



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking



An Roinn Breisnídeachais agus Ardoidenachais,
Taighde, Nuálaíochta agus Eolaíochta
Department of Further and Higher Education,
Research, Innovation and Science





ÉIRE2026.EU
IRELAND2026.EU

EuroHPC JU Access Opportunities


Klara Meštrović, Laura Martínez Calvo, Annika Kossack
Peer-Review Sector

EuroHPC JU Access Calls



EuroHPC JU Access Calls

HPC APPLICATIONS




Preparatory and
production access
calls types

AI APPLICATIONS



Scientific and industry
call types

QUANTUM APPLICATIONS



Pilot stage



Benchmark Access

For code and scalability tests

Access duration 2-3 months

Allocation process duration – 2 weeks after the cut-off date

Development Access

For code and algorithm development and optimization

Access duration 6-12 months

Allocation process duration – 2 weeks after the cut-off date

Regular Access

For research applications requesting large allocations

Access duration 12 months


Peer-Review process duration – 4 months after the cut-off date

Extreme Scale Access

For high-impact, high-gain research applications requesting extremely large allocations

Access duration 12 months

Peer-Review process duration – 6 months after the cut-off date



Playground Access

For new users or those with smaller computational needs for performing AI computations

Access duration 1, 2 or 3 months

Allocation process duration – 2 working days

Large Scale Access

For industrial AI applications with high impact needing large allocations

Access duration 3, 6 or 12 months

Peer-Review process duration – 2 weeks after the cut-off date

For industrial users familiar with HPC environments requiring fast access

Access duration 1, 2 or 3 months

Allocation process duration – 4 working days

Fast Lane Access

For AI applications performing scientific research activities

Access duration 6 months

Allocation process duration – 1 months after the cut-off date




AI for Science and Collaborative Projects Access


Impact of EuroHPC JU access opportunities

ACCESS CALL	PROPOSALS SUBMITTED	PROPOSALS AWARDED	NODE HOURS AWARDED
Extreme Scale Access (Dec 2022 - Apr 2026)	233 ↑143% vs last cut-off	162	100,856,071
Regular Access (Dec 2021 - Mar 2026)	625 ↑200% vs last cut-off	419	51,715,942
Development Access (Jan 2022 - Sep 2026)	2,461 ↑350% vs Jan 2026	1,896	12,135,863
Benchmark Access (Jan 2022 - Sep 2026)	1,266 ↑1044% vs Jan 2026	1,033	4,457,467
AI and Data-Intensive Applications Access (Apr 2024 - Apr 2025)	123	79	3,266,600
AI for Science and Collaborative EU Projects Access (Jun 2025 - Sep 2026)	231 ↑230% vs 2025	105	6,370,042
Playground Access (Apr 2025 - Sep 2026)	621 ↑423% vs 2025	467	469,375
Fast Lane Access (Apr 2025 - Sep 2026)	553 ↑345% vs 2025	362	2,180,686
Large Scale Access (Apr 2025 - Sep 2026)	155 ↑151% vs 2025	69	9,432,405
TOTAL	6,268	4,592	190,884,451

Awarded resources distribution across access calls



Awarded proposals distribution across access calls



Impact of EuroHPC JU access opportunities

How to create impact?



Cite the EuroHPC JU resources that have allowed you to carry out your work. And let us know when you publish! That way we can also promote your work.



Make publicly available your models, datasets, etc. In public repositories according to FAIR.



Be on our radar?? We publish success stories.



Please submit your **Final Reports**.



Participate in our events.


User survey results

Participants demographics


The survey was launched on **31 July 2026** and was opened until **4 Sept 2026**

2.399 invited users – 693 (29%) participated in the survey


Participants research domains distribution



Participants organization type distribution




Participants level of HPC experience




User survey results

Participants demographics


Access calls participation



Application types



Resource types



HPC applications: 42% Development Access, 33% Regular Access, 30 Benchmark Access, 13% Extreme Scale Access

74%

AI applications: 12% Fast Lane Access, 11% Playground Access, 5% AI for Science, 4% Large Scale

26%

User survey results

Overall experience


QUESTION	RATING
Please rate the accessibility and clarity of the information provided about the EuroHPC JU Access Calls and access possibilities in general within the EuroHPC communication channels (website, social media, etc.)	4.25
Please rate your experience (if any) in communicating with the EuroHPC JU and especially the Peer-Review Sector regarding access to HPC resources	4.27
Please rate your experience in using the EuroHPC JU Peer-Review Platform (submission platform)	4.2
Please rate your experience with the Hosting Entity (supercomputing centre) and the corresponding user support	4.27
Please rate your experience in using the EuroHPC JU HPC system(s)	4.24

Average rating


4.25



Access calls experience – from application to results



Access calls experience combined averages



User survey results

Qualitative feedback

New applications:

74% positive, 19% passive, 7% negative

Recommend EuroHPC:

72% positive, 19% passive, 9% negative

Qualitative feedback written by 165 participants (24%)



Login via EFP



Thank you message



SSH certificates

Downtime & outage issues



Rejections based on exhausted resources

Public statistics




Calls planning in advance

Complicated cut-off distribution across access modes

More frequent evaluations

For more information, feel free to visit our website
and social media or contact us at:



Thank you !

[/eurohpc-ju.europa.eu](https://eurohpc-ju.europa.eu)[/EuroHPC_JU](https://twitter.com/EuroHPC_JU)[/eurohpc-ju](https://www.linkedin.com/company/eurohpc-ju)[/eurohpc-ju](https://www.youtube.com/channel/UCaifCmC1b_AghvndGmPjL1g)


access@eurohpc-ju.europa.eu
aif.access@eurohpc-ju.europa.eu


New frontiers of Climate science enabled by EuroHPC

Daniel Klocke

MAX-PLANCK-INSTITUT
FÜR METEOROLOGIE








Destination Earth

- Climate DT delivers world's first multi-decadal, multi-model kilometre-scale scenarios.
- Extremes DT routinely detects and delivers global and regional on-demand extreme event simulations.
- Deliver the world's first AI-based Earth system model
- The Digital Twin Engine integrates simulations, AI-ready data production, training and inference.
- DestinE strengthens the maturing EuroHPC ecosystem.

**120+ institutions 26 countries 22 PB+
climate data portfolio**

World-leading performance & applications

Mean of annual maximum 1-day precipitation (RX1day) from 1990 to 2005



The effect of substantially enhanced resolution in the **global Climate DT**, versus **traditional climate models** and **regional climate models**

Additional challenge

When adding all necessary complexity, we do not want to compromise simulation throughput.

Temporal compression τ (simulated time per real time - non dimensional)

$\tau = 30$: easy to study the diurnal cycle, practical to study the annual cycle





$\tau = 300$: easy to study interannual variability and practical to perform multi-decadal simulations

Efficiency across EuroHPC platforms

Finalist 2025



ACM Gordon Bell Prize
for Climate Modelling




At the heart of the challenge

The ambition to incorporate all scales and processes relevant to physically couple the cycles of **Energy**, **Water** and Carbon

Energy and **Water** cycles: needs resolving the gravest modes of convective instability (weather)

Coupling those to the carbon cycle: adds biochemical and dynamical processes in the ocean and a lot of processes in the land biosphere

Earth System complexity



Non-hydrostatic dynamics on an icosahedron based grid

The Earth-System simulated at 1 km resolution

Degrees of Freedom

1 000 000 000 000

Two of the most powerful compute systems in Europe - both with GH200 chips

JUPITER at FZJ

24000 GH200 chips

InfiniBand NDR200


1 Exaflop (#4 in top500)

ALPS at CSCS

11000 GH200 chips


Slingshot-11

0.435 Exaflop (#8 in top500)




Earth System's Complexity and System Parallelism

Alps (CSCS) and JUPITER (JSC) supercomputers are based on NVIDIA Grace Hopper superchips




NVIDIA GH200 Superchip




The ocean and its biogeochemistry comes essentially for “free”

Adapted from Klocke et al., 2025


Scaling of the ICON Earth System Model on JUPITER



Scaling of the ICON Earth System Model on JUPITER



Scaling of the ICON Earth System Model on JUPITER




Phytoplankton

Nitrate

CO₂ flux

Zooplankton grazing

Zooplankton mortality



Conclusions and Implications

- km-scale Earth System models are realisable on EuroHPC infrastructure at useful time-compressions and open up new and exciting scientific frontiers
- Modern supercomputers are heterogeneous. Climate models can and should take advantage of that
- With the performance demonstrated, one could run two scenarios for the warming over the next 30 years (3 ensemble members each) in less than six months
- Potential to provide full global Earth system information on local scales about the implications of future warming for people and eco-systems
- Generate a physical based data set of the flow of **energy**, **water** and carbon through the Earth system.... Information which would otherwise not exist





EuroHPC
User Days

Dublin Ireland
22→24
September 2026



Proton structure with large-scale simulations

Constantia Alexandrou

University of Cyprus and The Cyprus Institute



AQTIVATE

European Joint Doctorate, grant agreement No. 101072344



EXCELLENCE/0421/0043

Empowering world-class research in Cyprus

- ✳ 2004: Cyprus becomes full member of the EU
- ✳ JSC opens its computational resources to new member countries. —> **for the first time we could apply for access on supercomputers** - the JUMP 9 Tflop/s machine, the most powerful computer in Europe
- ✳ 2008: The first supercomputer system of CPUs and GPUs of 30 Tflop/s becomes available for the Cypriot scientific community —> **allowed training and algorithm development locally**
- ✳ 2008: Cyprus became member of PRACE
- ✳ 2019: Cyprus joined the EuroHPC Joint Undertaking



Access to PRACE and EuroHPC infrastructure was and is a critical factor for our visibility on the world stage


Impact on talent building

- ✳ Participation in PRACE IP-projects included training and algorithmic development and interconnected Cyprus with advanced European know-how and opened funding opportunities
 - > One of first groups to have GPU-enabled codes that improve our ability to attract local and international talent
- ✳ This path continues with access to EuroHPC resources

Some of the young researchers in lattice QCD funded by IP projects who had access to European Infrastructure




G. Koutsou
Assoc. Prof. CyI



A. Strelchenko
Fermi Lab



K. Hadjiyiannakou
Financial sector,
Cyprus




A. Vaquero Assist.
Prof. Zaragoza



A. Abdel-Rehim
Medical Physicist ,USA



M. Constantinou
Assoc. Prof. Temple
University, USA



Ch. Kallidonis
GPU programmer
Accuracy, USA



J. Finkenrath ERC-
holder, Assist. Prof.
Wuppertal



S. Bacchio
MSC fellow,
Researcher, CyI



F. Pittler
Researcher, CyI



A. Athenodorou
Researcher, CyI



S. Spanoudes
Researcher, UCY

Quantum Chromodynamics (QCD)

$$\mathcal{L}_{QCD} = -\frac{1}{4}F_{\mu\nu}^a F^{a\mu\nu} + \sum_{f=u,d,s,c,b,t} \bar{\psi}_f (i\gamma^\mu D_\mu - m_f) \psi_f$$


Fritzsch, Gell-Mann and Leutwyler, Phys. Lett. 47B (1973) 365

✱ Unique properties:

- ★ Confinement
- ★ Asymptotic freedom
- ★ Mass generation via interaction

QED

Quantum theory of the **electromagnetic** force mediated by exchange of photons



Hydrogen atom


$$m_{\text{Hydrogen}} = \underbrace{0.51 \text{ MeV}}_{m_{e^-}} + \underbrace{938.29 \text{ MeV}}_{m_{p^+}} - \underbrace{13.6 \text{ eV}}_{E_{\text{binding}}}$$

A. Stodolna et al., PRL 110 (2013) 213001

QCD

Quantum theory of the **strong** force mediated by exchange of gluons

Proton



Artist impression

$$m_p = \underbrace{2.3 \text{ MeV}}_{2 \times m_u} + \underbrace{4.7 \text{ MeV}}_{m_d} + \underbrace{929 \text{ MeV}}_{E_{\text{binding}}}$$

99% of proton mass from interaction!

Lattice QCD provides an *ab initio* method to study a wide class of strong interaction phenomena


✱ Lattice QCD uses directly $\mathcal{L}_{QCD}$ or the action $S_{QCD} = \int d^4x \mathcal{L}_{QCD}$

$$\langle \mathcal{O} \rangle = \frac{1}{Z} \int \mathcal{D}[U] \mathcal{O}(D^{-1}[U], U) \left(\prod_{f=u,d,s,c} \text{Det}(D_f[U]) \right) e^{-S_{QCD}[U]}$$


Workflow in lattice QCD

$$U \in \frac{1}{Z} \int \mathcal{D}[U] \left(\prod_{f=u,d,s,c} \text{Det}(D_f[U]) \right) e^{-S_{\text{QCD}}[U]}$$


$$\langle \mathcal{O} \rangle = \frac{1}{N} \sum_i^N \mathcal{O}(D^{-1}[U_i], U_i)$$



Quark & gluon propagators




contractions



State-of-the-art gauge ensembles for hadron structure

$$U \in \frac{1}{Z} \int \mathcal{D}[U] \left(\prod_{f=u,d,s,c} \text{Det}(D_f[U]) \right) e^{-S_{\text{QCD}}[U]}$$

$$\text{Cost} \propto \left(\frac{L}{a}\right)^4 \left(\frac{1}{a^2}\right)^\delta \left(\frac{1}{m_\pi^2}\right)^\epsilon \quad \delta, \epsilon \in (0, 1)$$



CPU machines: SuperMUC and Hawk at HLRS
Frontera at TACC, US

Wilson-type



Staggered



Domain wall



ETMC ensembles generated **over a 10 years period** at physical pion mass

✳ Resources used for the production of ETMC ensembles: 16.14 Mi. CPU node-h and 1.76 Mi. GPU node-h

✳ Ensembles analysed for the results of this talk: 12.3 Mi. CPU node-h & 1 Mi. GPU node-h

$N_f = 2 + 1 + 1$, Twisted – mass fermions : $64^3 \times 128, a = 0.080 \text{ fm}, m_\pi = 140 \text{ MeV}$

B64

$80^3 \times 160, a = 0.069 \text{ fm}, m_\pi = 137 \text{ MeV}$

C80

$96^3 \times 192, a = 0.057 \text{ fm}, m_\pi = 141 \text{ MeV}$

D96

$112^2 \times 224, a = 0.049 \text{ fm}, m_\pi = 136 \text{ MeV}$


E112

GPU machines: Meluxina, LUMI-G & Jupiter




Proton spin puzzle

✳️ What is the proton spin puzzle? In 1987, high-energy physics experiments at CERN conducted by the European Muon Collaboration showed that quarks contribute only about half of the proton spin





✳️ With access to Tier-0 resources by JSC, **CSCS**, and CINECA allowed us to do world-class research

✳️ To address this puzzle one needs: 1) gauge ensembles and 2) matrix elements of the energy and momentum tensor of QCD for quarks and gluons



Used the B64 ensemble $64^3 \times 128$ at physical pion mass


Understanding the proton spin




Juelich Supercomputing Center
65 Million core-h



Piz Daint, CERN
200 Million core-h




Fermi, CINECA
37 Million core-h



Phys. Rev. Lett. 119 (2017) 14, 142002, 1706.02973 [hep-lat]

$$J_{q,g} = \frac{1}{2} [A_{20}^{q,g}(0) + B_{20}^{q,g}(0)] \quad \sum_q \left[\frac{1}{2} \Delta \Sigma_q + L_q \right] + J_g = \frac{1}{2}, \quad \sum_q \langle x \rangle_q + \langle x \rangle_g = 1$$





—> Headline: Researchers make progress toward solving the proton spin puzzle


https://phys.org/news/2017-10-proton-puzzle.html#google_vignette

Resolution of the proton spin

✱ Next step is to recover the continuous limit. We used our 4 ensembles with volume ~ 5.5 fm



Confirms what the European Muon Collaboration measured



Additional contribution to the proton spin

SuperMUC: 377 Mi. core-h
JUWELS: 23 Mi. core-h
Piz Daint: 1.7 Mi. node-h



0.5 Mi. node-h



0.8 Mi. node-h
61 Mi. core-h



1.75 Mi. node-h



0.6 Mi. node-h

Submitted to *Phys. Rev. Lett.* [2607.21230](#) [hep-lat]


Higher Mellin moments and 3D structure of the proton

✳ Higher moments allow the reconstruction of the function: $\int_{-1}^1 dx x^{n-1} H(x, 0, 0) = \langle x^{n-1} \rangle$

✳ The standard approach is to use a parametrization and fits to the moments

✳ Going beyond second moments to higher moment requires more computational resources

✳ Generalize to 3D functions : $\int_{-1}^1 dx x^{n-1} H(x, \xi, q^2) = \sum_{i=0}^{n-1} A_{n,i}(q^2) (-2\xi)^i + C_{n0}(q^2) (-2\xi)^n \Big|_{n \text{ even}}$



Jupiter Booster JSC
115 Million core-h



MN5, BSC
46 Million core-h


LUMI, CSC
58 Million core-h



Leonardo, Booster, CINECA
86 Million core-h


✳ Via extreme scale calls we had access to LUMI, Leonardo, Marenostrum and JUPITER

Higher moments of the proton



- ✱ Computed up to the fourth moment and reconstruct the unpolarized PDF using one of our ensembles



$$\langle x^{n-1} \rangle = \int_0^1 dx x^{n-1} N x^\alpha (1-x)^\beta$$





Reconstructed the proton unpolarized PDF

Submitted to Phys. Rev. D [2605.29998](https://arxiv.org/abs/2605.29998) [hep-lat]

Direct computation of GPDs



$$h_\Gamma(z, \tilde{\xi}, Q^2, P_3) = \langle N(P_3 \hat{e}_z + \vec{Q}/2) | \bar{\psi}(z) \Gamma W(z, 0) \psi(0) | N(P_3 \hat{e}_z - \vec{Q}/2) \rangle$$




The intrinsic transverse space distributions $\rho_X(x, \mathbf{b})$ of up and down quarks in a proton polarized in the X direction at $\mu = 2 \text{ GeV}$

Lattice Field Theory community

✱ The lattice community has been and continues to be at the frontier of HPC and an early adopter of HPC contributing also to computer co-development

Prominent examples are:

- APE machines (started in the late 80s)
- QCDOC 2004 that served as a prototype for the Blue-Gene machines of IBM
- QPACE2 2014 top green computer



APE




QPACE2




QCDOC

Regular Access call (Dec 2021- Nov 2023) -
Research domains distribution



Extreme Scale Access (Dec 2022-Oct 2023) -
Research domains distribution -
awarded projects



Courtesy L. Del Debbio

EuroLFT: An association of Lattice Field Theory researchers in Europe



<https://eurolft.github.io/index>

- ✳ In 2020: the European lattice community initiated the formation of a representative body, EuroLFT **to act as a point of contact for the community**
- ✳ In 2025: EuroLFT was launched
- ✳ Members include all lattice field theory researchers in the geographical area of Europe
- ✳ Lattice Field Theory researchers are playing pivotal role in EuroHPC advisory bodies and access committees and PRACE

See A. Rago's presentation

See H. Wittig's presentation on the latest PRACE Scientific and Innovation case for HPC in Europe

<https://prace-ri.eu/scientific-case/the-scientific-and-innovation-case-for-computing-in-europe-2026-2034/>





EuroHPC
User Days

Dublin Ireland
22→24
September 2026



Conclusions

- * Pan-European infrastructure is crucial for empowering talent across Europe
- * Cutting-edge infrastructure enables world-class research and attracts talent internationally
—> **access to such infrastructure and inclusiveness in EU-wide activities have been and are indispensable to be competitive in science and innovation**
- * Fields like lattice field theory attract top talent benefiting other fields and industry as researchers progress in their scientific careers
- * Lattice field theory needs continuous access to state-of-the art computational resources and computational savvy people in order to remain competitive on the world stage
—> **true for other scientific communities in Europe, see e.g. PRACE Scientific and Innovation case 2026-2032**



European Joint Doctorate, grant agreement No. 101072344



EXCELLENCE/0421/0043



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking



ÉIRE2026.EU
IRELAND2026.EU

Massive Black Hole Seeding in a JWST Era

John Regan – Maynooth University - Ireland

Regular Access User: Meluxina & Karolina CPU Systems



Overall Science Goals

- Trying to model structure formation in the very early Universe (first 500 Myrs after the Big Bang)
- First stars in the Universe
- First black holes in the Universe
- First galaxies in the Universe



Why now?

- JWST
- Electrified the field
- Insights into the early Universe
- Opportunities to test our theories and models against real data



nature

View all journals

[Explore content](#) ▾ [About the journal](#) ▾ [Publish with us](#) ▾ [Subscribe](#)[nature](#) > [news](#) > article

NEWS | 03 August 2023

'It's a dream': JWST spies more black holes than astronomers predicted

The James Webb Space Telescope's observations could help to answer questions about how the celestial objects formed early in the Universe.

By [Alexandra Witze](#)[Save article](#)

Last August, astrophysicist Dale Kocevski posted a paper to the arXiv preprint server with

JWST has shown us a brighter,
more chaotic and vibrant
Universe than we thought
possible


COSMIC EVOLU
BLACK H

NIRCam Imaging




Record-breaking distant galaxy

In 2022, JWST discovered four galaxies which had a redshift z of more than 10. They were designated JADES-GS-z10-0, JADES-GS-z11-0, JADES-GS-z12-0 and JADES-GS-z13-0. These galaxies are so far away that the light they emit has already completely shifted to the infrared part of the electromagnetic spectrum. JWST is the only telescope in the world capable of detecting it.

JADES/GOODS-S
JWST/NIRCam

Images of four galaxies found by JWST. Source: Emma Curtis-Lake et al. / Nature Astronomy, 2023


JWST so far: Little Red Dots, Overmassive Black Hole and super- bright galaxies



When the data from a variety of “little red dot” galaxies are broken up into their stellar mass component versus the component arising from an active supermassive black hole, the mass ratios of the galaxy’s total stellar mass compared with the supermassive black hole’s mass can be determined. Many, and perhaps even most, of these black holes are found to be significantly overmassive: at much more than 0.1% of the mass of the stellar component. (Credit: D. Kocevski et al., *Astrophysical Journal Letters* accepted/arXiv:2404.03576, 2025)

The goal of our **EuroHPC** project is to
RUN STATE-OF-THE-ART SIMULATIONS
MODELLING THE FIRST STAGES OF
GALAXY, STAR AND BLACK HOLE FORMATION
IN THE EARLY UNIVERSE

1 COMPARE OUR MODELS
AGAINST THE EMERGING
DATA FROM **JWST**



OUR SIMULATIONS

2 ALLOWING US TO BETTER
INTERPRET THE JWST DATA



JWST OBSERVATIONS




3 GIVING US A BETTER UNDERSTANDING
OF THE EARLY DAYS OF OUR UNIVERSE
AND OF HOW GALAXIES LIKE OUR OWN
MILKY WAY GALAXY EMERGED.



Tech Talk Lite


Software: Arepo Hydrodynamic Code

- Massively parallel code parallelised using MPI
- Moving mesh architecture
- Particles sample the dark matter flow
- Moving mesh tracks the hydro component allowing increased resolution in high density regions
- Subgrid prescriptions required to model processes that occur below the grid scale (e.g. star formation, black hole formation, radiation transport)





SEEDZ



- Our model for the early Universe is called SEEDZ
- SEEDZ uses HPC resources to model the formation of structure (stars, black holes and galaxies in the early Universe)
- We must then compare those virtual galaxies to the Universe we observe



Evolution of a small portion of a virtual Universe



Evolution of a small portion of a virtual Universe



Using SEEDZ we can
successfully reproduce
galaxies that look
remarkably similar to
"JWST" galaxies

N1-NoFB | Galaxy ID: 0

100pc

Galaxy Properties

Redshift	10.00
Halo Mass	$2.3 \times 10^{10} M_{\odot}$
Gas Mass	$1.8 \times 10^8 M_{\odot}$
Stellar Mass	$1.5 \times 10^8 M_{\odot}$
BH Mass	$2.6 \times 10^8 M_{\odot}$
HalfMass Radius	0.6 kpc

GasCold ($T < 10^4$) : 85.0%Hot ($T > 10^4$) : 15.0%**Metallicity**Mean Z : 4.0×10^{-3} Max Z : 1.5×10^{-1} **Stars**

Pop III : 0

Pop II (Young) : 1022

Pop II (Old) : 20417

Black Holes


Number : 918

Largest : $2.6 \times 10^8 M_{\odot}$ Gas Density (cm^{-3})

Temperature (K)


Metallicity ($Z_{\odot}$)

Radial Velocity (km/s)



And our virtual galaxy
population matches the
populations of galaxies
being detected by JWST

...and because this
is a simulation we
can study the virtual
assembly history
too!



The future....





FIG. 3.— Shown above are projected dark-matter density fields calculated using a cloud-in-cells method. The top row of images, from left to right, show density fields for a $z = 0$ snapshot from low and high resolution simulations of a $100 \text{ h}^{-1} \text{ Mpc}$ box, together with an enhanced snapshot respectively. The enhanced snapshot was created using the low resolution snapshot as input to the model. The bottom row of images show zoom-ins of the $25 \text{ h}^{-1} \text{ Mpc}$ regions highlighted in the top row by enclosing boxes. *Takeaway:* We see, even on scales where the resolution of the density field is apparent, it is difficult to distinguish between a high resolution density field and one created using the SR technique.

Impact of EuroHPC on my group's work

- EuroHPC awards: DECI 2019, Regular Access 2022, 2023 (Karolina), Regular Access 2025, 2026 (Meluxina), Developer access time on Marenostrum (2026) for ML models
- ~ 17 papers directly from EuroHPC compute time in last ~ 4 years
- Via EuroHPC our group has delivered the state-of-the-art SEEDZ simulations with the goal of complementing current JWST observations.





EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking



Voice for Purpose phase 2

Accelerated, multilingual end to end speech
production for non-verbal people

Fabio Minazzi – Translated Srl, Italy

**Voice
for Purpose.**

Medical need

- **200 ML** people in the world with **speech impairment conditions**
Motor neuron conditions (ALS, SMA), Larynx Cancer, Stroke, Cerebral Palsy, Autism.
- **500.000** people in the world with **ALS***
- **7,5** people with voice impairment in the US**

[*https://pubmed.ncbi.nlm.nih.gov/27185810/](https://pubmed.ncbi.nlm.nih.gov/27185810/)

[**https://dysphonia.org/voice-conditions/overview-of-vocal-disorders/](https://dysphonia.org/voice-conditions/overview-of-vocal-disorders/)



Voice for Purpose.




A stage set featuring a central projection screen displaying the text "There is nothing more precious than our voice." The screen is flanked by two identical side panels, each containing a lamp with a patterned shade and a black leather sofa. In the foreground, a vintage-style microphone on a stand is positioned on a light-colored wooden floor. The background wall has a dark, textured, arched pattern.

There is nothing more
precious than our voice.

Phase 1: Donate voices to those who cannot speak

- **Main project features**
 - User profiling and medical assistance
 - Voice Donation, Voice Self-Donation
 - Expressive AI voice model creation
 - Real time delivery of expressive AI speech to users
- **Innovations** (www.voiceforpurpose.com)
 - **Voice donation and self-donation:** social platform to support people in need of voice
 - **Cloud-based:** infrastructure for delivery to any connected device
 - **Expressive AI voices:** high quality, real time expressive voice rendering
- **Large social footprint:** 6000+ donors, 200+ users



Main issues to be overcome

- **Input speed**
 - avg 8 words/min (vs 40 w/min typing, 150 w/min normal speech)
- **Multilingual performance**
 - Address all European citizens
- **Sample size to create faithful voice models**
 - Reduce Donors' recordings from 1h to 5-10 minutes

Phase 2: Improve input speed, simplify voice model creation across the 24 European languages

Research Problems:

- Improve input speed/ reduce fatigue for people using AAC with eye-gaze
- Support 24 European languages
- Reduce the dataset recording for personal voices
- Improve models prosody
- Improve multilingual performance in 24 European languages
- Reduce number of voice models to maintain



=>

LLM track



=>

Audio track



Goal: Build an End-to-End text input-audio output system

Research team

- Fabio Minazzi, Project coordinator, Translated S.R.L.
- Disheng Qiu, PI, Translated S.R.L.
- Andrea De Angelis, Senior AI Engineer, Translated S.R.L.
- Lorenzo Pratesi, Data Scientist, Translated S.R.L.
- Thomas Passera, Junior Data Scientist, Translated S.R.L.
- Matteo Lippi, Junior Data Scientist, iFAB
- Jonas Mack, AI Engineer, Translated S.R.L. (while contributing to the project, now Smallest AI)
- Nicola Sobieski, Audiovisual Manager, Translated S.R.L.
- Constantin Lipcaneanu, Senior SW Developer Translated S.R.L.
- Maurizio Davoli, Data Scientist, Translated S.R.L.

End-to-end pipeline




Figure 1: Intended end-to-end AAC pipeline: the user enters shorthand, the LLM proposes top- k context-aware expansions in a predict-select-refine loop, and the selected text is synthesized by a zero-shot voice-cloning TTS model conditioned on a short user/donor reference clip.

End-to-end pipeline




Figure 1: Intended end-to-end AAC pipeline: the user enters shorthand, the LLM proposes top- k context-aware expansions in a predict-select-refine loop, and the selected text is synthesized by a zero-shot voice-cloning TTS model conditioned on a short user/donor reference clip.

End-to-end pipeline




Figure 1: Intended end-to-end AAC pipeline: the user enters shorthand, the LLM proposes top- k context-aware expansions in a predict-select-refine loop, and the selected text is synthesized by a zero-shot voice-cloning TTS model conditioned on a short user/donor reference clip.

End-to-end pipeline




Figure 1: Intended end-to-end AAC pipeline: the user enters shorthand, the LLM proposes top- k context-aware expansions in a predict-select-refine loop, and the selected text is synthesized by a zero-shot voice-cloning TTS model conditioned on a short user/donor reference clip.

**Voice
for Purpose.**

LLM Track

Models tested for prediction:

Open source

- Llama 8B, base and instruct
- Qwen 1B, 3B, 7B, base and instruct
- Microsoft Phi-3-small 7B
- OpenHermes 2.5 Mistral 7B

Close source

- GPT 4.1 and 4.1 mini



Finetune Llama on LEONARDO/Cineca.

- **Llama 8B vs 70B**
- **10 languages vs 24 languages**
- **2 interaction modes** for prediction: **word-level** refinement vs **letter-level** refinement


Goals:

- **Optimize Keystroke Saving Rate (KSR) and increase input speed**
- **Single, unified model for all Eu languages**

Voice
for Purpose.

LLM Track: Results

Keystroke Saving Rate per language (Llama 8B, letter, 24L)



Model Speed: 8B model runs at ~0.85–0.90 sec per call

Significant reduction in input actions: ~ 62.8% KSR avg on 24 langs

Subjective effort reduction: 4.61 ± 0.5 (on a 5 points scale)

Subjective Speed improvement: 200% - 300% (language, context)

Voice for Purpose.

TTS Track



- **New model: Koel-TTS** (LLM-based TTS)
 - Extended to **all 24 official EU languages** using only **openly available corpora**
 - Decoder-only scaling study from **~1.1B to ~2.2B parameters**
 - Replaced original codec with **NanoCodec**: causal/streaming-friendly generation, no licensing constraints
- **Baseline: XTTS-v2** (0.5B) **10 langs**, finetuned 150h of voice datasets. Condit. 2min of specific user's voice
- **Objective and subjective evaluation**: 100 unseen sentences per language (2400 sentences)
- **Tests voices 34 speakers**: incl. professional voice actors, healthy and voice impaired people

Objective evaluation (10 overlapping languages)

- **WER**: Koel-TTS: 9,68% - XTTS-v2: 10.38%, **CER**: Koel-TTS: 3.5% - XTTS-v2: 4.89%
- **SSIM (Speaker Similarity)**: Koel-TTS: 0.91% - XTTS-v2: 0.91%,

Subjective evaluation (18 listeners, 6 langs, 5 points scale):

- easy to understand 4.33 ± 0.49 ,
- speaker similarity: 3.78 ± 0.94 (on a 5 points scale)

**Voice
for Purpose.**

HPC Computing Resource usage



AI and Data Intensive Applications Access call

Awarded 200k hours on LEONARDO (Cineca, IT): study new, large voice and language models models.

- **Llama 8B**

- single training 32 gpus 1 day: 768 h
- 10 languages: $768 * 10 = 7,680$ h
- Hyperparameter optimisations: **2,000 h**

- **Llama 70B**

- single training 64 gpus 4 days: 6,144 h
- 10 languages: $6,144 * 10 = 61,440$ h
- Hyperparameter optimisations: **6,000 h**

- **total LLM track = 77,120 h**

- **XTTS:**

- Fine-tuning: 32 GPUs , 4 days (encoder+decoder) 3,072h
- 10 trainings: $3,072 = 30,720$ h
- Hyperparameter optimisations: **5000 h**


- **KoelTTS**

- Single training: 32 GPUs for 6 days: 4,608 h
- 10 lang-specific models: $10 * 34,608$ h = **46,080 h**
- One big model covering all languages: 64 GPUs for 10 days: **15,360 h**
- Hyperparameter optimisations: **25,000 h**

- **Total TTS track: 122,160 h**

Voice
for Purpose.

User interface demo



Voice for Purpose.

EuroHPC
User Days

Dublin Ireland
22→24 September 2026

ICHEC
EUROPEAN INSTITUTE FOR
COMPUTATIONAL HEALTH

ÉIRE2026.EU
IRELAND2026.EU

EuroHPC
Joint Undertaking

Giving a voice to those who can't speak

In this powerful testimony, Camilla Mercadante shares what it means to live without a voice - and how technology can help bring it back.

[Watch the video](#)

"Each donated voice is a piece of freedom restored, a bridge that brings people closer and breaks down barriers. Your gesture is not just a contribution, but a sign of humanity that truly makes a difference."

Camilla Mercadante, SMA patient
Voice for Purpose User

Thanks

Applying Successfully

Alberto Ramos

Extreme scale calls

- Single panel for all domains
- 2 calls per year
 - ▶ If rejected reapply!
- Typical project awarded ~ 1-2 M€
 - ▶ Take the application seriously!
- Role of the panel:
 - ▶ Strict ranking of proposals

Evaluation of proposals

- ▶ Two rapporteurs per proposal
- ▶ Three scientific reviewers (in your domain)
- ▶ One technical report by the computing center
 - ▶ Questions to PI / Answers
- ▶ Three reports in total
- ▶ Final panel evaluation (meeting in person)
- ▶ Most proposals pass the bar: The question is the ranking

Key ingredients of the evaluation

- ▶ Is the science **good**?
- ▶ Does the project have **concrete** and **sound** goals?
- ▶ Are the resources **indispensable** to achieve the goals?
- ▶ Can the goals be achieved with **less resources**?
- ▶ Are resources **well used**?
 - ▶ Scaling (weak/strong) of the code
 - ▶ Runs over all the allocation period
- ▶ Is the team **capable** of managing the resources?
 - ▶ Previous use of HPC

A few pieces of advice

- ▶ Write clear targets (precision/whatever) that justify the requested resources. **Be concise!**. Why X node-hours are needed (and not X/2 or 2X) should be transparent
- ▶ **Benchmark** your production codes (not a proxy!) on the target system (not a similar one).
- ▶ Answer questions of technical team/ARC members/scientific reviewers **concisely**.
- ▶ If rejected but science is strong: apply again taking **feedback into account**
- ▶ If needed ask help from the computing centers to **show scaling and efficiency**
- ▶ Tips on “minimal allocation” (what is needed to achieve “partial goals”)
- ▶ Incremental research projects need to show **progress/dissemination. Do not send the same “revolutionary” proposal again and again**
- ▶ **You are asking for 1-2M€: Take it seriously. We are not evaluating scientists: Leading experts working on open scientific questions need to give details**