Poland

² Academic Computer Center Cyfronet AGH, Nawojki 11, Kraków, 30-950, Poland

{awolk, capala, kzajac}@agh.edu.pl

Keywords: Quantum approximate optimization algorithm, Constrained optimization, Noisy intermediate-scale quantum, Quantum noise robustness, Hypercube operator

1. Introduction

Combinatorial problems are often encountered, with many classified as NP-complete. The emergence of quantum computing lead to the development of approximate algorithms that offer potential scalability improvements. Given the limited number of qubits and inherent noise in quantum devices, circuit size significantly affects solution accuracy. The Quantum Approximate Optimization Algorithm (QAOA) is a prominent method for deriving approximate solutions to unconstrained combinatorial problems.

2. Description of the problem

The QAOA is effective for unconstrained problems. Its part responsible for searching the solution space, called mixer, is simple in its standard form and unnecessarily allows for including infeasible answers. Most methods in the literature do not guarantee constraint satisfaction, but only reduce the probability of violations. However, there are also advanced types of mixers that reduce of the search space to feasible solutions. Unfortunately, they require a larger circuit, increasing vulnerability to noise. Minimizing circuit size is essential to enhance accuracy on noisy quantum hardware.

3. Related work

The baseline for our work, a mixer searching solutions on the hypercube graph, was introduced in the original work [1]. Its standard implementation approach is described in [2].

4. Solution of the problem

In the standard implementation, the constraint check is repeated multiple times, constituting the largest portion of the circuit. However, the hypercube structure ensures that only one of the n variables changes at each step. For constraints defined by linear functions of n variables, function values change solely by the coefficient corresponding to the modified variable. Based on this observation, a modified method was developed and analytically demonstrated that five variables represent the upper bound for the standard approach, potentially yielding a smaller circuit, as shown in Tab. 1. Additionally, experimental results indicate improved noise resilience for this method (Fig. 1).

Tab. 1. Sample problems and generated circuit sizes.

Problem no.	Constraints	Gate count		Difference
		Modified [r=7]	Standard [r=7]	
1	$3 \leq 2x_0 + 3x_1 + 2x_2 + x_3 + 4x_4 + x_5 \leq 11$	116193	166694	-30%
2	$4 \leq 7x_0 + 5x_1 + 3x_2 + 8x_3 + 6x_4 + x_5 \leq 20$	164631	230818	-29%

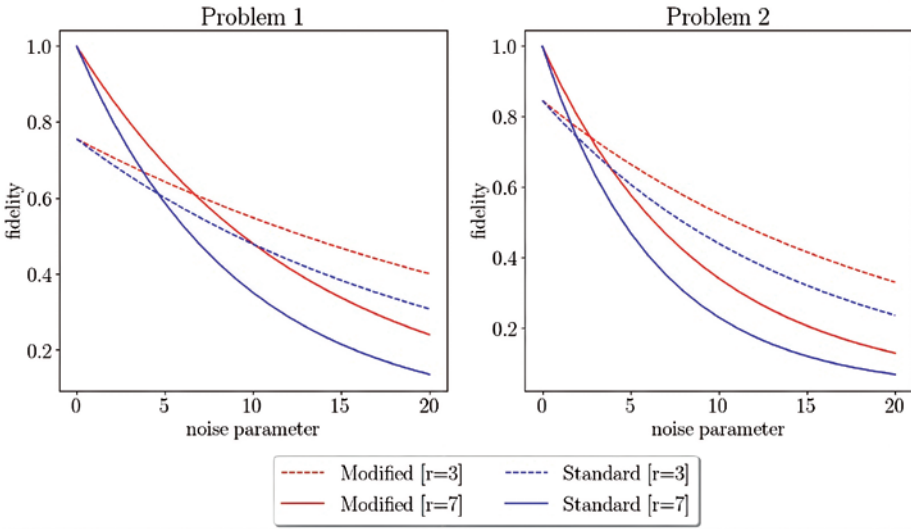


Fig. 1. Simulation results for problems in Tab. 1 subject to phase- and amplitude damping noise.

5. Conclusions

We have proposed a modification of the constrained hypercube mixer method that requires fewer gates and exhibits better noise robustness. Thus, it improves the accuracy of a method on quantum hardware.

Acknowledgements. The numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grant no. PLG/2025/018379.

References

1. E. Farhi, J. Goldstone, S. Gutmann, A quantum approximate optimization algorithm (2014). arXiv:1411.4028.
2. S. Marsh, J. B. Wang, A quantum walk-assisted approximate algorithm for bounded NP optimisation problems, Quantum Information Processing 18 (3) (2019) 61. doi:10.1007/s11128-019-2171-3.

Kinetic Simulations of Nonresonant Instability in the Precursor of Young Supernova Remnant Shock Containing Positrons

Oleh Kobzar

Astronomical Observatory of the Jagiellonian University, ul. Orla 171, 30-244 Kraków, Poland

oleh.kobzar@uj.edu.pl

Keywords: shock waves, plasmas, magnetic turbulence, particle acceleration, PIC simulations

1. Introduction

Shock waves at supernova remnants (SNRs) are typically considered as sources of Galactic cosmic rays (CRs). Recent observations indicated the presence of leptons in their content with energies up to ~ 1 TeV [1]. Besides, positron component demonstrates an excess with harder spectral indexes in the range of $\sim 20 - 300$ GeV. Currently, there is no generally accepted theory explaining the origin of positrons in CR flux and peculiarities of their energy distribution. This problem goes beyond of astrophysics and touches a fundamental understanding of nature.

2. Numerical simulations as a key solution

Since the astrophysical plasma is strongly nonlinear, it can be studied with numerical methods, typically requiring high-performance computing with thousands of CPUs and up to tens of TB of RAM. We use kinetic Particle-In-Cell (PIC) simulations, a self-consistent *ab-initio* method developed for collisionless plasma. It solves Maxwell's equations for EM-fields on a spatial grid and follows the motion of individual particles with finite time steps. The simulations have been realized with the use of advanced MPI-based parallel code TRISTAN [2, 3].

3. Related work

Our recent pilot study has demonstrated, that (pre-)acceleration of leptons through their interaction with magnetic turbulence in SNR shock precursor is charge-dependent: positrons are energized much more efficiently than electrons [4]. This is consistent with other investigations, also showing preferential positron acceleration at shocks compared to electrons [5].

4. Simulation setup and results

The results of our pilot study [4] were obtained with use of strongly reduced numerical parameters: 9 particles per cell for each specie, ion to electron mass ratio $m_i/m_e = 50$, and space resolution measured as electron skin depth per grid cell $\lambda_{se}/\Delta = 4$. Current investigation has been performed with better statistics using 16 particles per cell, increased mass ratio $m_i/m_e = 100$, and higher resolution $\lambda_{se}/\Delta = 10$, that results in significantly larger computational resources required. The main parameters of different simulation runs are listed in Table 1.

We performed the series of PIC-simulations using so called 2D-3V setup, where particles are positioned in the rectangular computational box only with 2 spatial coordinates, but all 3 components are treated for their velocities and fields. The ambient plasma is modeled to consist of protons and electrons (p, e⁻), as well as include small mixture of positrons varying from 1% to 5% of protons number (p, e⁻, e⁺). In all cases the monoenergetic population of CRs (2%) drifts through ambient plasma, being the source of nonresonant Bell's instability.

As a result of instability, circularly polarized magnetic waves are generated (Fig. 1). In the first case (p, e⁻) they are presented with only long-wave mode. In the second case (p, e⁻, e⁺) also

the short-wave mode appears, which has the opposite direction of polarization, as it follows from Fourier plots in Fig. 2. It remains unchanged in the range of positron number from 1% to 5%.

Tab. 1. List of the simulation parameters.

Setup	p, e ⁻ (small)	p, e ⁻ , e ⁺ (small)	p, e ⁻ , e ⁺ (large)
Box size (cells), $L_x \times L_y$	9 600 x 4 800	9 600 x 4 800	24 000 x 12 000
CPUs per simulation	960	960	2 400
CPU-hours used	~ 200 000	~ 200 000	~ 500 000
Disk space used	~ 5.7 TB	~ 7.5 TB	~ 46.5 TB




Fig. 1. Transverse MF resulting from nonresonant instability driven by CR current in: (a) p, e⁻ ambient plasma, and (b) p, e⁻, e⁺ ambient plasma containing small 1% mixture of positrons.




Fig. 2. Complex Fourier images of the MF waves shown in Fig. 1 (a) and (b), correspondingly.

5. Conclusions

Besides much more efficient (pre-)acceleration compared to electrons [4,5], positrons are found to modify the wave spectrum of the magnetic turbulence generated through CR current driven nonresonant instability in the precursor of SNR shock. It can influence the particle scattering processes, modifying in such a way the following stages of acceleration. The effect is kept almost unchanged in the wide range of positron concentrations.

Acknowledgements. The numerical experiment was possible through computing allocation on the Helios system at ACC Cyfronet AGH under the grant “heaptic25” (PLG/2025/018353).

References

1. Z. L. Weng, V. Vagelli, AMS Collaboration. Nuclear and Particle Physics Proceedings. v. 273–275, April–June 2016, p. 466–472.
2. J. Niemiec, M. Pohl, T. Stroman, K.-I. Nishikawa, 2008, *Astrophys. J.*, v. 684, p. 1174.
3. A. Dorobisz, M. Kotwica, J. Niemiec, O. Kobzar, A. Bohdan, K. Wiatr, 2018, *LNSC*, 10777, p. 156.
4. O. Kobzar, 2023, *KUKDM 2023 Proceedings*, p. 39.
5. H. Yu, Q. Xia, J. Fang, 2025, *Astrophys. J.*, v. 993, p. 100.

Towards Fast, Machine-Learning-Based Calorimeter Simulations in the ALICE Experiment at CERN

Emilia Majerz¹, Łukasz Dubiel¹, Jacek Otwinowski², Witold Dzwiniel¹, Jacek Kitowski¹

¹ AGH University of Krakow

² Institute of Nuclear Physics PAS in Kraków

majerz@agh.edu.pl, lukdubiel@student.agh.edu.pl

Keywords: high energy physics, fast simulation, generative neural networks, calorimeter simulation, zero degree calorimeter, forward calorimeter

1. Introduction

The rapid progress of artificial intelligence (AI) in recent years has created opportunities to integrate AI models into CERN software. In response, the AGH University of Krakow formed a team developing AI methods for the ALICE experiment. Current efforts focus on detector simulation, data reconstruction, and anomaly detection. Here, we concentrate on detector simulation, assessing whether generative neural networks can emulate calorimeter responses resulting from particle transport.

2. Description of the problem

Detector responses at CERN are traditionally simulated using the GEANT4 software. While accurate, these simulations are computationally intensive, limiting the scope of feasible studies. Machine learning addresses this: well-trained fast generative models can effectively emulate detector responses originating from particle feature vectors produced in beam-beam collisions. Our studies target two ALICE calorimeters: the Zero Degree Calorimeter (ZDC) and the hadronic section of the Forward Calorimeter (FoCal-H). Example detector outputs (our targets for generative modelling) are presented in Fig. 1.




Fig. 1. Representative visualisations of FoCal-H outputs simulated with GEANT4 (target data).

3. Related work

State-of-the-art (SOTA) generative frameworks, including Variational Autoencoders, Generative Adversarial Networks, Normalizing Flows, diffusion, and flow-matching models, have shown promise for calorimeter fast simulation, e.g. [1]. The rise of physics-based deep learning offers another avenue, emphasising physical consistency in simulations.

4. Solution of the problem

In our group’s preliminary ZDC study [2], we tested SOTA models and found that the architecture choice depends on the application: some offered speed at the cost of accuracy, others the reverse. Even the most accurate models, however, violated physical constraints without targeted optimisation. We thus developed a physics-based loss function which incorporated physical detector geometry, and combined with a response variability-induced scaling, improved the spatial distribution and morphology of generated particle showers [3].

For FoCal-H, we implemented SOTA models with an optional physics-based loss function that directly optimised particle shower positions [4]. We observed similar outcomes to the ZDC case, with a clear quality-speed tradeoff. Surprisingly, the physics-based loss did not improve the already satisfactory baseline results. We attribute this to the straightforward input parameters-shower position dependence. Added as a separate loss, it only made the process more complicated in terms of finding optimal hyperparameters for balancing loss functions.

In our most recent works [3, 4], we primarily used TensorFlow, PyTorch, NumPy, scikit-learn, and Uproot libraries, along with the Weights & Biases platform. Computations were performed on the Athena (NVIDIA A100-SXM4-40GB GPUs) and Ares (NVIDIA Tesla V100-SXM2 GPUs) clusters.

5. Conclusions

Our experiments demonstrate promising results for generative frameworks in calorimeter response simulation. In our most advanced ZDC study using Normalizing Flows with a custom physical loss, we achieved low generation time (0.38 ms per sample) while maintaining high fidelity to reference data [3]. Building on this success, we plan to prioritize a similar physical-constrained approach for FoCal-H, alongside ongoing work exploring contrastive, multi-layer, and generative approaches for muon track matching in reconstruction.

Acknowledgements. This work is co-financed and in part supported by the Ministry of Science and Higher Education (Agreements No. 2022/WK/01 and 2023/WK/07) by the program entitled “PMW” and by the Ministry funds assigned to AGH University of Krakow. The numerical experiment was possible through computing allocation on the Athena and Ares systems at ACC Cyfronet AGH under the grants PLG/2024/017264 and PLG/2025/018322.

References

1. Krause, C., Faucci Giannelli, M., Kasieczka, G., Nachman, B., Salamani, D., Shih, D., ... & Zhang, R. (2025). *CaloChallenge 2022: a community challenge for fast calorimeter simulation*. Reports on Progress in Physics, 88(11), 116201.
2. Wojnar, M., Majerz, E., & Dzwiniel, W. (2025). *Fast simulation of the Zero Degree Calorimeter responses with generative neural networks*. Computing and Software for Big Science, 9(1), 1.
3. Majerz, E., Dzwiniel, W., & Kitowski, J. (2025). *Inverse Autoregressive Flows for Zero Degree Calorimeter fast simulation*. 39th Annual Conference on Neural Information Processing Systems: Includes Machine Learning and the Physical Sciences (ML4PS).
4. Dubiel, Ł. (2025). *Improving the efficiency and precision of data-driven surrogates of Monte-Carlo models in the ALICE experiment* [Master’s thesis, AGH University of Krakow].

General-relativistic Magneto-hydro-dynamical Simulations of Relativistic Jets from Accreting Spinning Black Holes

Krzysztof Nalewajko

Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences,
Bartycka 18, 00-716 Warsaw, Poland

knalew@camk.edu.pl

Keywords: general relativity, magneto-hydro-dynamics, relativistic jets, black holes

1. Introduction

Various highly energetic astrophysical phenomena – active galaxies (e.g., quasars, blazars, radio galaxies), gamma-ray bursts (exploding massive stars or merging neutron stars), X-ray stellar binaries – are inferred to involve black holes. It is accepted that the activity of black holes is due to accretion of magnetised plasma. This process can be studied numerically by solving the conservation laws of magneto-hydro-dynamics (MHD) in the space-time described by general relativity (GR) – GRMHD.

2. Description of the problem

Formation of relativistic jets from black holes is governed by a mechanism that requires a spinning black hole that collects a large magnetic flux [1]. Magnetic flux accumulates across the black hole horizon during accretion of magnetised plasma, and accretion needs to be maintained in order to prevent shedding of the magnetic flux by means of magnetic reconnection. Magnetic flux accumulated on a black hole is subject to a saturation mechanism [2] – once the flux becomes large relative to the accretion rate, the magnetic pressure interrupts the accretion flow and triggers a magnetic flux eruption driven by reconnection.

3. Related work

First 3D GRMHD simulations demonstrated that the black hole magnetic flux is subject to saturation via black hole flux eruptions, this also corresponds to the most powerful relativistic jets [3]. GPU-accelerated GRMHD codes allowed to increase the numerical resolutions to $\sim 5000 \times 2000^2$ cells, and to recognise the fundamental role of magnetic reconnection in driving the black hole flux eruptions [5].

4. Solution of the problem

My team [plggknalew](https://github.com/plggknalew) uses different GRMHD codes for CPUs (Athena++) and GPUs (AthenaK, KHARMA, H-AMR). To investigate the onset mechanism of black hole flux eruptions [6], we used the public GRMHD code Athena++ [4] that we ran on Ares using 768 cores (16 nodes). We set up a 3D problem in quasi-spherical Kerr-Schild coordinates extending into the horizon with effective resolution of 288×256^2 cells and 3 levels of static mesh refinement. The two main simulations (cases of prograde and retrograde accretion on black hole spinning with $a = 0.9$) were run until at least 30k black hole light-crossing times. Using the 3D data on magnetic fields we integrated large samples of field lines and visualised them in 3D according to their connection with the horizon. This procedure was repeated for hundreds moments of time to achieve dynamic animations of the onset of magnetic flux eruption (<https://users.camk.edu.pl/knalew/aa50490-24/>). Figure 1 shows an example.

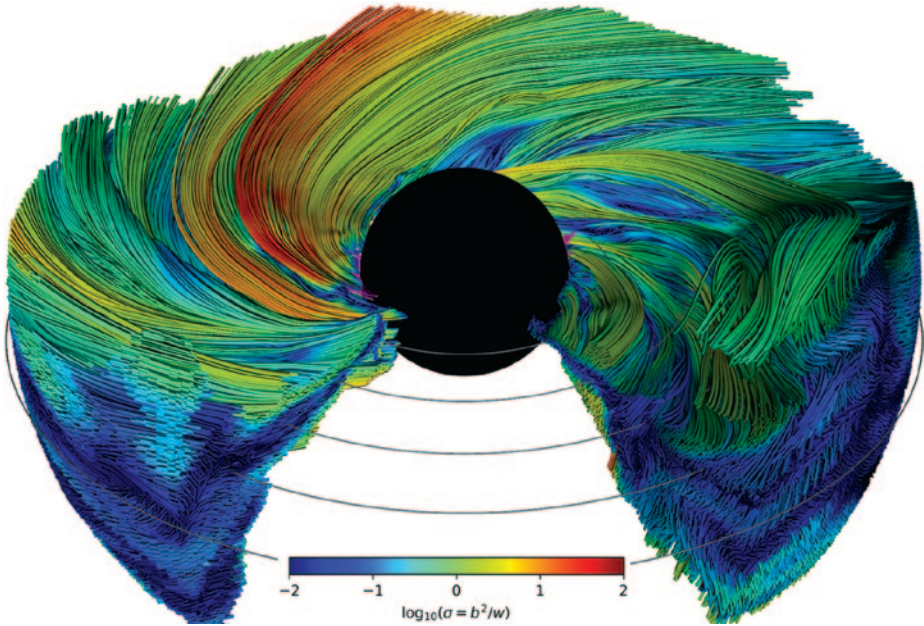


Fig. 1. Sample of magnetic field lines disconnected from the black hole horizon, coloured by plasma magnetisation. This represents an accretion flow constricted to a narrow wedge between the bases of two relativistic jets. Small magenta patches on the left side mark the onset of relativistic reconnection that is going to develop into a magnetic flux eruption [6]. This 3D rendering was made with Mayavi.

5. Conclusions

In-depth analysis of the results of 3D GRMHD simulations of magnetically saturated accretion flow onto a black hole allowed to explain the onset of magnetic flux eruption [6].

Acknowledgements. This research has been supported by the Polish National Science Centre grants 2021/41/B/ST9/04306 and 2024/53/B/ST9/03747. This numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grants PLG/2023/016444 and PLG/2024/017013.

References

1. R. D. Blandford & R. L. Znajek: „Electromagnetic extraction of energy from Kerr black holes”, 1977, MNRAS, 179, 433.
2. R. Narayan, I. V. Igumenshchev & M. A. Abramowicz: „Magnetically Arrested Disk: an Energetically Efficient Accretion Flow”, 2003, PASJ, 55, L69.
3. A. Tchekhovskoy, R. Narayan & J. C. McKinney: „Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole”, 2011, MNRAS, 418, L79.
4. C. J. White, J. M. Stone & C. F. Gammie: „An Extension of the Athena++ Code Framework for GRMHD Based on Advanced Riemann Solvers and Staggered-mesh ...”, 2016, ApJS, 225, 22.
5. B. Ripperda, M. Liska, K. Chatterjee et al.: „Black Hole Flares: Ejection of Accreted Magnetic Flux through 3D Plasmoid-mediated Reconnection”, 2022, ApJ, 924, L32.
6. K. Nalewajko, M. Kapusta & A. Janiuk: „Chaotic magnetic disconnections trigger flux eruptions in accretion flows channeled onto magnetically saturated Kerr black holes”, 2024, A&A, 692, A37.

VERONA: a GPU-Accelerated Special Relativistic Hydrodynamics (SRHD) Code for Astrophysical Applications

Piotr Płonka¹, Mateusz Kapusta², Agnieszka Janiuk¹

¹ Center for Theoretical Physics, Al. Lotników 32/46, 02-668 Warsaw, Poland

² Astronomical Observatory, University of Warsaw, Al. Ujazdowskie 4, 00-478 Warsaw, Poland

{pplonka, agnes}@cft.edu.pl, mr.kapusta@student.uw.edu.pl

Keywords: computational astrophysics, GPU acceleration, hydrodynamics

1. Introduction

In the era of rapid development of GPU computing technology, many computational astrophysics simulations can be significantly accelerated. We present *VERONA*, a GPU-accelerated, CUDA-aware MPI three-dimensional special relativistic hydrodynamics (SRHD) code in the *Julia* programming language with `KernelAbstractions.jl`, allowing for portability for different GPU and CPU architectures [1]. Our code has performance comparable to analogous codes in other programming languages, while offering improved usability and maintainability, thus reducing the barrier to entry for high-performance scientific computing.

2. Description of the problem

A wide range of extreme astrophysical phenomena in the universe can be studied using hydrodynamic models. To investigate these phenomena, we need to solve numerically the equations of relativistic hydrodynamics, i.e., the continuity equation and the covariant conservation of the energy-momentum tensor, supplemented by an equation of state [2].

3. Related work

In recent years, many numerical hydrodynamics codes have been developed for large-scale calculations on graphics cards, for example, *AthenaK* [3], *H-ARM* [4] or *AsterX* [5]. All the above codes use the programming language *C* or *C++*. For our code, we chose the *Julia* programming language.

4. Solution of the problem

The numerical scheme is conservative and shock-capturing, using the HLLE (Harten-Lax-van Leer-Einfeldt) Riemann solver with higher-order reconstruction methods, including PPM and WENO-Z. The recovery of primitive variables is performed using one-dimensional Brent's root-finding method. We validate our code using standard test problems such as relativistic shock tubes, Kelvin-Helmholtz instabilities, and two-dimensional Riemann problems, comparing the results with analytical solutions. We investigate the performance and scalability of the code in a multi-GPU environment on the Athena supercomputer using NVIDIA A100 GPUs, presenting strong and weak scaling tests. Finally, we illustrate a potential astrophysical application by simulating a highly relativistic jet propagating through a star in the process known as jet-breakout (see Fig. 1). This scenario is related to the collapse of massive stars at the end of their lives and is associated with long gamma-ray bursts, which are the most powerful explosions in the universe.

5. Conclusions

We have presented *VERONA*, a Julia-based, GPU-accelerated three-dimensional SRHD code for large-scale GPU simulations. The performance in multi-GPU computations on NVIDIA A100 hardware demonstrates good strong and weak scaling properties. The combination of CUDA-aware MPI parallelization and a high-level language shows that *Julia* can be an alternative to traditional *C/C++* implementations in computational astrophysics.




Fig. 1. A relativistic jet propagates through a star. After crossing the star, the jet breaks out into the surrounding space, and the ambient material forms a cocoon around the jet.

Acknowledgements. This work has been partially supported by grant No. 2023/50/A/ST9/00527 from the Polish National Science Center. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018186.

References

1. M. Kapusta and P. Płonka: *VERONA*, GitHub repository, 2026, github.com/Wesenheit/Verona.
2. L. Rezzolla and O. Zanotti: *Relativistic Hydrodynamics*, Oxford University Press, Oxford, 2013.
3. J.M. Stone et al.: *AthenaK: A Performance-Portable Version of the Athena++ AMR Framework*, arXiv e-prints, arXiv:2409.16053, 2024.
4. M.T.P. Liska et al.: *H-AMR: A New GPU-accelerated GRMHD Code for Exascale Computing with 3D Adaptive Mesh Refinement and Local Adaptive Time Stepping*. *Astrophysical Journal Supplement Series*, 263(2), 26, 2022.
5. J.V. Kalinani et al.: *AsterX: a new open-source GPU-accelerated GRMHD code for dynamical spacetimes*. *Classical and Quantum Gravity*, 42(2), 025016, 2025.

Finite Element Ion Recombination Simulations for Dosimetry and FLASH on PLGrid

Radosław Benedyckiński¹, Jakub Kaliński¹, Leszek Grzanka¹, Jeppe Brage Christensen²

¹ AGH University of Krakow, al. Mickiewicza 30, 30-059 Kraków, Poland

² Paul Scherrer Institute, Forschungsstrasse 111, 5232 Villigen PSI, Switzerland

{jkalinski, radbene}@student.agh.edu.pl, grzanka@agh.edu.pl,
jeppe.christensen@psi.ch

Keywords: ion recombination, ionization chamber, dosimetry, FLASH, finite element method, FEniCSx, MPI, SLURM, PLGrid, Ares, memfs

1. Introduction

Ionization chambers are primary reference detectors in clinical dosimetry. However, recombination of positive and negative charge carriers reduces the collected charge and may lead to dose underestimation. This effect becomes particularly significant for ultra-high dose rate (FLASH) beams, where classical analytical correction models may operate beyond their range of validity. We present a finite element method (FEM) approach that avoids these analytical approximations and supports complex chamber geometries as well as non-uniform ionization patterns.

2. Description of the problem

We solve coupled 3D transport-recombination partial differential equations, including electric fields and electrode boundary conditions. The recombination correction is determined from the collected charge fraction. Realistic meshes, time stepping, and the need to explore multiple chamber and beam conditions make these computations expensive. In our workflow, a typical production case used for experimental comparison requires between 1 and 5 GB of Storage. Depending on the problem scale (e.g. the bigger the values of: chamber radius, dose, time step, mesh density the longer it takes to finish), execution on a single PLGrid Ares node using 40 MPI ranks ranges from 30 minutes to several dozen hours per experiment. Systematic multi-scenario studies therefore quickly exceed the practical capabilities of a standard workstation.

3. Related work

Classical analytical ion recombination corrections (e.g., Boag-type models and voltage-based methods, including the two-voltage method) are widely used in routine dosimetry, but their assumptions may break down for ultra-short, high dose-per-pulse beams (FLASH) and more complex chamber geometries [2,3,4,10]. Recent studies propose refined models for high dose-per-pulse beams and specific chambers, and analytical extensions accounting for free-electron effects [5,11,12]. However, such approaches still have limited ability to capture fully three-dimensional, geometry-dependent field perturbations, motivating numerical 3D methods.

4. Solution of the problem

IonTracks-FEniCSx implements the finite element method using FEniCSx/DOLFINx with a PETSc backend and MPI parallelization. Calculations are performed on the PLGrid Ares clustered under SLURM, with large parameter scans managed through SLURM job arrays using multiple YAML configuration files. For reliable I/O operations, we use a RAM-backed scratch

filesystem (MEMFS - at least ~40 GB). To speed up the development and debugging process, we use a remote workflow (VS Code Remote running as a SLURM job through an SSH tunnel). Environment reproducibility is ensured through Conda and Docker containers.

5. Conclusions

Access to HPC infrastructure is a practical requirement for this class of 3D FEM recombination simulations. It turns single runs into a repeatable workflow and enables systematic parameter scans and validation cycles relevant to FLASH dosimetry.

Acknowledgements. This work was supported by the PLGrid infrastructure. Numerical experiments were performed on the Ares system at ACC Cyfronet AGH under the grant plgcbmc14.

References

1. G. Jaffé, *Ann. Phys.*, 347(12), 303-344, 1913.
2. J. W. Boag, *Br. J. Radiol.*, 23(274), 601-611, 1950.
3. K. Petersson et al., High dose-per-pulse electron beam dosimetry - A model to correct for the ion recombination in the Advanced Markus ionization chamber, *Med. Phys.*, 44(3), 2017, doi: 10.1002/mp.12111
4. M. McManus et al., The challenge of ionisation chamber dosimetry in ultra-short pulsed high dose-rate Very High Energy Electron beams, *Sci. Rep.*, 10, 9089, 2020, doi: 10.1038/s41598-020-65819-y
5. J. D. Fenwick and S. Kumar, Collection efficiencies of ionization chambers in pulsed radiation beams: an exact solution of an ion recombination model including free electron effects, *Phys. Med. Biol.*, 68, 2023, doi: 10.1088/1361-6560/aca74e
6. M. S. Alnæs et al., The FEniCS Project Version 1.5, *Archive of Numerical Software*, 3(100), 2015, doi: 10.11588/ans.2015.100.20553
7. S. Balay et al., Efficient Management of Parallelism in Object-Oriented Numerical Software Libraries, in *Modern Software Tools in Scientific Computing*, 1997.
8. Ion recombination measurements at the Danish Center for Particle Therapy (DCPT), publication: <https://iopscience.iop.org/article/10.1088/1361-6560/ab8579/meta>
9. IonTracks-FEniCSx source code: <https://github.com/ion-tracks/IonTracks-FEniCSx>
10. Medical Physics (AAPM), doi: 10.1118/1.3427411, <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.3427411>
11. *Physics in Medicine & Biology*, 57(21), 7161, doi: 10.1088/0031-9155/57/21/7161, <https://iopscience.iop.org/article/10.1088/0031-155/57/21/7161/meta>
12. *Physics in Medicine & Biology*, 51(24), doi: 10.1088/0031-9155/51/24/009, <https://iopscience.iop.org/article/10.1088/0031-9155/51/24/009/pdf>

Super-Eddington Accretion onto Neutron Stars, Ultra-compact Stars, and Black Holes: Radiative General Relativistic Magnetohydrodynamics Simulations

Fatemeh Kayanikhoo^{1,2*}, Maciej Wielgus³, Debora Lančová^{1,2}, Leszek Zdonik²,
Martin Urbanec¹, Włodzimierz Kluźniak²

¹ Research Centre for Computational Physics and Data Processing, Institute of Physics,
Silesian University in Opava, Czech Republic

² Nicolaus Copernicus Astronomical Center of the Polish Academy of Sciences, Warsaw, Poland
³ IAA-CSIC, Granada, Spain

* fatemeh.kayanikhoo@physics.slu.cz

Keywords: Ultraluminous X-ray sources, compact stars, super-Eddington accretion, GRRMHD simulations, black holes radiation beaming, HPC

1. Introduction

Ultraluminous X-ray sources (ULXs) are among the most luminous X-ray extra-galactic objects, with the luminosities beyond 10^{39} erg/s. Since 1980s, the first discovery of ULXs, many models are proposed to explain this critical luminosity. The pulsations detected in several ULXs initially by [1], point to neutron star or ultracompact star accretors. Explaining their extreme luminosities requires super-Eddington accretion with strongly beamed radiation [2].

2. Description of the problem

We investigate how the radius of a compact star affects super-Eddington accretion and observable ULX properties. Using 2.5D axisymmetric GRRMHD simulations, we model accretion onto compact stars with varying radii and surface magnetic field strengths, and onto black holes of the same mass for comparison. The strong magnetization near the stellar surface and the coupling between radiation and magnetic field require dedicated numerical treatment, making this problem tractable only on HPC systems.

3. Related work

The theoretical KLK model [3] proposed beamed emission of super-Eddington accreting neutron stars as the ULXs. In the previous GRRMHD simulations of accreting neutron stars we showed moderately magnetized neutron stars with magnetic field in order of 10 Giga Gauss are ULXs [4].

4. Solution of the problem

We conduct numerical simulations using GRRMHD code KORAL [6,7]. The neutron star accretion was implemented in the code by [5]. The code uses M1 radiation closure and force-free electrodynamics in the highly magnetized stellar magnetosphere. The code is parallelised using MPI [8]. In this work, we investigate the impact of radius of the accretor on the accretion geometry and luminosity of the source. We performed 14 compact star simulations with various radius and surface magnetic field strength and 4 black hole simulations with different spin.

6. Conclusions

We find that the compactness of the star is a key factor determining both the luminosity and beaming of the emitted radiation. More compact objects have smaller inflow rate (left panel of Fig. 1), consequently more outflow rate and stronger beaming. Accretion onto a smaller star (more compact) produces higher apparent luminosities due to stronger beaming. Accreting compact stars are more luminous than black holes of the same mass, supporting the ULX neutron star scenario (right panel of Fig. 1).

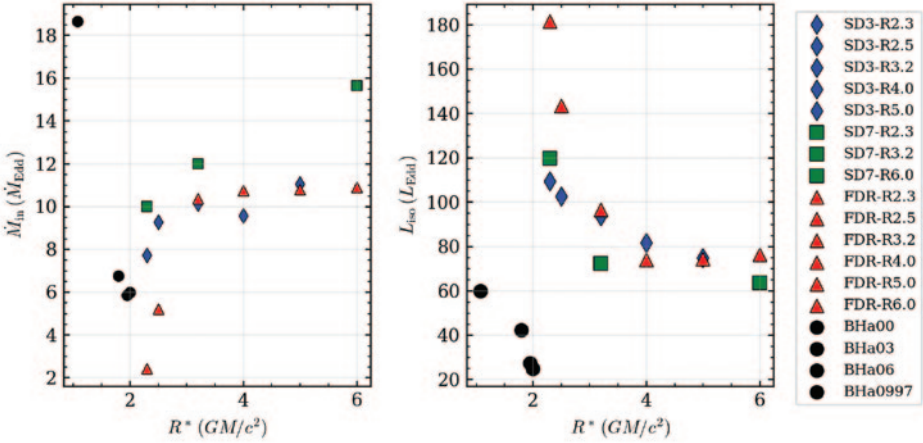


Fig. 1. The parameters inflow \dot{M}_{in} the apparent luminosity L_{iso} for 18 simulated models. The notation SD3–R# refers to simulations performed with a fixed surface magnetic dipole strength of 3×10^{10} G, where R# denotes the neutron star radius. Similarly, SD7–R# denotes simulations with a magnetic dipole strength of 7×10^{10} G. The notation FDR–R# refers to models with a fixed magnetic field strength at a specific radius. Black hole (BH) models are denoted with their spin a00, a03, a06 and a0997.

Acknowledgements. Simulations were run on the PLGrid HPC infrastructure, grant PLG/2023/016648 and PLG/2024/017525, and the research supported by CZ.10.03.01/00/23_042/0000390 (FK) and GAČR grant 25-16928O (DL).

References

1. Bachetti et al. 2014, Nature, 514, 202.
2. King et al. 2001, ApJ, 552, L109.
3. King et al. 2017, MNRAS 468, L59.
4. Kayanikhoo et al. 2025 ApJ 982 95.
5. Abarca et al. 2021, ApJ 917, L31.
6. Sądowski et al. 2013, MNRAS 429, 3533.
7. Sądowski et al. 2014, MNRAS 439, 503.
8. MPI Forum 2021, MPI Standard v4.0.

Anthropogenic Hazards Research with EPOS TCS AH EPISODES Platform

Joanna Kocot¹, Tomasz Balawajder¹, Maciej Leśniak¹, Mieszko Makuch¹, Hubert Siejkowski¹,
Jakub Stoł¹, Bartłomiej Wenda¹, Stanisław Lasocki², Anna Leśnodorska²,
Agnieszka Mtupa-Ndiaye², Beata Orlecka-Sikora²

¹ Academic Computer Centre Cyfronet AGH, ul. Nawojki 11, 30-950 Kraków, Poland

² Institute of Geophysics, Polish Academy of Sciences, ul. Księcia Janusza 64, 01-452 Warszawa, Poland

{joanna.kocot, tomasz.balawajder, maciej.lesniak, mieszko.makuch,
hubert.siejkowski, jakub.sto, bartlomiej.wenda}@cyfronet.pl,
{lasocki, alesnodorska, amtupa, orlecka}@igf.edu.pl

Keywords: EPISODES Platform, anthropogenic hazards, induced seismicity, EPOS, high-performance computing, virtual laboratory, geo-resource exploitation

1. Introduction

The EPISODES Platform [1] is a virtual laboratory for research on anthropogenic seismicity and related geohazards, integrating unique multidimensional datasets (“episodes”), analysis applications, and computational workspaces. The EPISODES Platform together with international data nodes constitutes the Thematic Core Service Anthropogenic Hazards (TCS AH) Research Infrastructure within the EPOS (European Plate Observing System) infrastructure (<https://www.epos-eu.org/>).

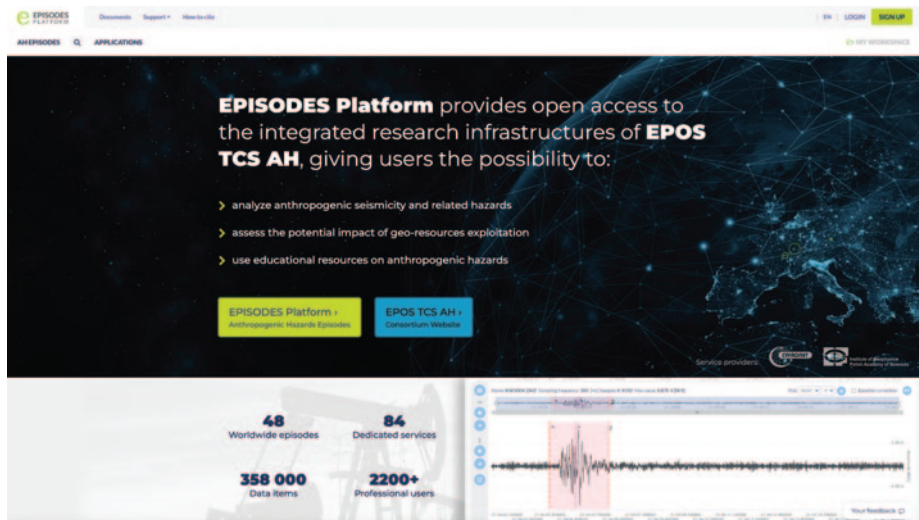


Fig. 1. EPISODES Platform main page.

2. Description of the problem

Human activities such as, e.g., mining, reservoir impoundment, geothermal energy extraction or hydrocarbon production can trigger rock deformation and induced seismicity. These events pose significant hazards to populations, infrastructure, and ecosystems, making

research into their mechanisms critically important for effective risk mitigation. However, this requires gathering extensive data related to the production and seismic process, along with adequate analytical and computational tools.

3. Related work

Existing seismological platforms, such as Orfeus (<https://www.orfeus-eu.org/>) provide waveform archives and standalone tools, but they rarely couple curated induced-seismicity datasets. There are no platforms coupling seismic data with production data, as the latter are often subject to many access restrictions.

4. Solution of the problem

EPISODES Platform – the central service of EPOS TCS Anthropogenic Hazards – federates 48 curated induced-seismicity episodes, combining seismic, industrial, and auxiliary geodata across diverse geo-resource exploitation contexts. Users design problem-oriented workflows constructed from both built-in applications (84 services currently offered) and their own codes for tasks such as AI-aided event detection, localization, spectral analysis and probabilistic hazard assessment. Computationally demanding workflows are executed on high-performance computing resources, ensuring scalable processing. EPISODES Platform has been successfully used in several European and Polish projects (e.g. DT-Geo, Geo-INQUIRE, EPOS PL, EPOS PL+) and supports community building through training and workshops.

5. Conclusions

EPISODES Platform demonstrates how a domain-specific, HPC-enabled virtual laboratory can transform induced-seismicity research into reproducible, shareable workflows, accelerating methodological innovation and supporting evidence-based risk mitigation.

Acknowledgements. This work was supported by project co-financed by the Minister of Science Republic of Poland under contract no. 2024/WK/05. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018430.

References

1. Orlecka-Sikora, B., Lasocki, S., Kocot, J. et al. (2020) An open data infrastructure for the study of anthropogenic hazards linked to georesource exploitation., *Sci Data* 7, 89, (2020), doi: 10.1038/s41597-020-0429-3.

Molecular Dynamics Study of Tetraglyme Solutions of Two Lithium Salts with Isomeric Anions: LiTFSI and LiFPFSI

Andrzej Eilmes¹, Piotr Kubisiak¹, Chiara Nicotri^{1,2}

¹ Jagiellonian University, Faculty of Chemistry, Gronostajowa 2, 30-387 Kraków, Poland

² Department of Applied Science and Technology, Politecnico di Torino, 10129 Torino, Italy

{eilmes, kubisiak}@chemia.uj.edu.pl

chiara.nicotri@student.uj.edu.pl

Keywords: Li-ion electrolytes, molecular dynamics, ion transport, conductivity

1. Introduction

Although Li-ion batteries achieved a great commercial success since they were introduced to the market in the 1990s of the 20th century, a large effort is being invested in the development of new, more effective, safe and environment-friendly devices. One of key components of a metal-ion battery is the ion-conducting electrolyte - typically a Li salt solution in a molecular liquid. The optimized electrolyte contributes to the overall performance of the battery.

2. Description of the problem

In the search for better Li-conducting electrolytes several promising salts with weakly coordinating anions are experimentally investigated, e.g. lithium bis(trifluoromethanesulfonyl) imide (LiTFSI). Recently, some asymmetric perfluorinated sulfonimide anions were studied, including the TFSI isomer, (fluorosulfonyl)(pentafluoroethanesulfonyl)imide (FPFSI) [1].




Fig. 1. Structure of TFSI and FPFSI anions.

3. Related work

The TFSI anion, more commonly used in Li-conducting electrolytes has been thoroughly theoretically investigated. On the other hand, not much computational modeling was done on the FPFSI isomer. Density functional (DFT) calculations were reported for several perfluorinated anions and their complexes with Li⁺ using an implicit solvent model [2]. In another work, DFT screening was performed for a set of salts in vacuum [3]. Recently, we studied the cation-anion interactions and vibrational spectra in LiTFSI or LiFPFSI solutions in tetraglyme using quantum chemical calculations and *ab initio* molecular dynamics simulations (AIMD) [4]. Nevertheless, the attainable length of AIMD simulations is not sufficient for obtaining diffusion coefficients, conductivities and transport numbers, which are of major practical importance.

4. Solution of the problem

We performed classical molecular dynamics simulations of LiTFSI and LiFPFSI solutions in tetraglyme using a polarizable force field developed based on quantum-chemical calculations [4]. Three salt:solvent ratios for each type of anion were simulated for 400-450 ns with a time step of 1 fs in the NVT ensemble at $T = 303$ K using Langevin dynamics. The results were averaged over 5 or 10 independent trajectories. NAMD v. 2.14 [5] package was used for MD simulations.

Based on the analysis of recorded MD trajectories, structure of electrolytes (radial distribution functions, coordination numbers, speciation of cations) and their dynamics (residence times, diffusion coefficients, conductivity, transport numbers) were determined.

5. Conclusions

For both salts, very stable $[\text{Li}(\text{tetraglyme})]^+$ solvates form in the electrolyte. For an equimolar salt:solvent composition, solutions exhibit properties of a solvate ionic liquid. In LiFPFSI electrolytes, the cation-anion interactions are slightly weaker than in LiTFSI solutions, resulting in more stable Li^+ solvates.

Although the diffusion coefficients of ions are similar for both salts, the ionic conductivities of LiFPFSI electrolytes estimated from the simulations are 40-70% larger than the conductivities of LiTFSI solutions. This enhancement originates from constructive contributions to the conductivity arising from anticorrelated motions of cations and anions, a feature characteristic for ionic liquids. However, the practical applicability of tetraglyme solutions of both salts appears to be constrained by the low Li^+ transference numbers under anion blocking conditions. Therefore, the detailed analysis of ion-ion correlations is necessary for a deeper understanding of ion transport in concentrated solutions [6].

Acknowledgements. The numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grant plgaena19. The MD simulations were running on 1-2 nodes (48-96 CPU cores) and consumed about 600k CPU hours.

References

1. P. Deng, H. Zhang, W. Feng, Z. Zhou, M. Armand, J. Nic: Lithium (Fluorosulfonyl) (pentafluoroethanesulfonyl)imide/poly(ethylene oxide) Polymer Electrolyte: Physical and Electrochemical Properties, in *Solid State Ionics* 2019, 338, 161-167.
2. H. Zhang, O. Arcelus, J. Carrasco: Role of Asymmetry in the Physicochemical and Electrochemical Behaviors of Perfluorinated Sulfonimide Anions for Lithium Batteries: A DFT Study, in *Electrochim. Acta* 2018, 280, 290-299.
3. D. A. Osborne, M. Breedon, T. R  ther, M. J. S. Spencer: Towards Higher Electrochemical Stability of Electrolytes: Lithium Salt Design Through in Silico Screening, in *J. Mater. Chem. A* 2022, 10, 13254-13265.
4. P. Kubisiak, A. Eilmes, D. Narkevi  cius, C. Nicotri: Comparative Study of Isomeric TFSI and FPF-SI Anions in Li-Ion Electrolytes Using Quantum Chemistry and Ab Initio Molecular Dynamics, in *J. Phys. Chem. B* 2025, 129, 2560-2572.
5. J. C. Phillips et al.: Scalable Molecular Dynamics on CPU and GPU Architectures with NAMD, in *J. Chem. Phys.* 2020, 153, 044130.
6. P. Kubisiak, C. Nicotri, A. Eilmes: Molecular Dynamics Study of Tetraglyme Solutions of Two Lithium Salts with Isomeric Anions: LiTFSI and LiFPFSI, in *J. Phys. Chem. B* 2026, 130, 1902-1914.

Applications of Inter-reactant Interaction Surfaces in the Study of Membrane Degradation in Alkaline Fuel Cells

Maria Rózga, Artur Michalak

Department of Theoretical Chemistry, Faculty of Chemistry,
Jagiellonian University, Gronostajowa 2, 30-387 Krakow, Poland

maria.rozga@doctoral.uj.edu.pl, michalak@chemia.uj.edu.pl

Keywords: non-covalent interactions, effective rate constant, anion-exchange membrane fuel cells, interaction surfaces

1. Introduction

Anion Exchange Membranes (AEMs) are key components in the development of efficient alkaline fuel cells (AFCs), capable of rivaling traditional proton-exchange membrane fuel cells. These membranes incorporate quaternary ammonium headgroups bound to polymer backbones, enabling alternative non-precious metal catalysts in place of rare platinum group metals. The longevity of AEMs depends heavily on the stability of cationic headgroups in alkaline conditions, where they undergo degradation through multiple mechanisms and pathways [1].

2. Description of the problem

Membranes in which isoindolinium cations with varying substituents on the aromatic ring act as cationic headgroups have been studied experimentally [2]. Our aim is to investigate the stability of novel isoindolinium cations with emphasis on two compounds with large conformational lability. This aspect introduces a great challenge as alternative reaction paths, binding sites with varied interactions, and multiple conformers have to be accounted for.

3. Related work

The Effective Rate Constant method [3] averages multiple pathways. Activation energies for some isoindolinium cations were reported in [2], however such an approach was not applied. Reactivity analysis requires first to identify binding sites. Using simple conceptual DFT methods [4] which consider only one reactant, limits comparison across varied charged and radical species.

4. Solution of the problem

The Effective Rate Constant method, which averages reactive pathways, requires activation free energies and substrate complex Gibbs energies. We built a workflow combining semiempirical MD and DFT calculations performed with Amsterdam Density Functional [5] software to evaluate up to 700 substrate-transition-state pairs. An in-house program for generating interaction surfaces marked by interaction energies, enabled an assessment and analysis of possible cation-OH binding modes.




Fig. 1. Inter-reactant interaction surfaces of isoindolinium cation derivative probed by OH⁻ anion through O atom (left) and H atom (right) with characterized and named types of interactions.

5. Conclusions

Isoindolinium cations in alkaline fuel cells are complex to model due to accuracy limits, multiple reactive species, polymer backbones, and microsolvation. Our approach provides insight into their chemistry without major simplifications, demonstrating the generality and utility of the applied methodology.

Acknowledgements. This work was funded by the National Science Centre Poland under the M-ERA.NET project “Novel asymmetric anion-exchange membranes for fuel cells”. We also thank the Research Support Module awarded by Faculty of Chemistry (UJ), and ACC Cyfronet AGH for Ares computing resources allocated on the Ares system under plgmichalak9, plgmichalak10 grants.

References

- Gjineci, N.; Aharonovich, S.; Willdorf-Cohen, S.; et al. The Reaction Mechanism Between Tetraarylammonium Salts and Hydroxide. *Eur. J. Org. Chem.* 2020, 3161–3168.
- Aggarwal, K.; Gjineci, N.; Kaushansky, A.; et al. Isoindolinium Groups as Stable Anion Conductors for Anion-Exchange Membrane Fuel Cells and Electrolyzers. *ACS Mater. Au* 2022, 2, 367–373.
- Kim, T.-J.; Kim, S.-K.; Kim, B.-J.; et al. Half-Metallocene Titanium(IV) Phenyl Phenoxide for High-Temperature Olefin Polymerization: Ortho-Substituent Effect at the Ancillary o-Phenoxy Ligand for Enhanced Catalytic Performance. *Macromolecules* 2009, 42, 6932–6943.
- Geerlings, P.; De Proft, F.; Langenaeker, W. Conceptual Density Functional Theory. *Chem. Rev.* 2003, 103, 1793–1874.
- ADF2023, SCM, Theoretical Chemistry, Vrije Universiteit, Amsterdam, The Netherlands, <http://www.scm.com>.

Theoretical Study of Alkaline Degradation Pathways of Carbazolium Cations in Asymmetric Anion-Exchange Membranes

Olga Żurowska^{1,2}, Mercedes Kukułka¹, Artur Michalak¹

¹ Jagiellonian University, Faculty of Chemistry, Kraków, Poland

² Jagiellonian University, Doctoral School of Exact and Natural Sciences, Kraków, Poland

olga.zurowska@doctoral.uj.edu.pl

Keywords: carbazolium cations, alkaline degradation, SET mechanism, asymmetric anion-exchange membranes, fuel cells

1. Introduction

The stability of ammonium salts is a key factor that determines the performance and durability of alkaline anion exchange membrane fuel cells (AEMFCs). Classical degradation routes such as Hofmann elimination and SN2 nucleophilic substitution are well described in the literature [1]. However, recent experimental studies [1,2] have shown that N,N-diaryl-carbazolium salts degrade through a different pathway that involves single electron transfer combined with radical coupling. Substituent effects, especially those introduced at the para position of the phenyl ring, were found to significantly influence the stability of these systems.

2. Description of the problem

A central challenge in understanding carbazolium-cation degradation is the lack of a unified mechanistic model describing how SET competes with SN2 nucleophilic substitution under alkaline conditions. While experimental data suggest that SET-initiated pathways may dominate in aryl-substituted systems, the relationship between redox properties, electronic structure, reaction barriers and substituent effects has not been systematically quantified.

3. Related work

Experimental studies on aryl-substituted ammonium salts have reported SET-driven degradation competing with classical SN2 pathways [1,2]. Carbazolium units are typically incorporated as structural fragments within polymer backbones rather than used as standalone cations in ionomers [3], yet understanding their intrinsic alkaline stability remains important for assessing their potential in AEMFC-related applications.

4. Solution of the problem

To investigate the competition between the degradation pathways of substituted carbazolium cations, we developed a computational workflow focused on two processes: nucleophilic SN2 attack and charge-transfer-related reactivity. All calculations were performed using AMS 2024.106, applying DFT with various exchange–correlation functionals, D3(BJ) dispersion correction, and ZORA scalar-relativistic. SN2 pathways induced by OH⁻ were determined. Transition states were optimized and verified through frequency and intrinsic reaction coordinate (IRC) analysis, yielding substituent-dependent SN2 activation barriers. Charge-transfer energy was approximated based on ETS energy decomposition analysis, for the donor-acceptor systems relevant to SET processes. Together, these results describe how substituent-dependent electronic structure influences the balance between nucleophilic and charge-transfer-driven degradation pathways.

All computations were performed on high-performance computing resources provided by the PLGrid infrastructure, using 24 CPU cores per job.

5. Conclusions

The combined analysis of SN2 reaction barriers and ETS derived charge transfer energies provides a coherent mechanistic picture of carbazolium cation degradation under alkaline conditions. Substituent effects strongly modulate both nucleophilic reactivity and charge transfer propensity, revealing clear structure–reactivity trends. The results identify electronic features that either suppress or promote each pathway, enabling rational selection of substituents for more stable cation in ionomers for AEMFCs applications.

Acknowledgements. The research was funded by the National Science Centre, Poland, Grant UMO-2022/04/Y/ST4/00154 (M-ERA.NET 3 Call 2022). The numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grants no. PLG/2024/017261, PLG/2025/018396.

References

1. Nansi Gjineci, Sinai Aharonovich, Dario R. Dekel, and Charles E. Diesendruck *Eur. J. Org. Chem.* 2020, 3161–3168.
2. Nansi Gjineci, Sinai Aharonovich, Sapir Willdorf-Cohen, Dario R. Dekel, and Charles E. Diesendruck, *ACS Appl. Mater. Interfaces* 2020, 12, 49617–49625.
3. Jiayao Yang, Jialin Zhao, Na Li, Yijia Lei, Jingyi Wu, Jian Gao, Shiyao Sun, Kuirong Feng, Yan Wang, and Zhe Wang, *Chem. Eng. J.* 2024, 489, 151446.

Computational Investigation of Temperature - Dependent Luminescence in Transition Metal Complexes

Paweł J. Bonarek^{1,2}, Robert Jankowski¹, Junhao Wang³, Szymon Chorąży¹

¹ Faculty of Chemistry, Jagiellonian University, Gronostajowa 2, Kraków, Poland

² Doctoral School of Exact and Natural Sciences, Jagiellonian University,
Prof. St. Łojasiewicza 11, Kraków, Poland

³ Department of Material Science, Faculty of Pure and Applied Science, University of Tsukuba 1-1-1
Tennodai, Tsukuba, Ibaraki 305-8573, Japan

pawel.bonarek@doctoral.uj.edu.pl, wang.junhao.gb@u.tsukuba.ac.jp,
{robert14.jankowski, simon.chorazy}@uj.edu.pl

Keywords: luminescence, transition metal complexes, vibronic progression, TD-DFT

1. Introduction

Among luminescent materials, those built of transition metal complexes bearing aromatic organic ligands reveal highly efficient charge-transfer-based luminescence (phosphorescence), which makes them applicable in, e.g., OLEDs [1]. Furthermore, as the light emission of such complexes is often highly sensitive to external stimuli (temperature (T), pressure, solvent vapors, gases), they gain considerable scientific interest as promising molecular building blocks for the solid luminophores applicable in sensing and data storage, as well as advanced multifunctional systems [2,3]. Therefore, the development of theoretical investigation methods to rationalize experimental observations related to the optical responsiveness of transition metal complexes to chemical and physical stimuli is of high importance.

2. Description of the problem

A simulation of the vibrationally-resolved and T -dependent charge-transfer emission spectra of transition metal complexes remains challenging, as it must consist of both the excited electronic states and the molecular vibrations analyses, while also considering the charge separation and the relativistic effects in the heavy atoms. Nevertheless, such a holistic approach is crucial to meaningfully determine the emissive properties *in silico*.

3. Related work

Previous works on the mentioned problem established some guidelines on how to approach the simulation of the emission spectrum based on density functional theory (DFT) [4,5]. However, the T -dependence of the vibronic progression remains mostly unexplored, even though some computational modules allow for the exact accounting of the temperature [6].

4. Solution of the problem

The calculations were performed for two selected metal complexes, i.e., $[\text{Pt}^{\text{II}}(\text{CN})_2(\text{ppy})]^-$ and $[\text{Ru}^{\text{II}}(\text{CN})_4(\text{bpy})]^{2-}$ (ppy = carbanion of 2-phenylpyridine, bpy = 2,2'-bipyridine), and consisted of unrestricted DFT geometry optimization of ground (S_0) and excited (T_1) states, a time-dependent (TD-DFT) analysis of the excited states with spin-orbit coupling, and the phosphorescence spectra simulation within an adiabatic hessian approach (Fig. 1). The relativistic effects were evaluated with 0th-order regular approximation (ZORA) using dedicated basis sets. All calculations were done with ORCA 6.0.1 quantum-chemistry software package [7], on the ARES system, usually using 1 node (48 cores) and 450 k computing hours in total.

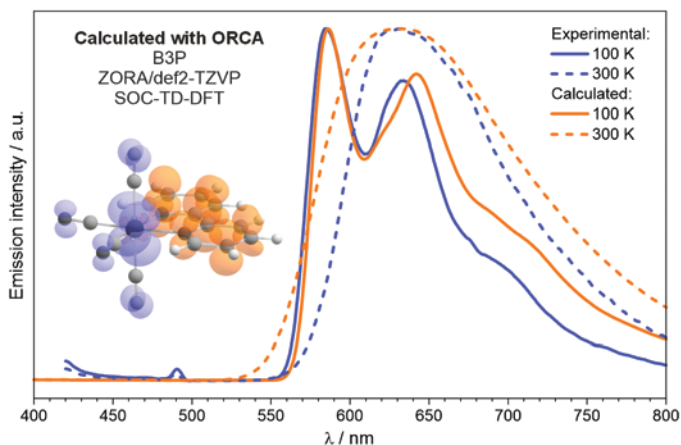


Fig. 1. A comparison between the experimental and calculated phosphorescence spectra of $[\text{Ru}^{\text{II}}(\text{CN})_4(\text{bpy})]^{2-}$ anion at low and room temperatures, together with the electron-density difference plot for $S_0 \leftarrow T_1$ transition (drawn at 0.05 isosurfaces level).

5. Conclusions

The comparison between the computed and experimental results has proven that the chosen TD-DFT approach is suitable for simulating T -dependent emission spectra of transition metal complexes, at least for the studied heteroleptic complexes with cyanido and organic ligands, and allows for rationalization of the results for the better understanding of the optical phenomena.

Acknowledgements. This work has been supported by the National Science Centre of Poland within the SONATA BIS-9 grant (2019/34/E/ST5/00148) and the MINIATURA-8 grant (2024/08/X/ST5/01289). The numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grant no. PLG/2024/017429.

References

1. H. Yersin, A. F. Rausch, R. Czerwieniec, T. Hofbeck, T. Fisher: "The triplet state of organo-transition metal compounds. Triplet harvesting and singlet harvesting for efficient OLEDs", in *Coord. Chem. Rev.*, 2011, 225, 2622–2652.
2. J. Rzepiela, M. Liberka, M. Zychowicz, J. Wang, H. Tokoro, K. Piotrowska, S. Baś, S. Ohkoshi, S. Chorazy: "SHG-active luminescent thermometers based on chiral cyclometalated dicyanidoiridate(III) complexes", in *Inorg. Chem. Front.*, 2024, 11, 1366–1380.
3. P. Bonarek, M. Zychowicz, J. Rzepiela, M. Liberka, S. Baś, J. J. Zakrzewski, S. Chorazy: "Design of Dy^{III} single-molecule magnets with molecularly installed luminescent thermometers based on bridging $[\text{Pt}^{\text{II}}(\text{CN})_2(\text{C}'\text{N})]^-$ complexes", in *Inorg. Chem. Front.*, 2024, 11, 7966–7978.
4. R. Grotjahn and M. Kaupp: "Reliable TDDFT Protocol Based on a Local Hybrid Functional for the Prediction of Vibronic Phosphorescence Spectra Applied to Tris(2,2'-bipyridine)-Metal Complexes", in *J. Phys. Chem. A*, 2021, 125, 7099–7110.
5. R. Schira and C. Latouche: "DFT vs. TDDFT vs. TDA to simulate phosphorescence spectra of Pt- and Ir-based complexes", in *Dalton Trans.*, 2021, 50, 746–753.
6. B. de Souza, F. Giliandro, F. Neese, R. Izsák: "Predicting Phosphorescence Rates of Light Organic Molecules Using Time-Dependent Density Functional Theory and the Path Integral Approach to Dynamics", in *J. Chem. Theory Comput.*, 2019, 15, 1896–1904.
7. F. Neese: „Software Update: The ORCA Program System–Version 6.0”, in *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 2025, 15, e70019.

Palmitoylation as a Structural Switch for the CFTR Ion Channel

Jakub Mróz^{1,2}, Wojciech Kopec^{3,4}, Mateusz Sikora^{1,2}

¹ Małopolska Centre of Biotechnology, Jagiellonian University, Kraków, Poland

² Max Planck Institute of Biophysics, Frankfurt, Germany

³ Department of Chemistry, Queen Mary University of London, London, UK

⁴ Computational Biomolecular Dynamics Group, Max Planck Institute for Multidisciplinary Sciences, Göttingen, Germany

`jakub.mroz@doctoral.uj.edu.pl, mateusz.sikora@uj.edu.pl`

Keywords: CFTR, molecular dynamics, multiscale simulations, ion channels

1. Introduction

Cystic fibrosis is a severe genetic disease that affects the lungs and digestive systems and significantly shortens patients' lives. It is caused by malfunction of CFTR, a protein that regulates salt transport across cell membranes [1]. Although new drugs improve CFTR function, many aspects of how the protein switches between active and inactive states remain unclear. Understanding these mechanisms is essential for developing better therapies.

2. Description of the problem

CFTR activity depends on complex structural changes that are difficult to observe experimentally [2]. One suspected regulatory mechanism is the attachment of lipid molecules that may influence how the protein interacts with the cell membrane. It is unknown whether such modifications can stabilize specific functional states of CFTR. Large-scale molecular simulations allow to explore these effects and propose mechanistic explanations.

3. Related work

Experimental studies have shown that CFTR undergoes lipid attachment, but the functional and structural consequences of this modification remain unclear [3,4]. Similar lipid-mediated membrane anchoring has been observed in other proteins [5], yet its role in CFTR has not been investigated computationally.

4. Solution of the problem

We performed multiscale molecular dynamics simulations of full-length CFTR embedded in lipid membranes using all-atom and coarse-grained force fields (CHARMM36m and Martini 3) in GROMACS 2025.1. Simulations were run on the Helios and in-house HPC clusters. Our results suggest that lipid attachment to one CFTR domain can anchor it to the membrane, stabilizing a specific structural state of the protein.




Fig. 1. CFTR protein in hypothesized, membrane-bound state.

5. Conclusions

Our results suggest that lipid attachment to one CFTR domain can anchor it to the membrane, stabilizing a specific structural state of the protein. This anchoring reduces domain mobility and may represent a previously unrecognized regulatory mechanism relevant for therapeutic strategies. This study also improves our understanding of how this medically important protein works and demonstrates how large-scale molecular simulations can help reveal mechanisms underlying human disease.

Acknowledgements. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Centers: ACC Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018090 and PLG/2026/019112.

References

1. Bell S.C., Mall M.A., Gutierrez H, et al. Cystic Fibrosis: A Review. *JAMA*. 2023.
2. El Hiani, Y., & Linsdell, P. Conformational changes opening and closing the CFTR Chloride Channel: Insights from cysteine scanning mutagenesis. *Biochemistry and Cell Biology*. 2014.
3. McClure M, McCaffrey M, Sorscher E.J., et al. S-palmitoylation regulates biogenesis of core glycosylated wild-type and F508del CFTR in a post-ER compartment. *Biochemical Journal*. 2014.
4. Abu-Arish A., Pandžić E., Luo Y., Sato Y., Turner M. J., Wiseman P. W., & Hanrahan J. W. Lipid-driven CFTR clustering is impaired in cystic fibrosis and restored by Corrector Drugs. *Journal of Cell Science*, 2022.
5. Bhattacharyya R., Barren C., & Kovacs D. M. Palmitoylation of amyloid precursor protein regulates amyloidogenic processing in lipid rafts. *Journal of Neuroscience*. 2013.

Mechanistic Plasticity of the Diels–Alder Reaction: the Role of Substituent and Solvent Effects

Agnieszka Łapczuk

Cracow University of Technology, Faculty of Chemical Engineering and Technology,
Warszawska 24, 31-155 Kraków Poland

agnieszka.lapczuk@pk.edu.pl

Keywords: Cycloaddition; Molecular Mechanism; Nitronorbornene

1. Introduction

The Diels–Alder (DA) reaction remains one of the most powerful and widely applied strategies for constructing the norbornene framework. This structural motif is a fundamental component of many biologically and technologically important molecules, including pharmaceuticals, polymer building blocks, and advanced functional materials. Moreover, its unique reactivity and rigidity make it highly attractive for the design of novel molecular architectures and emerging applications in modern organic and materials chemistry.

2. Description of the problem

It is well established that the cycloaddition mechanism of methylcyclopentadiene (CpMe) may proceed via multiple pathways. The present study aims to determine whether one of the feasible routes involves the formation of a zwitterionic intermediate and to elucidate how solvent variation and the presence of substituents in the alkene influence this mechanism.




Fig. 1. Possible reaction pathways of a cyclopentadiene analog with a conjugated nitroalkene leading to nitronorbornene products.

3. Related work

Recent studies have investigated the reactions of CpMe with conjugated nitroalkenes at the B3LYP/6-311G(d) level of theory [1]. In addition, a bicyclic 1,2-oxazine 2-oxide intermediate has been characterized, which is formed in the reaction of CpSiMe₃ with nitroalkenes [2].

4. Solution of the problem

All calculations were performed using the GAUSSIAN 16 software package. Molecular geometries were visualized using GaussView 6.0. Thermodynamic and kinetic parameters were calculated at the ω B97X-D/6-311G(d) level of theory in the presence of various solvents, simulated using the Polarizable Continuum Model (PCM). All reactants, intermediates, and products were optimized using the Berny algorithm, ensuring convergence to stationary points on the potential energy surface. Transition state (TS) geometries were obtained through the QST2 procedure, followed by harmonic frequency analyses to verify their nature. Topological analyses of the Electron Localization Function (ELF) were carried out with the TopMod 09 program, based on monodeterminantal wavefunctions computed on a grid with a spatial resolution of 0.1 atomic units.

5. Conclusions

We investigated the cycloaddition of CpMe with conjugated nitroalkenes. Pathways A and C lead to nitronorborene via a bicyclic 1,2-oxazine 2-oxide intermediate, whereas pathways B and D in nonpolar media afford the product directly. Increasing solvent polarity does not alter the mechanism until highly polar conditions are reached, where pathway B becomes stepwise and involves a zwitterionic intermediate. NPA, MEP, NCI, and BET analyses confirmed its polar, highly asynchronous character. Electron-withdrawing substituents reduce activation barriers, supporting a general electronically controlled mechanism.




Fig. 2. Simplified representation of the molecular mechanism of the reaction between 1 and 2c (path B) using ELF-based Lewis structures.

Acknowledgements. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018731.

References

1. Karaś, A. Łapczuk: “Computational model of the formation of novel nitronorborene analogs via Diels–Alder process”, *Reaction Kinetics Mechanisms and Catalysis*, 2025, 138(4), pp. 2671–2689.
2. D. Kapuściński, A. Łapczuk: “Bicyclic 1,2-oxazine 2-oxides as the discrete intermediates in the reactions between regioisomeric trimethylsilylcyclopentadienes and (2E)-2-nitro-3-phenylprop-2-enitrile”, *Chemistry of Heterocyclic Compounds*, 2025, 61(7-8), pp. 372–380.

Abduction-Guided Experimental Design for Biomimetic CaP Synthesis via MD Screening of Ca–P Ionic Interactions

Krzysztof Stafin^{1,2}, Paweł Śliwa¹, Marek Piątkowski²

¹ Cracow University of Technology, Faculty of Chemical Engineering and Technology, Department of Organic Chemistry and Technology, ul. Warszawska 24, PL 31-155, Kraków, Poland

² Cracow University of Technology, Faculty of Chemical Engineering and Technology, Department of Biotechnology and Physical Chemistry, ul. Warszawska 24, PL 31-155, Kraków, Poland

krzysztof.stafin@jezuici.pl

Keywords: MD simulation-guided design, abductive reasoning, biomimetic calcium phosphate, chitosan-based matrix

1. Introduction

In bone tissue engineering, a primary objective is to synthesize biomimetic calcium phosphate (CaP) with crystallinity and mineral maturity that closely resemble those of biological apatite [1]. Here, we promote CaP mineralization on a polyelectrolyte matrix composed of polycationic chitosan (CS) and polyanionic carboxymethyl chitosan (CCS), emulating the bone organic template.

2. Description of the problem

CaP synthesis parameters are often selected through trial and error, while molecular dynamics (MD) simulations are typically applied *post hoc* to interpret outcomes. The methodological challenge is whether MD-based screening of Ca–P ionic interactions can prospectively guide parameter selection for CaP synthesis. Addressing this requires identifying simulation-derived descriptors, establishing validation criteria, and defining a reasoning framework that links simulated molecular signals to macroscopic, experimentally measurable outputs.

3. Related work

Previous studies have used MD simulations to investigate Ca–P ion association [2], while experimental studies have shown that polysaccharide-based matrices can support biomimetic CaP mineralization [3]. However, these two approaches have usually remained separate. Studies by our group [4,5] began to bridge this gap. The present study uses MD descriptors to guide the pre-experimental selection of synthesis parameters according to rational criteria within an abductive framework.

4. Solution of the problem

MD results were used as heuristic guidance for protocol design. Simulations of the prenucleation stage of CaP synthesis quantified kinetic descriptors (MSD, D^{xyz}) for Ca^{2+} , HPO_4^{2-} , and H_2PO_4^- in the presence of CS/CCS, as well as the thermodynamic descriptor, the Ca^{2+} – HPO_4^{2-} interaction energy (IE). These metrics were evaluated at pH 7.4, 8.0, and 9.0, while temperature, ion concentrations, and CS–CCS protonation were adjusted according to experimental requirements and the literature. At pH 8.0, stable Ca–P clusters with Ca/P ratios of 1.67 and 1.50 emerged, consistent with apatite-related precursor populations.

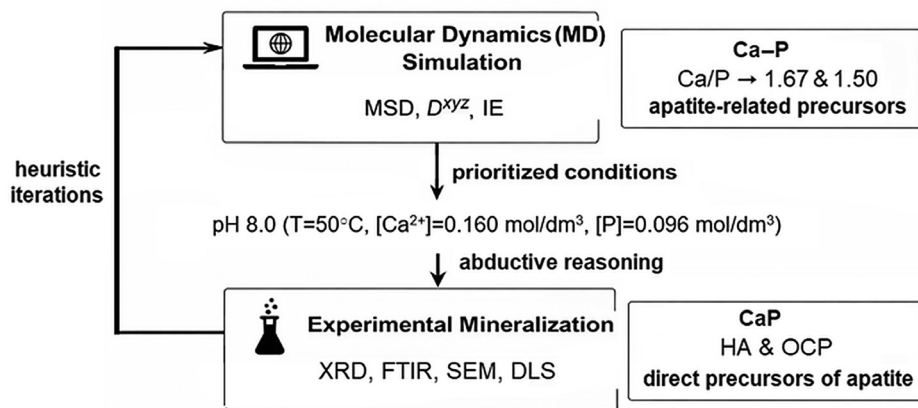


Fig. 1. Abduction-guided experimental design toward biomimetic CaP mineralization.

Using abductive reasoning, we prioritized pH 8.0 (Fig. 1) and hypothesized that it would promote biomimetic CaP formation. XRD and FTIR confirmed the presence of hydroxyapatite (HA) and octacalcium phosphate (OCP), phases commonly associated with biological apatite formation, thereby supporting the MD-guided protocol.

5. Conclusions

MD-based screening of Ca–P ionic interactions can provide decision-relevant guidance for selecting CaP synthesis conditions, reducing reliance on trial-and-error parameter tuning and supporting a more prospective design of biomimetic mineralization experiments.

Acknowledgements. This work was supported by The National Centre for Science and Development, Poland, under grant number LIDER/42/0149/L-9/17/NCBR/2018. Computational resources were provided by ACC Cyfronet AGH (Ares) under allocations **PLG/2023/016772** and **PLG/2024/017832**.

References

1. K. Stafin, P. Śliwa, and M. Piątkowski: Towards polycaprolactone-based scaffolds for alveolar bone tissue engineering: A biomimetic approach in a 3D printing technique. *Int. J. Mol. Sci.* 24, 2023, 16180.
2. G. Mancardi, C.E. Hernandez Tamargo, D. Di Tommaso, and N.H. De Leeuw: Detection of Posner's clusters during calcium phosphate nucleation: A molecular dynamics study. *J. Mater. Chem. B* 5, 2017, pp.7274–7284.
3. V.C. Dumont, A.A.P. Mansur, S.M. Carvalho, F.G.L.M. Borsagli, M.M. Pereira, and H.S. Mansur: Chitosan and carboxymethyl-chitosan capping ligands: Effects on the nucleation and growth of hydroxyapatite nanoparticles for producing biocomposite membranes. *Mater. Sci. Eng. C* 59, 2016, pp.265–277.
4. K. Stafin, P. Śliwa, M. Piątkowski, and D. Matysek: Chitosan as a templating agent of calcium phosphate crystalline phases in biomimetic mineralization: Theoretical and experimental studies. *ACS Appl. Mater. Interfaces* 16, 2024, pp.63155–63169.
5. K. Stafin, P. Śliwa, M. Piątkowski, and D. Matysek: Synergistic role of chitosan and carboxymethyl chitosan in regulating phase maturation and crystallinity of biomimetic calcium phosphate: From precursors to apatite. *Carbohydr. Polym.* 375, 2026, 124763.

Theoretical Modeling of Two D-sp³-A Type Donor-Acceptor Systems with Extreme Solvatofluorochromism

Wojciech Radzik^{1,2}, Marcin Andrzejak¹

¹ Doctoral School of Exact and Natural Sciences, Jagiellonian University, Prof. St. Łojasiewicza 11, Kraków, Poland

² Faculty of Chemistry, Jagiellonian University, Gronostajowa 2, Kraków, Poland

wojciech.radzik@doctoral.uj.edu.pl, m.andrzejak@uj.edu.pl

Keywords: fluorescence, dipole moment, DFT, CC2, ADC(2), PCM, solvatofluorochromism

1. Introduction

Donor-acceptor systems (D/A) are widely used in optoelectronics. Applicability of D/A system is determined by the nature of the low-energy electronic states. Among these, relatively less studied are systems, in which two fragments are linked by a single sp³ carbon atom, like the malachite green lactone analogues: MGLA and MGLA-C5. In both lactones DMA moiety is donor and PD moiety is acceptor. Based on the experimental evidence, Karpiuk et al. postulated that the MGLA and MGLAC5 in their lowest singlet excited state show near complete intramolecular electron transfer [1], [2], which was to explain the extremely large dipole moment in the S₁ state, exceeding 24 D. This value was estimated from the very strong dependence of the fluorescence energy on solvent polarity (solvatofluorochromism).




Fig. 1. Chemical structures of studied compounds.

2. Description of the problem

Quantum chemistry treatment of solvation effects offers limited credibility when the solute is in an excited state. The solvation energies are strongly dependent on the quantum chemistry method and on the solvation model, especially when the wavefunction of the solute significantly changes with the electronic excitation. On this view, properties of MGLA and MGLA-C5 make them an excellent playground for testing the performance of the available theoretical approaches.

3. Related work

Guido et al. [3] showed that credibility of methods based on time dependent density functional theory (TDDFT) strongly depends on type of described transition (local or charge transfer). Promising alternatives are restricted open-shell Kohn-Sham (ROKS) [4] or *ab initio* methods [5].

4. Solution of the problem

Different methods describing solvent and solute in S_1 were used to estimate fluorescence energy and excited state dipole moment. These calculations were performed with time dependent functional theory (TDDFT) with linear response (LR) and state specific (SS) polarizable continuum model (PCM), restricted open-shell Kohn-Sham (ROKS) method with SS-PCM and 2 *ab initio* methods: CC2 and ADC(2) (both with SS-PCM). For acetonitrile fluorescence energies were also estimated with explicit solvent using ROKS. Different approaches were used combining explicit solvent with PCM or classically approximated bulk solvent molecules. Calculations were performed with: TURBOMOLE 7.8.1 and 7.9; Q-Chem 6.2, 6.3 and 6.4; Gaussian 16 and NAMD 2.14.

5. Conclusions

Theoretical study confirms experimental conclusions that S_1 states of MGLA and MGLAC-5 exhibit the near complete intramolecular electron transfer and an extremely high dipole moment exceeding 24 D. For TDDFT the routinely used LR-PCM model of solvation completely fails in recovering the fluorescence energy dependence on the solvent polarity. The strong solvato-fluorochromism is partly restored by using the 1-st order perturbative correction to the S_1 state energy, while good agreement with experiment is achieved only by using the iterative, self-consistent procedure of Improta et al. [6]. An interesting alternative to TDDFT is the ROKS model, which offers not only the excellent reproduction (comparable with the correlated *ab initio* methods) of the experimental solvent-induced shifts of the fluorescence band, but also (in stark contrast with TDDFT) is only weakly dependent on the choice of functional. Preliminary results based of our QM/MM studies of using the explicit solvent show physically sound (and statistically meaningful) trends in the solvation energy accounted for by using increasing amount of the solvent molecules.

Acknowledgements. The numerical experiment was possible through computing allocation on the Ares and Helios systems at ACC Cyfronet AGH under the grants PLG/2024/017931, PLG/2024/017822, PLG/2025/019019 and PLG/2026/019099.

References

1. J. Karpiuk, P. Gawryś, E. Karpiuk and K. Suwińska. "Electron transfer across a spiro link: extreme solvatofluorochromism of a compact spiro-bridged N,N-dimethylaniline-phthalide dyad", in Chem. Commun. 55 (2019), pp. 8414–8417.
2. J. Karpiuk, E. Karolak, and J. Nowacki. "Photoinduced Intramolecular Charge Separation and Recombination in a Donor-Acceptor Dyad Linked via Tetrahedral Carbon Atom. Photophysics of a Malachite Green Lactone Analogue", in Polish J. Chem. 82 (2008), pp. 865–882.
3. C. A. Guido, D. Jacquemin, C. Adamo, and B. Mennucci. "Electronic Excitations in Solution: The Interplay between State Specific Approaches and a Time-Dependent Density Functional Theory Description" in J. Chem. Theory Comput. 11 (2015), pp. 5782–5790.
4. T. Froitzheim, L. Kunze, S. Grimme, J. M. Herbert, and J.-M. Mewes "Benchmarking Charge-Transfer Excited States in TADF Emitters: Δ DFT Outperforms TD-DFT for Emission Energies" in J. Phys. Chem. A 128 (2024), pp. 6324-6335.
5. T. Froitzheim, C. Hättig, and J.-M. Mewes "Mind the gaps: what the STGABS27 set can teach about second-order excited state methods, solvent models, and charge transfer" in Phys. Chem. Chem. Phys. 27 (2025), pp. 18870-18886.
6. R. Improta, V. Barone, G. Scalmani, and M. J. Frisch. "A state-specific polarizable continuum model time dependent density functional theory method for excited state calculations in solution". In: The Journal of Chemical Physics 125.5 (2006), s. 054103.

Development of AMBER Force Field Parameters for W-cofactors and Application to Study W-enzymes' Catalytic Properties

Maciej Szalaniec¹, Victor Baerle¹, Tommaso Attucci², Claudia Andreini², Agata Raczyńska^{3,4}

¹ Jerzy Haber Institute of Catalysis and Surface Chemistry, Kraków, Poland

² Department of Chemistry, University of Florence, Italy

³ Tunneling Group, Biotechnology Centre, Silesian University of Technology, Gliwice, Poland

⁴ Department of Bioinformatics and Artificial Intelligence Applications, Polish-Japanese Academy of Information Technology, Warsaw, Poland

maciej.szalaniec@ikifp.edu.pl

Keywords: tungsten enzymes, AMBER, MD simulations, aldehyde oxidoreductase, W-formate dehydrogenase

1. Introduction

The incorporation of tungsten pterin complexes (W-co, Fig. 1) into the active site of enzymes enormously expands the available repertoire of possible chemical transformations. The metal facilitates two-electron redox processes due to stable IV and VI oxidation states and acts as a Lewis acid. Meanwhile, the other ligands participate in the reactions by (i) providing additional means to activate recalcitrant bonds, (ii) enabling hydroxylations without molecular oxygen (by O transfer to activated intermediates), (iii) introducing strong acid/base catalysis (accepting protons or enabling proton transfer electron transfer processes), (iv) tuning the W redox potentials by varying their positions, or (v) even participate directly in the redox processes by providing a reservoir for additional electrons.




Fig. 1. The example of W-cofactor from aldehyde oxidoreductase (AOR) [1].

2. Description of the problem

Although the capabilities granted by tungsten (or molybdenum) cofactors to enzyme catalysis are of high importance, their complicated and non-standard structure make them difficult to describe by molecular force fields such as AMBER [2]. This makes MD simulations and QM:MM modeling for W-enzymes difficult. To address this problem, we have developed parameters describing W-co at VI and IV oxidation states and with different ligands and tested them with MD and QM:MM.

3. Related work

In recent years AMBER parameters were developed for molybdenum enzymes, mainly by our group (e.g. [3]) while for W-enzymes QM:MM multiscale modeling was used where QM calculations provided description of W-co [4].

4. Solution of the problem

The following protocol was used to develop and test the parameters:

- DFT geometry minimization of the cofactors models at different oxidations states (IV-VI), with different ligands coordinating metal and different oxidation states of pterins (tetrahydropterin and dihydropterin)
- Vibrational analysis of the obtained geometries and retrieval of the force constants
- Calculation of electron density followed by RESP fitting of point charges to reproduce electrostatic potential
- Testing isolated cofactors in water box (if structure not good addition of new parameters) in short MD simulation
- Testing parameters for the whole protein models in water box (150-200 ns simulations)
- Clustering of the trajectory and minimization of representative geometries:
 - Using AMBER molecular mechanics
 - Using QM:MM (DFT:AMBER)
- Application of the tested parameters to study structure and catalytic dynamic behavior of enzymes.

We have used the obtained parameters to study W-formate dehydrogenase (Fdh), the enzyme that catalyzes reversible formate-CO₂ interconversion as well as AOR, that catalyzes reversible interconversion of aldehydes and carboxylic acids. The results of simulations of Fdh were analyzed with AQUA-DUCT^[5] providing insight into substrate and product channels while results for AOR were compared with crystallographic and spectroscopic data.

5. Conclusions

The parametrization of W-co enables MD simulation and multiscale modeling of biologically and industrially important processes. The protocol developed by us can be also adapted for other Mo/W-enzymes expanding the molecular modeling toolbox.

Acknowledgements. This research has been supported by HORIZON-EIC-2023- PATHFINDE-ROPER-01 project W-BioCat. The numerical experiment was possible through computing allocation on the Ares and Helios system at ACC Cyfronet AGH under the grants PLG/2025/018297 and PLG/2025/018449.

References

1. M. Szaleniec, J. Heider, *Biochemistry* 2025, 64, pp. 2154–2172.
2. R. Salomon-Ferrer, D. A. Case, R. C. Walker, *Wiley Interdisciplinary Reviews: Computational Molecular Science* 2013, 3, pp. 198-210.
3. a) A. Rugor, A. Wojcik-Augustyn, E. Niedzialkowska, S. Mordalski, J. Staron, A. Bojarski, M. Szaleniec, *J. Inorg. Biochem.* 2017, 173, pp. 28-43; b) N. C. Giri, L. Wedasingha, N. Manicke, M. Szaleniec, P. Basu, *J. Am. Chem. Soc.* 2025, 147, pp. 13243-13250.
4. M. Culka, S. G. Huwiler, M. Boll, G. M. Ullmann, *J. Am. Chem. Soc.* 2017, 139, pp. 14488-14500.
5. T. Magdziarz, K. Mitusińska, M. Bzówka, A. Raczyńska, A. Stańczak, M. Banas, W. Bagrowska, A. Góra, *Bioinformatics* 2020, 36, pp. 2599-2601.

Mathematical Model Explaining the Structural Transformation of Proteins Causing Neurodegenerative Diseases

Irena Roterman-Konieczna¹, Dawid Dułak², Katarzyna Stapor³, Leszek Konieczny¹

¹Jagiellonian University – Medical College, Kraków, Poland

²ASSA ABLOY Opening Solutions Poland S.A. Ul. Magazynowa 4, 64-100 Leszno, Poland

³Silesian University of Technology Gliwice, Poland

myroterm@cyf-kr.edu.pl, katarzyna.stapor@polsl.pl

Keywords: simulation of protein folding, misfolding, amyloids causing dementia

1. Introduction

Proteins causing dementia are examples of proteins not obeying the biological dogma treating the folding process as deterministic. The misfolded proteins causing dementia are not mutation-related [1]. The specificity of protein structures causing the neurodegenerative diseases is the formation of long fibrils of unlimited propagation. Their presence damage the cellular constructions. The base for proposed model explaining the structural transformation introduces the environment-dependent nature of folding process directing it toward the structure preferable for external conditions [2]. Thus the source of the disease-related process is generated by changes in surrounding the proteins are acting. The misfolding changes the protein structure from soluble globular form into Beta-structural forms complexed one to another causing the generation of fibrillar form (as shown in Fig. 1).

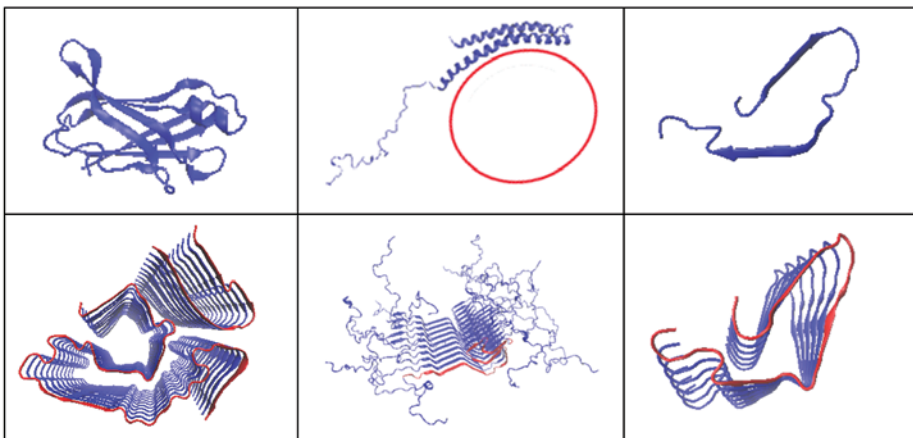


Fig.1. The native (upper row) and misfolded (lower row) structures: Left: transthyretin – disease causing, central: Alpha-Synuclein – disease causing, right: endorphin – natural, biologically active protein if form of fibril with self-control regarding length of fibril (limited to 8 chains).

2. Description of the problem

The question is focused on the conditions supporting the transformation. The main assumption regarding structural transformation in fuzzy oil drop model (FOD-M) is related to changes of environmental conditions. Model allows quantitative expression of the degree (expressed by

parameter K) of environmental changes in respect to natural conditions for biology, which is water. The problem of amyloid transformation is presented in details in open access system [3,4]. The background and application of the fuzzy oil drop model discussing the early stage step of folding, partial unfolding (important for amyloid transformation) and definition of environmental changed inducing the structural alteration is presented in [3,4]. Using the FOD-M model for analysis it is suggested, that the function expressing energetic level related to particular structure shall be treated as multiple object analysis. It is postulated, that the energy function depends on two functions expressing the internal and external force fields influencing the directional structural changes during folding process. The method called front Pareto is suggested to be used to solve the problem of protein folding. The conditions as they appear in in vitro experiments can be explained. However the specificity of that conditions are difficult to be transformed to in vivo conditions.

3. Results

The algorithm developed and implemented is available in open access system: <https://hphob.sano.science/>. Program uses the data collected by publicly available Protein Data Bank [5]. This tool allows calculation of K value for any protein present in PDB.

4. Conclusions

The disease-related proteins in form of amyloids appear to follow three scenarios of structural transformations: 1. significant influence of other than water factors directing the amyloid transformation (VL domain of IgG and transthyretin) 2. adaptation to water environment (α -synuclein) lowering the K value after loosing contact with target object (red circle on Fig.1.) stabilising the structure of high K value and 3. stable adaptation to water environment of biologically active amyloid (beta-endorphin). The example 3 represents self-control of the fibril length (maximum 8 chains in fibril) avoiding the unlimited elongation of the fibril. The examples representing each of proposed scenarios is shown on Fig.1.

Acknowledgements. The authors wish to thank to Anna Śmietańska and Zdzisław Wiśniowski for technical support. This research was supported by the Polish Minister of Science and Higher Education under the program “Support for the activity of Centers of Excellence established in Poland under Horizon 2020” (MEiN/2023/DIR/3796). The project also received funding from the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 857533). We also gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACC Cyfronet AGH) for providing computer facilities and support within computational grant „plgrantfordrippy”. Special thanks to Piotr Nowakowski and Krzysztof Gądek from the Sano Centre for Computational Medicine, Krakow, Poland.

References

1. F. Chiti, CM. Dobson. Protein Misfolding, Amyloid Formation, and Human Disease: A Summary of Progress Over the Last Decade. *Annu Rev Biochem.* 2017; 86: 27-68. doi: 10.1146/annurev-biochem-061516-045115
2. I. Roterman, I. Konieczny. Protein Is an Intelligent Micelle. *Entropy (Basel).* 2023; 25(6): 850. doi: 10.3390/e25060850
3. Simulating protein folding in variable environmental conditions. Ed: Irena Roterman-Konieczna, Elsevier, 2026. ISBN: 9780443404757.
4. <https://shop.elsevier.com/books/simulating-protein-folding-in-variable-environmental-conditions/roterman-konieczna/978-0-443-40475-7>
5. HM. Berman, J. Westbrook, Z. Feng, G. Gilliland, T.N. Bhat, H. Weissig, I.N. Shindyalov, P.E. Bourne. The Protein Data Bank *Nucleic Acids Research* 2000, 28: 235-242: <https://doi.org/10.1093/nar/28.1.235>

Spatial Model of the Presynaptic Bouton for Neurotransmitter Transport Simulation by Using Cellular Automaton

Maciej Gierdziewicz¹, Robert-Mihai Stama²

¹ Department of Applied Computer Science, AGH University of Krakow, Al. Mickiewicza 30, Kraków, Poland

² National University of Science and Technology POLITEHNICA Bucharest, Splaiul Independenței 313, Sector 6, Bucharest, 060042, Romania

gierdzma@agh.edu.pl, robert_mihai.stama@stud.fim.upb.ro

Keywords: presynaptic bouton, neurotransmitter, mesh quality, cellular automaton

1. Introduction

Design of models of working synapses has a long history [1]. It has been noticed that cellular automaton (CA) could be used to model diffusion and exocytosis-like process [2]. High quality mesh is required for finite element method (FEM) [3]; 3D mesh cells should be (almost) regular tetrahedra but, in reality, some of them are of worse quality. Such a mesh, with internal structure sufficient, however, for a CA algorithm, was designed in this paper.

2. Description of the problem

It is difficult to model organelles or inner cell complexes with cubic meshes. On the other hand, FEM applied to a tetrahedral mesh needs a system of equations to be solved during each iteration, accounting also for vectors tangential to surfaces. This is the reason for simulating flow phenomena by using CA: this method needs neither very good quality meshes nor complex algorithms. Modeling neurotransmitter (NT) reaction and diffusion by CA may improve calculations. So far, it has been studied in a simplified presynaptic bouton model [2].

3. Related work

Tetrahedral mesh of a presynaptic bouton was used for a complex NT flow modeling with differential equations [1]. On the other hand, a realistic mesh, especially with synaptic vesicles, used too much computer memory [4]. CA was applied in [2] but there the bouton was too simple to study NT density in a realistic way. In conclusion, more detailed mesh is needed.

4. Solution of the problem

We made the model of a presynaptic bouton ($V=0.717 \mu\text{m}^3$), consisting of upper part with optional reuptake zone, middle part, primary reuptake zone, and active zone, by using gmsh software. Another mesh was made by passing the surface mesh on to TetGen. Calculations in TetGen and python were performed on the Ares computer in Cyfronet. The first mesh had 4611 elements with mean volume $1.55 \times 10^{-4} \mu\text{m}^3$, mean surface $0.0229 \mu\text{m}^2$ and quality parameters: EH (maximum edge / minimum tetrahedron height) (1), ER (maximum edge / inscribed sphere radius) (2), and SV (based on surface and volume) (3).

$$EH = e_{\max} / H_{\min} \quad (1)$$

$$ER = e_{\max} / r_{\text{ins}} \quad (1)$$

$$SV = S^{1/2} / V^{1/3} \quad (2)$$

The EH quality measure (1) had a quality comparable to ER [5]; the equation $EH=0.334ER-0.235$ was fitted with $R^2>0.9$. For our bouton mesh it was $EH=0.327ER-0.109$ (gmsH) and $EH=0.420ER-0.920$ (TetGen), both with $R^2>0.95$. For some types of tetrahedra, however, ER and EH behave differently. The plot in Fig.1 (coordinates in micrometers) is the projection of ER/EH on xy plane. It suggests that wedge-like tetrahedra with two close vertices are better detected with EH than with ER. The SV measure proved less sensitive.




Fig. 1. Distribution of ER/EH for tetrahedra with base on xy plane and 4th vertex in hemisphere.

5. Concluding remarks

Our 3D model of a presynaptic bouton of a neuron contains enough details to be used in simulation of synaptic processes. However, its moderate quality, that only to a limited extent enables its application in analysis performed with partial differential equations, makes it a suitable data structure for calculations performed with cellular automata.

Acknowledgements. The numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grants plgneuron2024 and plgneuron2025.

References

1. M.M. Knoedel et al.: Synaptic bouton properties are tuned to best fit the prevailing firing pattern. *Frontiers in Computational Neuroscience*, 8:101, pp. 1-16, 2014.
2. A. Bielecki and M. Gierdziewicz: Cellular automaton approach to estimation of neurotransmitter flow parameters in a presynaptic bouton of a neuron, in *Proc. KUKDM 2025*, pp. 91-92.
3. M. Gierdziewicz: Simulation of processes and structures in the synapse in the context of tetrahedral mesh quality. *Computers and Mathematics with Applications* 145:58-64, 2023.
4. A. Bielecki, M. Gierdziewicz, P. Kalita, and K. Szostek: *Construction of a 3D geometric model of a presynaptic bouton for use in modeling of neurotransmitter flow*. *Proc. Of the XVth ACC Cyfronet AGH HPC Users' Conference*, 13-15 March 2024, 33-34, 2024.
5. M. Gierdziewicz: Relations between geometric parameters and numerical simulation accuracy in modeling signal transmission in the presynaptic bouton. *Applied Sciences* 11:2811, 2021.

Mechanistic Insights into the Cycloaddition of (2E,4E)-2,5-Dinitro-2,4-hexadiene with Diazomethane from Density Functional Theory

Karolina Kula^{1,*}, Radomir Jasiński¹

¹ Cracow University of Technology, Faculty of Chemical Engineering and Technology, Warszawska 24, 31-155 Kraków, Poland

karolina.kula@pk.edu.pl, radomir.jasinski@pk.edu.pl

Keywords: Nitrodiene, Cycloaddition, Molecular Mechanism, Density Functional Theory

1. Introduction

The chemistry of conjugated nitrodienes has attracted increasing interest in recent years. These compounds are widely used in cycloaddition reactions, notably in Diels-Alder processes, enabling the efficient synthesis of six-membered rings. In addition, nitrodienes can participate in (3+2) cycloadditions, leading to the formation of bis-adducts. Moreover, the presence of a nitro group offers further opportunities for post-synthetic modification of the products [1].

2. Description of the problem

In 2015, *Sharko et al.* [2] reported an experimental study on the reaction between (2E,4E)-2,5-dinitro-2,4-hexadiene (1) and diazomethane (2) (Fig. 1). A double (3+2) cycloaddition (CA) selectively affords only one bis-pyrazoline, which is unstable and undergoes subsequent [1,5]-H shift (HAT) and HNO₂ extrusion, probably determined by elimination of nitro group at the first stage. However, the mechanistic aspects of these transformations remain unclear.




Fig. 1. Theoretically possible reaction paths between nitrodiene (1) and diazo compound (2).

3. Related work

In recent years, studies have begun to emerge highlighting the interesting electronic properties of conjugated nitrodienes and their unique reactivity [3,4]. For example, CA of (1E,3E)-1,4-dinitro-1,3-butadiene with an azomethine ylide showed the formation of a single cycloadduct rather than a double cycloaddition product [5].

4. Solution of the problem

All calculations were carried out within the framework of Molecular Electron Density Theory (MEDT), using ω B97XD/6-311G(d,p) and B3LYP/6-31G(d) computational levels.

All localized stationary points were characterized by vibrational analysis. Starting materials and products exhibited positive Hessians, while all transition states displayed a single negative eigenvalue. IRC calculations were performed for all optimized transition states to confirm their connection to the corresponding reactant and product minima. Solvent effects were simulated using SCRF approach within the PCM model.

First, analyses of the Electron Localization Function and Conceptual Density Functional Theory for the reagents were carried out. Next, theoretically possible reaction pathways were explored, together with an analysis of key critical structures.

5. Conclusions of the problem

Analysis of the reactivity descriptors of molecules for presented reactions suggests that nitrodiene as well as both of nitrovinyl pyrazolines will participate as electrophile, while studied diazomethane will play a role of nucleophilic agent. Kinetic and thermodynamic aspects and analysis of critical structures indicates that the creation of bis-pirazoline in pmr-type double cycloaddition is realised *via* one-step polar asynchronous mechanism. In turn, the conversion to bis-pirazol occurs *via* sequence of competitive reactions of one-step HNO₂ elimination and one-step [1,5]-H shift. Both of these transformations are performed *via* non-polar asynchronous one-step mechanisms, without ionic intermediates [6].

Acknowledgements. This research has been supported by PLGrid. The numerical experiment was possible through computing allocation on the Ares system at ACC Cyfronet AGH under the grants PLG/2024/017868.

References

1. M. Sadowski, K. Kula: "Nitro-functionalized analogues of 1,3-butadiene: An overview of characteristic, synthesis, chemical transformations and biological activity", in *Curr. Chem. Lett.*, 2023, pp. 15.
2. A.V. Sharko, G.A. Senchyk, E.B. Rusanov, K.V. Domasevitch: "Preparative synthesis of 3(5),3'(5')-dimethyl-4,4'-bipyrazole", in *Tetrahedron Lett.*, 2015, pp. 6089.
3. M. Sadowski, B. Synkiewicz-Musialska, K. Kula: "(1E,3E)-1,4-Dinitro-1,3-butadiene - Synthesis, Spectral Characteristics and Computational Study Based on MEDT, ADME and PASS Simulation", in *Molecules*, 2024, 542.
4. K. Kula, E. Kuś: "In Silico Study About Substituent Effects, Electronic Properties, and the Biological Potential of 1,3-Butadiene Analogues", in *Int. J. Mol. Sci.*, 2025, 8993.
5. M. Sadowski, K. Kula: "Unexpected Course of Reaction Between (1E,3E)-1,4-Dinitro-1,3-butadiene and N-Methyl Azomethine Ylide - A Comprehensive Experimental and Quantum-Chemical Study", in *Molecules*, 2024, 5066.
6. K. Kula, R. Jasiński: "Synthesis of bis(het)aryl systems via domino reaction involving (2E,4E)-2,5-dinitrohexa-2,4-diene: DFT mechanistic considerations", in *Chem. Heterocycl. Compd.*, 2024, 6089.

Language Driven Therapy Design in Predictive Oncology

Filip Ręka

AGH University of Krakow

filipreka@agh.edu.pl

Keywords: Conversational Modeling, Large Language Models, Computational Oncology, Treatment Planning

1. Introduction

Modern oncology faces a fundamental challenge: clinical knowledge is articulated in narrative form, while treatment planning demands executable computational models. We introduce a web-based tool implementing the Language-Driven Therapy Design (LDTD) framework, enabling researchers to build and interact with predictive tumor models using natural language - supporting multiple LLM backends from Anthropic, Google, OpenAI, and open-source providers via vLLM.

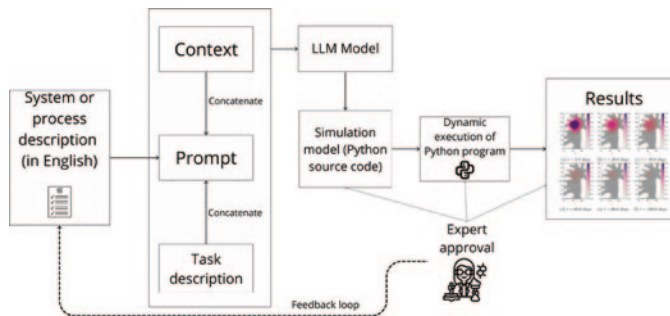


Fig. 1. LDTD workflow: narrative description → LLM translation → Python simulation → results.

2. Description of the problem

Translating biological narratives into spatiotemporal tumor simulations has historically required rare, cross-disciplinary expertise in cancer biology and applied mathematics. A clinician's intuition must be converted into coupled partial differential equations with numerical solvers, boundary conditions, and code. This translation bottleneck limits access to advanced computational oncology and slows hypothesis exploration, particularly in data- and knowledge-constrained settings. LMMs can potentially bridge the gap between those two domains.

3. Related work


Computational modeling has become an established tool in cancer research [3]. Mechanistic tumor models such as TumorTwin [1] achieve high spatial accuracy but require expert parameterization and remain inaccessible to most clinicians. Hybrid neural-physical approaches partially address scalability yet do not resolve the knowledge-translation barrier [5]. Maeda et al. [2] demonstrated that LLMs can generate SBML kinetic models from natural language, establishing a precedent for LLM-driven formalization. More broadly, LLMs have been applied to clinical oncology decision support [4], but their use as autonomous translators of biological narratives into executable spatiotemporal simulation code remains unexplored.

4. Solution of the problem

We developed a multi-LLM web interface implementing the LDTD pipeline (Fig. 1). The algorithm proceeds in four stages: (i) the user provides a natural-language description of a tumor system and therapeutic protocol; (ii) the LLM proposes governing equations (reaction-diffusion PDEs coupled with ODE-based pharmacokinetics); (iii) it generates executable Python simulation code using NumPy; (iv) results are executed server-side and returned as spatiotemporal outputs. The tool supports iterative refinement through conversational feedback and is LLM-agnostic - integrating APIs from proprietary models and any OpenAI-compatible endpoint. Open-source models were served locally via vLLM on the Athena supercomputer (ACC Cyfronet AGH) using NVIDIA A100 GPUs.

We benchmarked the approach against TumorTwin - a validated 3D reaction-diffusion framework on two cancer types: high-grade glioma (HGG) with concurrent chemoradiation and triple-negative breast cancer (TNBC) with chemotherapy. The LLM (Qwen3 family) was provided only with clinical narratives, treatment protocols, and patient-derived MRI initial conditions. As shown in Fig. 2, the LLM-generated models closely match the total tumor cell count trajectories of both TumorTwin and ground-truth measurements throughout treatment.

Fig. 2. Cell count comparison between real-world data, TumorTwin model and LLM-generated model.



5. Conclusions

These findings demonstrate that the long-standing translation barrier between narrative oncology and computational modeling is already fracturing. The LDTD tool lowers this barrier substantially, enabling domain experts with limited computational backgrounds to construct, run, and refine mechanistic tumor models through natural conversation.

Acknowledgements. Supported by the National Science Centre (NCN), Poland, grant OPUS-29, DEC-2025/57/B/ST6/04377 and PLGrid Cyfronet grant PLG/2025/018971. LLM tools were used to assist in editorial refinement of this paper and as experimental instruments for narrative-driven model generation.

References

1. Kapteyn et al. TumorTwin: A Python framework for patient-specific digital twins in oncology (2025).
2. Maeda et al. Automatic generation of SBML kinetic models from natural language texts using GPT. *IJMS* 24(8), 7296 (2023).
3. Scibilia et al. Mathematical oncology: How modeling is transforming clinical decision making. *Cancer Research* 85(24) (2025).
4. Chen et al. Large language models in oncology: a review. *BMJ Oncology* 4(1) (2025).
5. Dzwiniel et al. Supermodeling in predictive diagnostics of cancer under treatment. *Computers in Biology and Medicine* 137 (2021).

DFT Calculations of Flavin Derivatives in Their Anionic Forms

Iwona Gulaczyk¹, Dorota Prukala¹, Ekaterina Zubova², Radek Cibulka² and Marek Sikorski¹

¹ Faculty of Chemistry, Adam Mickiewicz University in Poznań, Uniwersytetu Poznańskiego 8, 61-614 Poznań, Poland

² Department of Organic Chemistry, University of Chemistry and Technology, Prague, Technická 5, 16628 Prague 6, Czech Republic

gulai@amu.edu.pl, dorota.prukala@amu.edu.pl, zubovae@vscht.cz, cibulkar@vscht.cz, sikorski@amu.edu.pl

Keywords: Flavin anions, 3-methyllumichrome, Photoredox catalysis, Absorption, Fluorescence, Fluorescence lifetime, Azaalloxazines

1. Introduction

Flavin derivatives are widely regarded as efficient catalysts in photoredox chemistry. However, their application in reductive photocatalysis is limited, because the transient flavin radicals exhibit low stability. The use of excited organic anions as photoredox catalysts provides several advantages over the neutral molecules. It is important to investigate the photophysical and chemical properties of the anionic forms of flavin derivatives using both experimental and theoretical approaches. Based on these findings, new photocatalytic systems that employ excited flavin anions can be developed, with particular emphasis on photoreductions that extend beyond the current limits of photoredox catalysis. These goals can be accomplished through integrated efforts that combine flavin synthesis and photocatalysis with studies of flavin photophysics and quantum-chemical calculations.




Fig. 1. Left: Isoalloxazines and alloxazines involved in the study. 2a-H represents a hypothetical structure, which was investigated using DFT calculations. Right: Electrostatic surface potential (ESP) for the anion formed from 1a-H (calculated from the optimized geometries using the DFT-B3LYP method and the basis function aug-cc-pVTZ).

2. Description of the problem

The hypothesis was put forward that employing the anion of the oxidized form could provide a new strategy for incorporating flavins into photoreductive chemistry. The use of excited anions in reduction processes is advantageous, as the negative charge facilitates oxidation and enhances the reducing power of the catalyst¹⁻³. In addition, the red shift in the absorption spectrum upon deprotonation enables the use of visible light instead of high-energy UV radiation. In this study, the behaviour of 3-methyllumichrome 1a-H and its derivatives was examined in detail, with the goal of identifying conditions that enable its conversion into the desired anion

2a⁻ under photoreduction-compatible conditions. This investigation employed a combination of techniques, including UV–Vis and fluorescence spectroscopy, NMR, electrochemistry, and quantum-chemical (DFT) calculations. This integrated approach confirmed the predominant isoalloxazinic character of the anionic species and to demonstrate that its photophysical and electrochemical properties are well suited for applications in photoredox catalysis.

3. Related work

This work is the first to demonstrate how to generate and simultaneously utilise the anion of a hypothetical isoalloxazine unsubstituted at position 10 in photoredox catalysis. It is also a demonstration of a new way, how to switch the oxidized flavins, still exclusively used as oxidants, to a stable reductive species.

4. Solution of the problem

For use in photoredox catalysis, 3-methylalumichrome (1a-H) was chosen as the model compound. It contains a single acidic hydrogen at the N(1)–H position, which enables tautomerization to 2a-H. For comparison, the alloxazine (3a-H) and isoalloxazine (4a-H) structures were also examined; both possess an acidic N(3)–H group. In addition, derivative 5a, which lacks any acidic hydrogen, was included (Figure 1). Tetrabutylammonium acetate (TBAOAc) was employed as a non-nucleophilic base. Its addition to an acetonitrile solution of 1a-H led to pronounced changes in the UV–Vis spectra. The characteristic absorption bands of the alloxazinic form were replaced by a new band typical of isoalloxazine structures. A substantial alteration was observed in the fluorescence spectra. In contrast, only minor changes in the UV–Vis spectra and almost no shift in the fluorescence maximum were detected for 3a-H after the addition of TBAOAc. No spectral changes were observed for the less acidic isoalloxazine 4a-H and 5a. Based on DFT calculations focusing on electron density distribution, an approximately equal contribution of the relevant resonance forms was proposed.

5. Conclusions

It was confirmed that an alloxazine such as 1a-H undergoes deprotonation to form the anion 2a⁻. This species displays isoalloxazinic spectral characteristics, as reflected in its absorption, fluorescence spectra and fluorescence lifetime. Owing to its anionic nature, 2a⁻ becomes a strongly reducing species when excited with blue or cyan light. This work⁴ represents the first demonstration of how to generate and apply the anion of a hypothetical isoalloxazine unsubstituted at position 10 in photoredox catalysis. It introduces a new strategy for converting oxidized flavins into stable, strongly reducing species. The high reducing power of the excited anion 2a⁻ makes photocatalytic systems based on this species promising candidates for practical applications.

Acknowledgements. This work was supported by the Czech Science Foundation (Grant No 24-11386K) and by the research grant WEAVE-UNISONO UMO-2023/05/Y/ST4/00062, from The National Science Centre of Poland (NCN). The numerical experiment was possible through computing allocation on the Ares supercomputer system at Cyfronet under the grant no. PLG/2024/017625.

References

1. B. Bartolomei, G. Gentile, C. Rosso, G. Filippini and M. Prato, *Chem. Eur. J.*, 2021, 27, 16062–16070.
2. S. Wu, F. Schiel and P. Melchiorre, *Angew. Chem. Int. Ed.*, 2023, 62, e202306364.
3. M. Schmalzbauer, M. Marcon and B. König, *Angew. Chem. Int. Ed.*, 2021, 60, 6270–6292.
4. Prukała, D., Zubova, E., Svobodová, E., Šimková, L., Varma, N., Chudoba, J., Ludvík, J., Burdzinski, G., Gulaczyk, I., Sikorski, M., & Cibulka, R. *Chem. Sci.*, 2025, 16, 11255–11263.

Domain Bias in Deep Learning for Cytology

Jan Krupiński¹, Ernest Jamro², Maciej Wielgosz², Agnieszka Dąbrowska-Boruch²

¹ Cracow University of Technology, Warszawska 24, 31-155 Kraków, Poland

² AGH University of Krakow, al. Mickiewicza 30, 30-059 Kraków, Poland

jan.krupinski@pk.edu.pl, {jamro, wielgosz, adabrow}@agh.edu.pl

Keywords: domain bias, cytology, deep learning, skin cancer

1. Introduction

Cytology enables diagnosis from cell samples by analyzing individual cell morphology within a whole image context. Though less accurate than histopathology, it is minimally invasive, fast, and practical. Deep learning (DL) can automate this process through cell detection and classification using convolutional neural networks, showing potential for clinical use.

2. Description of the problem

Although strong results were achieved on the initial dataset (96.7% accuracy), gathered by a limited number of veterinarians, performance dropped substantially on external samples (51.4% accuracy). This is consistent with DL models trained on homogeneous, in-house datasets often underperforming in clinical settings due to variability in staining protocols, imaging equipment, and cellular composition [1]. Such behavior is known as domain bias: models overfit to domain-specific features such as color, texture, or morphology, limiting generalizability across institutions and imaging conditions.

3. Related work

The most common ways to improve model generalizability include data augmentation and tuning the model architecture. Domain adaptation methods such as Maximum Mean Discrepancy (MMD) [2] and CORAL [3] aim to align feature distributions across domains.

4. Solution of the problem

Table 1. Domain bias estimated via Proxy A-distance [4] across models (higher = greater bias).

Model	No Augmentation	Data Augmentation
Histogram Features	1.40	0.32
Vit-L (DINOv2) [5]	0.93	0.54
YOLOv8 [6]	1.42	0.80
YOLOv8 + MMD	0.99	0.42

Multiple deep learning models were compared and evaluated for domain bias (selected models shown in table 1). Standard data augmentation pipelines were employed but found insufficient: while influencing in-domain performance, they failed to bridge the inter-domain gap. To address this, MMD was applied at multiple network depths to align feature distributions during training. Similarly, multi-task learning combining cell detection and image-level classification proved beneficial (Figure 1). Models were implemented in PyTorch and trained on NVIDIA A100 GPUs, with development conducted via VSCode remote sessions over SSH tunnels.

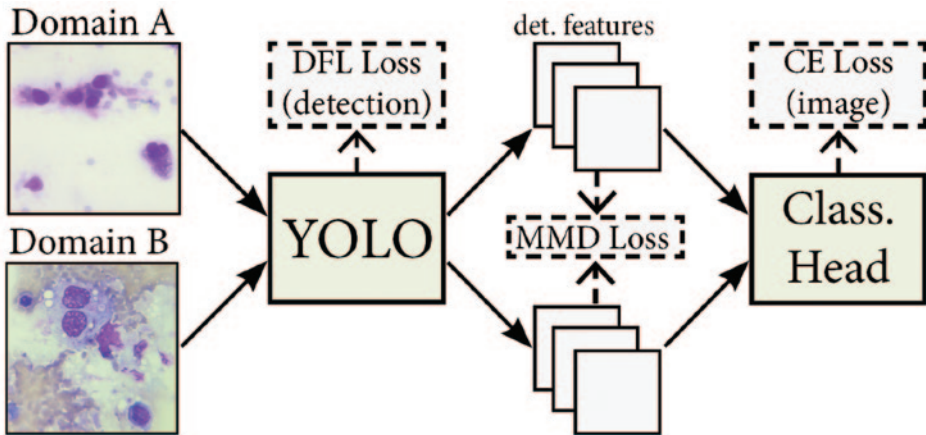


Fig. 1. Multi-task learning pipeline, combining YOLOv8 and an image classification head.

5. Conclusions

Conventional augmentations cannot correct for domain bias in cytological image analysis. Multi-task learning with MMD loss proved necessary to achieve meaningful domain alignment, highlighting the need for explicit distribution-matching methods in cytology tasks. Image classification on the in-house dataset improved to 97.3% and on the external data to 57.9%.

Acknowledgements. The numerical experiment was possible through computing allocation on the Athena system at ACC Cyfronet AGH under the grant plgdyplomanci7.

References

1. H. Guan, and M. Liu: "Domain Adaptation for Medical Image Analysis: A Survey.", IEEE Transactions on Biomedical Engineering, 2022, 69(3), 1173-1185.
2. A. Gretton et al., "A Kernel Two-Sample Test.". Journal of Machine Learning Research, 2012, 12(25), 723-773.
3. B. Sun and K. Saenko, "Deep CORAL: Correlation Alignment for Deep Domain Adaptation.", ECCV 2016 Workshops, Lecture Notes in Computer Science, vol. 9915, Springer.
4. S. Ben-David et al., "A Theory of Learning from Different Domains", Machine Learning, 2010, 79, 151-175.
5. M. Oquab et al., "DINOv2: Learning robust visual features without supervision." arXiv preprint arXiv:2304.07193.
6. G. Jocher, J. Qiu and A. Chaurasia: "Ultralytics YOLO (Version 8.0.0)": <http://github.com/ultralytics/ultralytics>

Computer Vision Pipelines for High-Content Microscopy Reveal Heritable Single-Cell Immune States

Truong Co Nguyen*, Juan Alfonso Redondo, Marek Kočańczyk, Paweł Paszek

Department of Biosystems and Soft Matter, Institute of Fundamental Technological Research, Polish Academy of Sciences, Pawińskiego 5B, 02-106 Warsaw, Poland

conguyen@ippt.pan.pl

Keywords: high-content microscopy, image analysis pipeline, single-cell phenotyping, macrophages, bootstrap resampling, heritability

1. Introduction

High-content microscopy quantifies cellular responses at single-cell resolution, but reproducible measurements from large image datasets require efficient, standardized, and well-controlled computational workflows [1]. We present a scalable computer-vision pipeline for multi-image processing and phenotyping of clonal macrophage populations under innate immune stimulation and bacterial infection.




Fig. 1. Single cell representation of the clonal population.

2. Description of the problem

Large-scale microscopy produces heterogeneous images with illumination artifacts, segmentation errors, and outlier pixels, complicating robust intensity measurement across plates and

generations. A central challenge is to extract reliable single-cell features and statistically sound measures of immune-state inheritance from mixed and clonally expanding populations (Fig.1).

3. Related work

Existing high-content analysis (HCA) workflows primarily emphasize segmentation, tracking, and high-dimensional feature extraction [2,3]. Although newer platforms have expanded support for spatial analysis, integrated support for clone-aware phenotyping and density-aware spatial modeling remains comparatively underdeveloped across general-purpose workflows [4,5].

4. Solution of the problem

Operetta CLS images are flat-field corrected and stitched, then processed by nucleus-based segmentation and fluorescence quantification. Post-processing includes morphology-based quality control, perinuclear region generation, clone delineation, and robust percentile-based intensity estimation. Downstream analyses comprise plate-specific responder classification from untreated controls, local cell-density estimation, spatial identification of daughter-cell pairs for odds-ratio analysis, and bootstrap-based virtual clone generation to quantify heritable expression patterns across generations. The workflow scales to thousands of clonal populations and links heritable immune traits to infection susceptibility.

5. Conclusions

The proposed pipeline provides an end-to-end, clone-aware computer-vision workflow that yields robust single-cell measurements and enables statistically rigorous inference of immune-state inheritance at scale, supporting mechanistic links between heritable expression programs and infection outcomes.

Acknowledgements. This work is supported by the Polish National Agency for Academic Exchange (grant no. BPN/PPO/2022/1/00002/DEC/1) and the National Science Center, Poland (grant no. 2022/45/B/NZ6/01643). We gratefully acknowledge the Polish high-performance computing infrastructure PLGrid (HPC Centers: ACC Cyfronet AGH, CI TASK) for providing computer facilities and support within computational grant no. PLG/2025/018818.

References

1. Redondo, J. A., & Paszek, P. (2025). AI-Assisted Microscopy for Infection Biology: Advances in High-Content Imaging of Host-Pathogen Interactions. *Computer Assisted Methods in Engineering and Science*, 32(4), 355–365.
2. Stirling DR, Swain-Bowden MJ, Lucas AM, Carpenter AE, Cimini BA, Goodman A. CellProfiler 4: improvements in speed, utility and usability. *BMC Bioinformatics*. 2021;22(1):433.
3. Bray MA, Singh S, Han H, Davis CT, Borgeson B, Hartland C, et al. Cell Painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes. *Nat Protoc*. 2016;11(9):1757-1774.
4. Magness A, Colliver E, Enfield KSS, Lee C, Shimato M, Daly E, et al. Deep cell phenotyping and spatial analysis of multiplexed imaging with TRACERx-PHLEX. *Nat Commun*. 2024;15(1):5135.
5. Morgan D, Jost TA, De Santiago C, Brock A. Applications of high-resolution clone tracking technologies in cancer. *Curr Opin Biomed Eng*. 2021; 19:100317.

Cross-Polarized Skin Image Synthesis from Non-Polarized Inputs Using Deep Generative Models

Paweł R. Popielski^{1,2}, Sławomir Wilczyński³, Jarosław Żmudzki⁴, Kamil Kwieciński⁵

¹ University of Silesia in Katowice, Faculty of Science and Technology, Institute of Biomedical Engineering, Bankowa Street 12, 40-007 Katowice, Poland

² AI-Med Lab – Laboratory of Artificial Intelligence in Medical Imaging and Diagnostics, Bedzinska Street 39, 41-200 Sosnowiec, Poland

³ Medical University of Silesia, Faculty of Pharmaceutical Sciences, Katowice, Poland

⁴ Silesian University of Technology, Faculty of Mechanical Engineering, Department of Engineering Materials and Biomaterials, 18a Konarskiego Street, 44-100 Gliwice, Poland

⁵ Aipharma Sp. z o.o., Katowice, Poland

{pawel,popielski}@us.edu.pl

Keywords: GAN, pix2pix, skin imaging, cross-polarization, medical image translation, deep learning

1. Introduction

Cross-polarized skin imaging reveals diagnostically relevant subsurface structures that are often obscured in conventional non-polarized photographs. Recent progress in generative deep learning, especially conditional GANs and image-to-image translation frameworks, has shown strong potential for reconstructing missing visual information from related image domains [1,3,8]. These advances motivate the use of learned transformation models for dermatological imaging tasks [9,10].

2. Description of the problem

The problem addressed in this work is the synthesis of cross-polarized skin images from standard non-polarized inputs. The main challenge is to recover diagnostically meaningful visual patterns, including structures masked by specular reflection, although such information is not explicitly available in the source image.

3. Related work

Image-to-image translation methods such as pix2pix and CycleGAN have been widely used in visual domain transfer problems [1,2]. In medical imaging, deep neural networks have also shown strong performance in skin analysis and related biomedical tasks [5,9,10]. However, transformation between non-polarized and cross-polarized skin images remains insufficiently explored.

4. Solution of the problem

We propose a supervised image-to-image translation pipeline based on the pix2pix framework [1]. The model is trained on paired datasets containing corresponding non-polarized and cross-polarized skin photographs. The solution combines adversarial learning [3], L1 reconstruction loss, and feature-oriented optimization inspired by perceptual approaches [6]. The implementation is developed in PyTorch and executed on GPU-based high-performance computing infrastructure available through PLGrid. The computational environment enables efficient training, hyperparameter tuning, and large-scale experimentation on dermatological image datasets.

5. Conclusions

Preliminary experiments indicate that GAN-based models can generate cross-polarized skin images with promising visual fidelity and clinically relevant structural detail. The proposed approach may support wider access to enhanced dermatological imaging without requiring dedicated polarization hardware.

Acknowledgements. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Center: ACC Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2025/018905.

References

1. P. Isola, J.-Y. Zhu, T. Zhou, A.A. Efros: “Image-to-Image Translation with Conditional Adversarial Networks”, CVPR, 2017.
2. J.-Y. Zhu, T. Park, P. Isola, A.A. Efros: “Unpaired Image-to-Image Translation using CycleGAN”, ICCV, 2017.
3. I. Goodfellow et al.: “Generative Adversarial Nets”, NeurIPS, 2014.
4. K. He, X. Zhang, S. Ren, J. Sun: “Deep Residual Learning for Image Recognition”, CVPR, 2016.
5. O. Ronneberger, P. Fischer, T. Brox: “U-Net: Convolutional Networks for Biomedical Image Segmentation”, MICCAI, 2015.
6. J. Johnson, A. Alahi, L. Fei-Fei: “Perceptual Losses for Real-Time Style Transfer and Super-Resolution”, ECCV, 2016.
7. M. Arjovsky, S. Chintala, L. Bottou: “Wasserstein GAN”, ICML, 2017.
8. T.-C. Wang et al.: “High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs”, CVPR, 2018.
9. A. Esteva et al.: “Dermatologist-level classification of skin cancer with deep neural networks”, Nature, 2017.
10. N.T. Pham, J. Park, S. Lee: “Deep Learning-Based Skin Lesion Analysis: A Review”, IEEE Access, 2020.

Author Index

- Abramowicz M. 33
Andreini C. 81
Andrzejak M. 79
Attucci T. 81
- B**
Baerle V. 81
Balawajder T. 35, 63
Baliś B. 15, 37
Benedykciński R. 59
Białas P. 45, 47
Bonarek P. J. 71
Bubak M. 10
- Capała K. 49
Cholewa M. 35
Choraży S. 71
Christensen J. B. 31, 59
Cibulka R. 91
Ciupek D. 25
- D**
Dąbrowska-Boruch A. 93
Dąbrowska-Wierzbicka P. 35
Długosz M. 41
Dubiel Ł. 53
Dułak D. 83
Dutka Ł. 37
Dzwinel W. 53
- Eilmes A. 65
- Flis Ł. 35
- G**
Gierdziewicz M. 85
Głomski S. 39
Gómez del Pozo E. 29
Grzanka L. 31, 59
Gulaczyk I. 91
- Halliday I. 19
- J**
Jamro E. 93
Janiuk A. 13, 57
Janiuk I. 13
Jankowski R. 71
Jasiński R. 87
Jeziorek K. 41
- K**
Kaliński J. 59
Kapusta M. 57
Karczewski J. 37
Kayanikhoo F. 61
Kica P. 17, 27, 37
Kitowski J. 37, 53
Klepczarek Ł. 13
Kluźniak W. 33, 61
Kobzar O. 51
Kochańczyk M. 95
Kocot J. 35, 63
Komorkiewicz M. 41
Konieczny L. 83
Kopeć W. 73
Korcyl P. 45, 47
Kosińska J. 39
Kościółek T. 17
Kowalczyk M. 41
Krupiński J. 93
Kryjak T. 41
Kryza B. 37
Kubisiak P. 65
Kukułka M. 69
Kula K. 87
Kurdziel M. 15
Kwieciński K. 97
- L**
Lančová D. 33, 61
Lasocki S. 63
Lelek T. 15
Leśniak M. 35, 63
Leśnodorska A. 63
Liput J. 37
- Ł**
Łapczuk A. 75
- M**
Majerz E. 53
Makuch M. 35, 63
Malawski M. 17, 19, 21, 25, 27, 29
Marchewka D. 41
Mazurek S. 35
Meizner J. 23
Michalak A. 67, 69
Mróz J. 73
Mtupa-Ndiaye A. 63
Mytnik A. 35

Nalewajko K. 55
 Narracott A. 19, 21
 Nguyen T. C. 95
 Nicotri C. 65
 Nikolow D. 37
 Nowak A. 23
 Nowakowski P. 23

Opioła Ł. 37
 Orlecka-Sikora B. 63
 Orzechowski M. 17, 37
 Otta M. 21, 27
 Otwinowski J. 53

Paszek P. 95
 Piątkowski M. 77
 Pięciak T. 25
 Płonka P. 13, 57
 Popielski P. R. 97
 Pukała D. 91

Raczek A. 37
 Raczyńska A. 81
 Radzik W. 79
 Redondo J. A. 95
 Rejda M. 31
 Ręka F. 89
 Roterman-Konieczna I. 83
 Rózga M. 67
 Różańska A. 33
 Rybski S. 31
 Rycerz K. 49

Siejkowski H. 35, 63
 Sikora M. 73
 Sikorski M. 91
 Skruch P. 41
 Słota R. G. 37
 Stafin K. 77

Stama R.-M. 85
 Stapor K. 83
 Stebel T. 45, 47
 Stefański A. 45
 Sto J. 35, 63
 Sułkowski B. 43
 Szaleniec M. 81
 Szelest M. 41
 Szmelich W. 37

Śliwa P. 77

Tlałka K. 19, 29
 Török G. 33

Urbanec M. 61

Vondrák V. 11

Wang J. 71
 Wenda B. 35, 63
 Wielgosz M. 93
 Wielgus M. 33, 61
 Wilczyński S. 97
 Wołk A. 49
 Wzorek P. 41

Zajac K. 17, 21, 29
 Zajdel P. 39
 Zapolski D. 45, 47
 Zasada M. 15, 37
 Zdunik L. 61
 Zhyhulin T. 23
 Zielińska K. 17
 Zimmermann O. 9
 Zubova E. 91

Żmudzki J. 97
 Żurowska O. 69

Published by

ACC Cyfronet AGH
ul. Nawojki 11
30-950 Kraków
www.cyfronet.pl



ISBN 978-83-61433-51-4

