








# High-Performance Computing at HSE University

Pavel Kostenetskiy<sup>(✉)</sup> , Vyacheslav Kozyrev , Roman Chulkevich ,  
Mariia Salikova , and Roman Mishenin 

HSE University, Moscow, Russia  
pkostenetskiy@hse.ru

**Abstract.** High-performance computing (HPC) has emerged as a critical tool for accelerating research across diverse scientific domains, enabling the efficient processing of large datasets and complex simulations. This article offers a comprehensive overview of the HPC resources available at HSE University in Moscow. We outline the university's current HPC infrastructure, detailing its computational capabilities, software environments, and key recent upgrades implemented to address the evolving needs of researchers. Furthermore, we showcase how these resources are being leveraged to support cutting-edge research projects across multiple disciplines. By highlighting available opportunities, this review seeks to inform the academic community and inspire innovative applications of HPC in both research and education.

**Keywords:** High-performance computing · HSE University · University research · HPC cluster · Supercomputer

## 1 Introduction

In today's rapidly evolving research landscape, access to advanced computational resources is crucial for driving innovation and solving complex challenges. As a leader in this domain, the National Research University Higher School of Economics (HSE University) has established a high-performance computing (HPC) cluster to support cutting-edge academic and educational initiatives. With computational power becoming increasingly vital, these resources empower researchers to conduct sophisticated simulations and analyze large-scale datasets across diverse disciplines.

At the heart of HSE University's HPC ecosystem lies a high-performance computing cluster that powers interdisciplinary research and education across the university. This article presents a comprehensive overview of HSE's HPC infrastructure, describing its architecture, technical capabilities, and diverse applications in supporting cutting-edge research and teaching. Through this examination, we demonstrate how HPC serves as a catalyst for scientific discovery, innovation, and academic excellence at HSE.

## 2 HPC Resource Configuration

The cHARISMa supercomputer is a high-performance computing cluster consisting of 48 compute nodes with an integrated parallel storage system (complete specifications provided in Table 1). The cluster infrastructure includes:

- 6 nodes with eight Nvidia A100 80 GB SXM GPUs each;
- 29 nodes with four Nvidia V100 32 GB SXM GPUs each (featuring high-memory configurations of 768 GB–1536 GB);
- 2 nodes with two Nvidia H100 80 GB PCIe GPUs each;
- 11 CPU-only nodes optimized for non-GPU workloads.

The system incorporates an 848 TB Lustre parallel file system with Nvidia GPUDirect Storage support, significantly accelerating I/O operations for neural network training. High-speed node interconnects are provided through a  $2\times 100$  GB/s InfiniBand EDR network, with workload management handled by the SLURM scheduling system.

**Table 1.** Technical specifications of cHARISMa HPC cluster

Parameter	Value
Total nodes	48
CPU sockets/CPU cores	96/2648
GPUs/GPU cores	168/984 064
CPU model	Intel Xeon Gold / AMD EPYC
GPU model	116×Nvidia V100 32 GB SXM (NVLink) 48×Nvidia A100 80 GB SXM (NVLink) 4×Nvidia H100 80 GB PCIe
RAM	53.9 TB
Disk storage	1256 TB
Parallel data storage system	Lustre (848 TB usable)
Computer network type	InfiniBand EDR (2×100 Gb/s)
Management network type	Gigabit Ethernet
Peak performance (FP64)	2.24 PFLOPS
LINPACK performance	975.6 TFLOPS
Peak performance (FP32)	3.27 PFLOPS
Peak AI performance (FP16)	36.2 AI-PFLOPS
Job scheduler	SLURM
Operating system	CentOS Linux

The cHARISMa supercomputer features six specialized node types, whose technical specifications are detailed in Table 2.

**Table 2.** Computational Node Specifications of cCHARISMa Cluster

Type	Number	CPU (cores)	GPUs	RAM (GB)	SSD (GB)	InfiniBand (GB/s)
A	16	2×Intel Xeon Gold 6152 (2×22)	4×Nvidia Tesla V100 32 GB SXM NVLink	768	2×240	2×100
B	10	2×Intel Xeon Gold 6152 (2×22)	4×Nvidia Tesla V100 32 GB SXM NVLink	1536	2×240	2×100
C	3	2×Intel Xeon Gold 6240R (2×24)	4×Nvidia Tesla V100 32 GB SXM NVLink	768	2×240	2×100
D	11	2×Intel Xeon Gold 6248R (2×24)	—	768	2×240	100
E	6	2×AMD EPYC 7702 (2×64)	8×Nvidia A100 80 GB SXM NVLink	1024	2×960	2×200
F	2	2×Intel Xeon Gold 6426Y (2×16)	2×Nvidia H100 80 GB PCIe NVLink bridge	512	2×960	100

The HSE Supercomputing Modeling Unit regularly updates the HPC cluster and application software configurations while performing scheduled maintenance and ongoing repairs. Below are the major achievements of the past year:

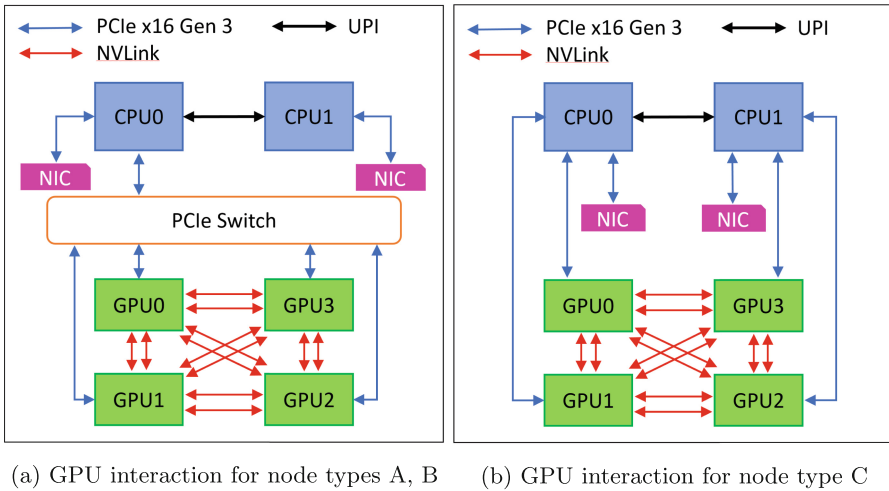
- implemented 24/7 supercomputing complex availability with round-the-clock technical support (including nights and weekends);
- installed new scientific software packages: Abinit, Stata, Julia, Apertium, CP2K, Schmutzi, Elegant, Amber, Matlab, ML-agents, DeepMD-Kit, Smilei, Vasp;
- installed OpenMPI 5.0.3;
- updated Nvidia drivers with CUDA 12.2 framework;
- updated SLURM job scheduler up to version 23.02;
- maintained storage of popular datasets based on user requests;
- deployed GPUDirect technologies (Copy + RDMA), achieving a 238% increase in GPU data exchange speed and a sevenfold reduction in data transfer latency;
- established connection to Russia’s National Research Computer Network;
- updated firmware for all control/compute nodes;
- configured new partitions in the SLURM job scheduler;
- developed two custom patches for the SLURM source that improved filesystem usage statistics and enabled tracking of GPU allocation to jobs;
- added two next-generation compute nodes equipped with Nvidia H100 GPUs;
- replaced more than 1000 lead-acid UPS batteries.

The HPC cluster currently supports more than 700 active scientific projects, enabling HSE University researchers to publish over 400 scientific papers—including 176 articles in top-tier scientific journals—through supercomputer-powered research.

To optimize computational efficiency across diverse workloads, the cHARISMa supercomputer features a heterogeneous GPU architecture designed for Deep Learning, Big-data analytics, computational physics simulations, and materials modeling. Table 3 provides a technical comparison of the installed GPU accelerators, detailing their specifications.

The efficient utilization of computation nodes is critical for optimizing high-performance computing workflows. On shared nodes, where multiple users execute tasks concurrently, a well-designed internal architecture ensures that GPU resources are fully leveraged to meet research objectives.

Figures 1, 2 and 3 present the GPU interaction schematics for node types A, B, C, E, and F. (Note: Type D nodes do not contain GPUs.)



**Fig. 1.** GPU interaction for node types A, B, and C.

Figure 1a presents the architectural schematics of type A and B compute nodes. In these configurations, GPU–CPU interaction is mediated through a PCIe switch (CPU1 lacks direct GPU access). The additional data routing layer introduces a modest performance overhead. Performance benchmarking revealed this architectural constraint has minimal practical impact. Exclusive execution on CPU0 cores showed a performance degradation not exceeding 2%.

Figure 1b details the architecture of type C compute nodes, which feature direct CPU–GPU interconnectivity for all processors, thereby achieving the highest performance among all V100-equipped nodes in the cluster.

**Table 3.** Technical specifications of installed GPU accelerators

Parameter	H100 80 GB PCIe	A100 80 GB SXM	A40 48 GB PCIe	V100 32 GB SXM	P40 24 GB PCIe
Architecture	Hopper	Ampere	Ampere	Volta	Pascal
FP64 (TFLOPS)	25.6	9.7	—	7.8	—
FP64 tensor core (TFLOPS)	51.2	19.5	—	—	—
Real LINPACK performance FP64 (TFLOPS)	25.71	13.58	—	6.11	—
FP32 (TFLOPS)	51.2	19.5	10.752	15.7	12
TF32 tensor core (TFLOPS)	378	156	74.8	—	—
BF16 (TFLOPS)	102.4	39	—	—	—
BF16 tensor core (TFLOPS)	756	312	149.7	—	—
FP16 (TFLOPS)	102.4	78	37.42	31.4	183.7
FP16 tensor core (TFLOPS)	756	312	149.7	125	—
FP8 tensor core (TFLOPS)	1513	—	—	—	—
INT32 (TOps)	25.6	19.5	5.376	15.7	—
INT8 tensor core (TOps)	1513	624	299.3	62	47
INT4 tensor core (TOps)	3026	1248	598.7	—	—
GPU memory (GB)	80	80	48	32	24
GPU memory bandwidth (GB/s)	2039	1555	696	900	346
FP32 cores (count)	14 592	6912	—	5020	—
FP64 cores (count)	7296	3456	—	2560	—
Tensor cores GPU (count)	456	432	336	640	—
Max Thermal Design Power (TDP)	350 W	400 W	300 W	300 W	250 W
Multi-Instance GPU (MIG)	7×10 GB, 4×20 GB, 2×40 GB	7×10 GB, 3×20 GB, 2×40 GB	—	—	—
Form factor	PCIe5	SXM4	PCIe4	SXM2	PCIe3
Interconnect	NVLink Bridge: 600 GB/s, PCIe Gen5: 128 GB/s	HGX-8xGPU NVLink: 600 GB/s, PCIe Gen4: 64 GB/s	PCIe Gen4: 64 GB/s	HGX-4xGPU NVLink: 300 GB/s, PCIe Gen3: 32 GB/s	PCIe Gen3: 32 GB/s

Figure 2 illustrates the architecture of type E nodes, which implement Nvidia’s DGX A100 reference design. In this configuration, all GPUs are mounted on a specialized Nvidia HGX board and interconnected via high-speed NVSwitch links. This architecture allows the GPUs to directly access InfiniBand adapters and NVMe storage through a PCIe switch, eliminating the need to transfer data via the node’s main memory.

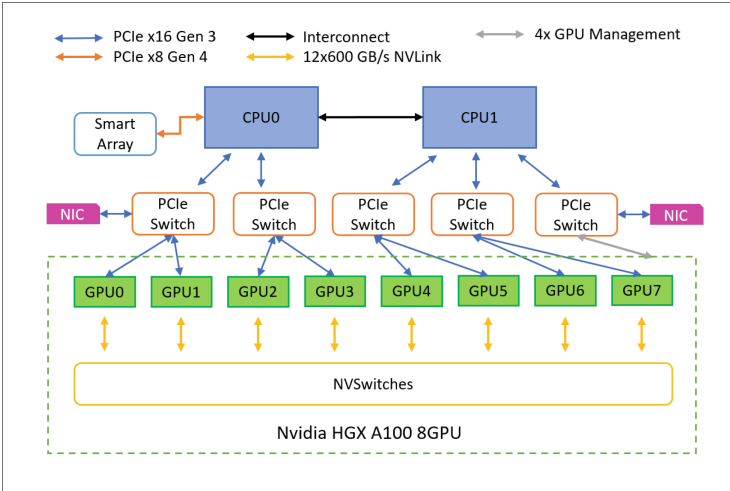


Fig. 2. GPU interaction for node type E.

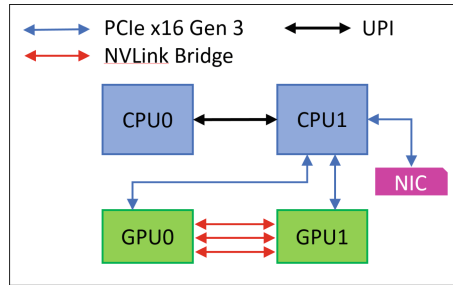


Fig. 3. GPU interaction for node type F.

Figure 3 presents the architecture of type F nodes, They are similar to type C nodes, but are faster, owing to direct CPU-to-GPU interconnectivity.

### 3 HPC TaskMaster: Task Efficiency Monitoring System

The HPC TaskMaster system [1], developed and deployed by the HSE Supercomputing Modeling Unit for the CHARISMa cluster, monitors task execution statistics and identifies inefficient resource usage, accounting for multiple simultaneous tasks running on each node. By analyzing jobs from the SLURM queue, the system detects suboptimal resource usage. Through detailed performance reports, users can fine-tune task parameters to optimize future resource allocation, enabling more efficient utilization of compute resources.

Figure 4 illustrates the architecture of the HPC TaskMaster system. Task submission is managed by the SLURM scheduler, while task data is stored in the

MariaDB database system. The Telegraf daemon collects CPU and GPU usage metrics, which are stored in the InfluxDB time-series database. Visualization is provided through Grafana, offering flexible graph creation and customization via API, display of both completed and ongoing tasks, and real-time progress monitoring. The system is built on Django, a Python-based web framework that provides robust tools for web application development.

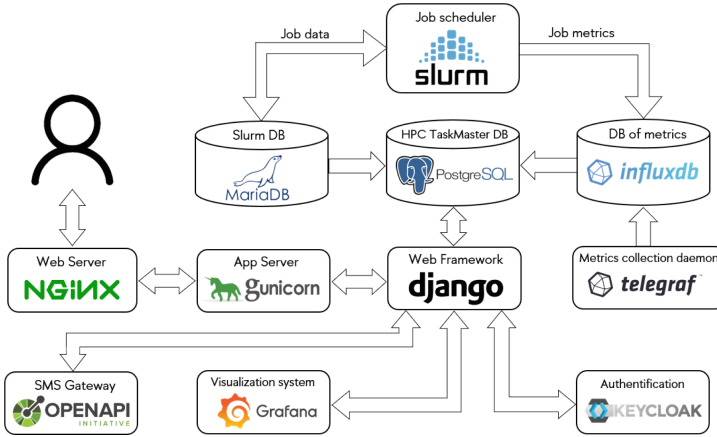


Fig. 4. HPC TaskMaster system architecture.

A typical HPC TaskMaster workflow involves:

1. Metric collection from compute nodes via Telegraf and storing retrieved data in InfluxDB.
2. PostgreSQL database updates through comparison with SLURM database records.
3. Continuous metric collection from InfluxDB during task execution.
4. Final metric aggregation upon task completion.
5. Comprehensive analysis of collected metrics to evaluate task efficiency.

Upon completion of the task analysis, the system generates the following efficiency optimization aids to help users improve the performance of their tasks:

- *Tag*. Binary label (0/1) identifying task properties or characteristics, enabling simplified analysis.
- *Indicator*. Normalized efficiency metric ranging from 0 to 1, inversely proportional to component utilization levels.
- *Inference*. Condition-based findings derived from specific tag/indicator combinations, assigned when all requirements in the respective set are met. Each inference has predefined requirements for indicator values and tag presence.

The tag-indicator-inference system serves two key purposes: making task performance statistics more accessible and suggesting specific efficiency improvements for future tasks. Additionally, the HPC TaskMaster system actively identifies inefficient tasks through its resource utilization analysis subsystem, which detects non-parallel computational tasks and tasks causing compute resource idling. Inefficiency determinations are based on low correlation metrics for resource usage and excessive resource idle times. The system implements a progressive response: it notifies users of suboptimal resource usage and automatically terminates tasks if no corrective actions are taken within one hour.

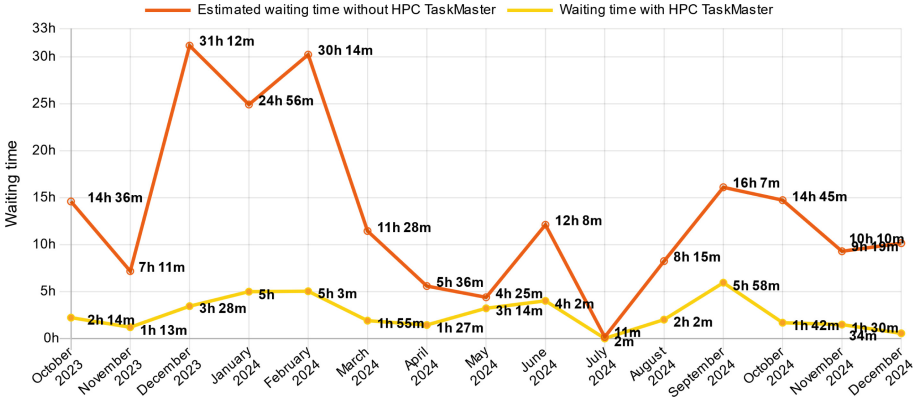


Fig. 5. Queue waiting time analysis with HPC TaskMaster. (Color figure online)

The average task queue waiting time serves as a key user satisfaction metric for the computing cluster. The average task queue waiting time—a key user satisfaction metric—has shown a fivefold reduction since deploying the inefficient-task termination system. Figure 5 displays the dynamics of average queue waiting times. The red line corresponds to theoretical, projected wait times without HPC TaskMaster, while the yellow line shows actual, observed wait times with the system active. Theoretical calculations account for all tasks terminated by the inefficiency detection system.

## 4 Monitoring Panel

Efficient monitoring of computing resources and task queues remains a core responsibility for supercomputer centers. Effective monitoring enables early detection or prediction of potential issues, timely corrective interventions, and minimized user disruptions. While universal monitoring systems exist (e.g., Nagios, Zabbix, Ganglia), most HPC clusters require customized solutions to address their unique operational requirements.

The HSE Supercomputing Modeling Unit monitoring approach extends beyond basic node status checks. It includes per-node real-time workload evaluation and queue participation tracking, as well as per-task resource utilization analysis. Unlike traditional HPC clusters with node-level allocation, our architecture supports concurrent multi-task execution per node and partial resource utilization per task. Dynamic performance optimization is essential. This unique environment required the development of a custom monitoring system with flexible architecture for continuous improvements. The concept behind this system is straightforward: aggregating and logically linking data from IPMI (hardware metrics), OS-level software counters, and the job queue system. This integration enables unified visibility into resource efficiency, node status, task performance metrics, and other critical parameters. A dedicated daemon operates on each compute node to continuously collect system metrics and IPMI data. This data is periodically transferred to the head server. The system's fault tolerance includes local data storage during network interruptions and automatic synchronization upon connection restoration.

The head server operates through coordinated daemon activities. A data collection daemon receives and logs raw metrics from compute nodes. The MSystem daemon aggregates data with job queue metrics, conducts analysis for visualization, and stores processed results in a PostgreSQL database.

Additionally, queue replication data from login servers ensures the system can validate and update data from compute nodes during synchronization delays. This validation enables accurate task evaluation by comparing requested and utilized resources and calculating overall resource efficiency.

The system's web interface (Fig. 6) provides comprehensive monitoring through:

- a. *Small pie charts* (left). Depict CPU/GPU resource allocation and node queue states.
- b. *Horizontal bar chart* (right). Displays real-time resource usage by active users, with interactive features including hover-to-highlight functionality for the central pie chart and click-to-navigate access to individual user profiles.
- c. *Large pie chart* (center). Provides a detailed overview of node status:
  - Inner ring: Node identifiers;
  - Middle ring: GPU status per node;
  - Outer ring: CPU core allocation.
- d. *Server real-time camera image* (bottom right). Provides real-time physical monitoring.

Figure 7 shows how data is designated in the pie chart. Divided into four parts, each sector displays: the node number (at the center of the chart), RAM usage, GPU usage (free/occupied), and (at the edge of the chart) CPU core allocation (free/occupied cores).

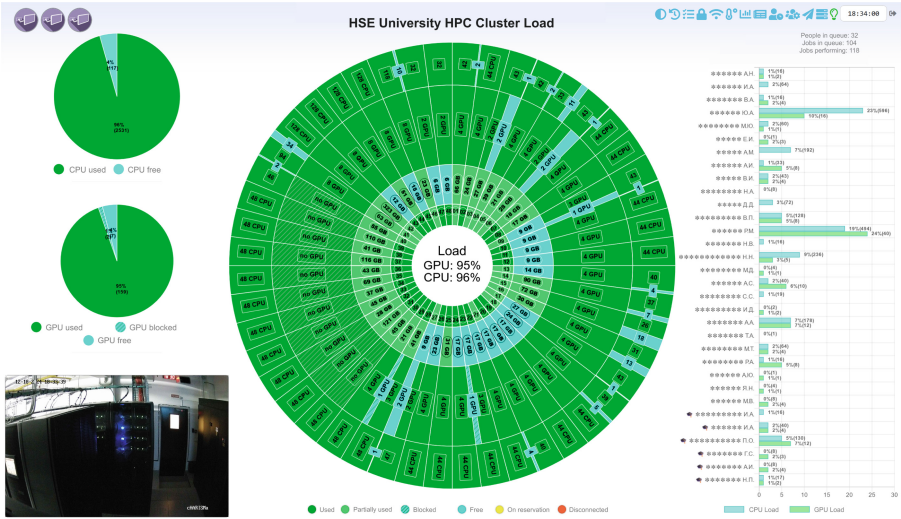


Fig. 6. View of the main monitoring panel.




Fig. 7. Data representation in the large pie chart.

The charts refresh in real time with maximum 10-second delays, allowing quick identification of queue optimization opportunities. Unlike typical monitoring systems that use multiple separate graphs requiring complex interpretation, our solution enables instant cluster status overview, detailed node inspection, and simultaneous task/user analysis.

The system maintains comprehensive historical records of cluster usage, enabling analysis of queue states and generation of load reports.

The system relies on various core indicators: CPU utilization, GPU utilization, and total system load. These metrics support granular reporting at multiple levels: faculty-wide trends, department-level analysis, and individual user statistics.

Another advantage of the system is the integration with both internal cluster services (such as user management system and email notifications) and external

systems (HSE university portal). This integration enables automated reporting (for instance, instant resource usage reports upon user/project registration).

## 5 Scientific Project Execution

The cHARISMa supercomputer serves as a critical resource for more than 64 departments at HSE University. Over 1500 users rely on it for both scientific research and educational purposes. A key operational feature is the concurrent allocation of node resources (multi-core/GPU sharing). These time and resource-sharing capabilities enable the execution of more than 580 thousand computational experiments annually. The cHARISMa supercomputer has processed over 2.4 million jobs to date.

The system supports cutting-edge research across diverse disciplines at HSE, including:

- machine learning [2,3];
- big-data analytics [4];
- complex systems modeling [5];
- neurocognitive studies [6,7];
- chemical and materials modeling [8–10];
- structural and applied linguistics [11];
- genetics and biology [12];
- particle, quantum, and high-energy physics [13,14];
- astrophysical simulations [15].

The cluster’s computational capabilities have enabled high-impact publications across diverse domains, including:

1. comparative analysis of novel physics simulation algorithms [16];
2. population genomics admixture proportion estimation [17];
3. algorithm development for angle distribution density in polyhedron random sections [18];
4. GPU-accelerated force calculation for microparticles in plasma flows [19];
5. parallel database system modeling [20];
6. finite-size analysis in neural network classification of critical phenomena [21];
7. neural network approaches for transient light curve approximations [22];
8. neural response analysis to naturalistic persuasive messages [23];
9. dyslexia identification in school pupils from eye-tracking and demographic data [24];
10. intrinsic dimensions of language fractal structures [25];
11. visual perception robustness analysis [26].

The breadth of research themes demonstrates HPC’s vital role in modern science. While advanced hardware is essential, a responsive and efficient support framework is equally crucial for a seamless user experience. Prompt resolution of inquiries and technical issues minimizes workflow disruptions and builds user trust in system reliability. The Supercomputer Modeling Unit has successfully addressed over 2350 user requests. This support infrastructure ensures researchers can focus on their work precisely when needed, reflecting our commitment to precision in problem-solving and timely intervention.

## References

1. Kostenetskiy, P., et al.: HPC TaskMaster - task efficiency monitoring system for the supercomputer center. In: *Parallel Computational Technologies*, pp. 17–29. Springer, Cham (2022). [https://doi.org/10.1007/978-3-031-11623-0\\_2](https://doi.org/10.1007/978-3-031-11623-0_2). ISBN: 978-3-031-11623-0
2. Koltcov, S., et al.: Topic models with elements of neural networks: investigation of stability, coherence, and determining the optimal number of topics. In: *PeerJ Comput. Sci.* **10**, e1758 (2024)
3. Dogonasheva, O., et al.: Multistability and evolution of chimera states in a network of type II MorrisLecar neurons with asymmetrical nonlocal inhibitory connections. *Chaos Interdisc. J. Nonlin. Sci.* **32**(10), 101101 (2022). <https://doi.org/10.1063/5.0117845>, [https://pubs.aip.org/aip/cha/article-pdf/doi/10.1063/5.0117845/16452021/101101\\_1\\_online.pdf](https://pubs.aip.org/aip/cha/article-pdf/doi/10.1063/5.0117845/16452021/101101_1_online.pdf). ISSN 1054-1500.
4. Nikitiuk, B.I., et al.: Pair entropy and universal viscosity scaling for molecular systems via molecular dynamics simulations. *J. Mol. Liq.* **368**, 120714 (2022). <https://doi.org/10.1016/j.molliq.2022.120714>, <https://www.sciencedirect.com/science/article/pii/S016773222202253X>. ISSN: 0167-7322.
5. Matveenko, S.I., et al.: Rotons and their damping in elongated dipolar Bose-Einstein condensates. *Phys. Rev. A* **106**, 013319 (2022). <https://doi.org/10.1103/PhysRevA.106.013319>, <https://link.aps.org/doi/10.1103/PhysRevA.106.013319>
6. Panidi, K., et al.: Posterior parietal cortex is causally involved in reward valuation but not in probability weighting during risky choice. *Cerebral Cortex* **34**(1), bhad446 (2023). <https://doi.org/10.1093/cercor/bhad446>, <https://academic.oup.com/cercor/article-pdf/34/1/bhad446/56132257/bhad446.pdf>. ISSN: 1460-2199.
7. Baranov, A., Kniazhevsky, V., Braslavski, P.: You told me that joke twice: a systematic investigation of transferability and robustness of humor detection models. In: Bouamor, H., Pino, J., Bali, K. (eds.) *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Singapore: Association for Computational Linguistics, pp. 13701–13715 (2023). <https://doi.org/10.18653/v1/2023.emnlp-main.845>, <https://aclanthology.org/2023.emnlp-main.845>
8. Bakulin, I.K., Kopanichuk, I.V., Kondratyuk, N.D.: Molecular-level insights to structure and hydrogen bonds network of 1,4-dioxane aqueous solution. *J. Mol. Liq.* **393**, 123523 (2024). <https://doi.org/10.1016/j.molliq.2023.123523>, <https://www.sciencedirect.com/science/article/pii/S01677322230023292>. ISSN: 0167-7322.
9. Khnkoian, G.V., Nikolaev, V.S., Stegailov, V.V.: Towards atomistic modelling of solid Pb-O formation and dissolution in liquid lead coolant: Interatomic potential development. *J. Nucl. Mater.* **594**, 155016 (2024). <https://doi.org/10.1016/j.jnucmat.2024.155016>, <https://www.sciencedirect.com/science/article/pii/S0022311524001193>. ISSN: 0022-3115.
10. Kondratyuk, N., et al.: First-principles calculations of the viscosity in multi-component metallic melts: Al-Cu-Ni as a test case. *J. Mol. Liq.* **380**, 121751 (2023). <https://doi.org/10.1016/j.molliq.2023.121751>, <https://www.sciencedirect.com/science/article/pii/S01677322230005548>. ISSN: 0167-7322
11. Klyuchnikov, N., et al.: NAS-Bench-NLP: Neural Architecture Search Benchmark for Natural Language Processing. *IEEE Access* **10**, 4573645747 (2022). <https://doi.org/10.1109/ACCESS.2022.3169897>
12. Liang, M., et al.: Understanding admixture fractions: theory and estimation of gene-ow. *J. Math. Biol.* **89** (2024). <https://doi.org/10.1007/s00285-024-02146-0>.

13. Bronin, S.Y., et al.: Electronic resonances in expanding non-neutral ultracold plasma. *Phys. Plasmas* **31**(3), 033507 (2024). <https://doi.org/10.1063/5.0194420>, <https://pubs.aip.org/aip/pop/article-pdf/doi/10.1063/5.0194420/19850158/033507.1.5.0194420.pdf>. ISSN: 1070-664X
14. Bahovadinov, M.S., et al.: Enhanced Drude weight in a one-dimensional system of fermions with pair-hopping events. In: *Phys. Rev.* **109**, 195159 (2024). <https://doi.org/10.1103/PhysRevB.109.195159>, <https://link.aps.org/doi/10.1103/PhysRevB.109.195159>
15. Korolkov, S., Izmodenov, V.: The global structure of astro- spheres: effect of Knudsen number. *Publ. Astron. Soc. Austral.* **41**, e074 (2024). <https://doi.org/10.1017/pasa.2024.44>.
16. Mozolenko, V., Fadeeva, M., Shchur, L.: Comparison of the microcanonical population annealing algorithm with the Wang-Landau algorithm. *Phys. Rev. E* **110**, 045301 (2024). <https://doi.org/10.1103/PhysRevE.110.045301>, <https://link.aps.org/doi/10.1103/PhysRevE.110.045301>
17. Liang, M. et al.: Understanding admixture fractions: theory and estimation of gene-ow. *J. Math. Biol.* **89**(5) (2024). <https://doi.org/10.1007/s00285-024-02146-0>, <https://link.springer.com/10.1007/s00285-024-02146-0>
18. Konstantinova, E.P., Shchur, L.N.: Algorithm for density of angle distribution in random sections of polyhedron. *Expert Syst. Appl.* **235**, 121195 (2024). <https://doi.org/10.1016/j.eswa.2023.121195>, <https://www.sciencedirect.com/science/article/pii/S0957417423016974>. ISSN: 0957-4174
19. Kolotinskii, D., Timofeev, A.: OpenDust: a fast GPU-accelerated code for the calculation of forces acting on microparticles in a plasma ow. *Comput. Phys. Commun.* **288**, 108746 (2023). <https://doi.org/10.1016/j.cpc.2023.108746>, <https://www.sciencedirect.com/science/article/pii/S0010465523000917>. ISSN: 0010-4655
20. Rekachinsky, A.I., Chulkevich, R., Kostenetskiy, P.: Modeling parallel processing of databases on the central processor Intel Xeon Phi KNL. In: (May 2018), pp. 1605–1610. <https://doi.org/10.23919/MIPRO.2018.8400288>
21. Chertenkov, V., Burovski, E., Shchur, L.: Finite-size analysis in neural network classification of critical phenomena. *Phys. Rev. E* **108**, L032102 (2023). <https://doi.org/10.1103/PhysRevE.108.L032102>, <https://link.aps.org/doi/10.1103/PhysRevE.108.L032102>
22. Demianenko, M., et al.: Understanding of the properties of neural network approaches for transient light curve approximations. *A&A* **677**, A16 (2023). <https://doi.org/10.1051/0004-6361/202245189>
23. Ntoumanis, I., et al.: Deciphering the neural responses to a naturalistic persuasive message. *Proc. Nat. Acad. Sci.* **121**(43) e2401317121 (2024). <https://doi.org/10.1073/pnas.2401317121>, <https://www.pnas.org/doi/pdf/10.1073/pnas.2401317121>
24. Shalileh, S., et al.: Identifying dyslexia in school pupils from eye movement and demographic data using artificial intelligence. *PLOS ONE* **18**(11), 1–26 (2023). <https://doi.org/10.1371/journal.pone.0292047>
25. Gromov, V.A., Borodin, N.S., Yerbolova, A.S.: A language and its dimensions: intrinsic dimensions of language fractal structures. *Complexity* 2024(1), 8863360 (2024). <https://doi.org/10.1155/2024/8863360>
26. Shirnin, A., et al.: Analyzing the robustness of vision & language models. *IEEE/ACM Trans. Audio Speech Lang. Process.* **PP**, 1–15 (2024). <https://doi.org/10.1109/TASLP.2024.3399061>