

AI & Quantum Use Cases

Draft QEX: from quantum use cases to European adoption

Evidence-driven experimentation across energy, water, mobility, Earth observation, finance and quantum simulation.

Dr. Ir. Karin Chávez Caballero, Msc. MBA.
QEX Coordinator & MC Lead | QDNL, the Netherlands

A special thank you to the QEX “Use cases and applications” team (WP5) for their contributions and support. WP5 Contact: belliard@imada.sdu.dk, giaccap@imada.sdu.dk

Dublin, Ireland | 23 September 2026 | Parallel session: AI and Quantum Use Cases



QEX in one minute

A European capability for end-users - coordinated by QDNL and delivered through a distributed consortium.

MISSION

European mission

QEX is designed as a one-stop shop for potential end-users of quantum technology - connecting application needs with quantum expertise, software, training and access to relevant QCS / HPC-QCS infrastructure.

QDNL coordinates

ROLE	Project coordination, governance and consortium alignment.
NL LINK	Connects QEX with the Dutch quantum ecosystem and end-user network.
FOCUS	A coherent route from stakeholder engagement to tested use cases.

15 Apr 2025 → 14 Apr 2029

Horizon Europe / EuroHPC JU • Project 101194491

Distributed European capability

The consortium combines:


- quantum and HPC research centres
- application and domain expertise
- software and algorithm development
- infrastructure and benchmarking capability
- stakeholder and training networks

Local anchor in Dublin: ICHEC
contributes quantum-HPC software and use cases such as QSE and QSVM4EO.



One front door - one evidence - driven journey

The user should experience QEX as one coherent pathway, not as a collection of work packages.



Leadership principle:

do not promise a “quantum solution”; create a transparent decision path and credible evidence.

Around fourteen selected use cases and reusable assets

Different sectors, modalities and maturity levels - deliberately connected through the same delivery discipline.

ENERGY, WATER & INFRASTRUCTURE

WNTR-Quantum • N-1 grid security • PyPSA-Quantum • Prosumer Coalition Formation • Railway timetabling

EARTH OBSERVATION & MARITIME

QSVM4EO land-cover classification • Quantum-Bathymetry / for seabed mapping

FINANCE

QQuantLib • Multi-objective portfolio optimisation

QUANTUM SIMULATION & PROTOCOLS

QSE • Preparation of Spinons • Holographic error correction • Analogue QC regression • Unsupervised phase discovery

Portfolio principle

Maturity is not a score of importance. It tells us what evidence is needed next - formulation, benchmark, hardware test, integration or scaling.

Maturity is intentionally mixed

REUSABLE ASSET

e.g. QQuantLib, QSE

PROOF OF CONCEPT

e.g. WNTR-Quantum, portfolio optimisation

EXPLORATORY


e.g. railway, holography

Spotlight: quantum computing for water distribution networks

A strong example of why hybrid evidence - not hype - should guide the next step.

THE END-USER CHALLENGE

Large water networks require repeatedly solving non-linear hydraulic equations. At scale, simulations and optimisation can become computationally expensive.



Water distribution network

Simulation + optimisation

HYBRID + ANNEALING

What was tested

Quantum subroutines for the Newton–Raphson Global Gradient Algorithm, including HHL, VQLS and QUBO-based linear solving; plus annealing approaches for higher-order equations and least-cost design.

SMALL-NETWORK POC

What the evidence says

On the small test network, VQLS was the most promising hybrid subroutine. Pure QUBO approaches showed growing error and significant translation overhead.

NEXT DECISION

Leadership conclusion

The near-term route is hybrid experimentation and better benchmarking - not a claim of production-ready quantum advantage.

WHAT QEX CAN ADD NOW

- 1 Define stronger classical baselines and larger benchmark instances
- 2 Test hybrid workflows on suitable EuroHPC / dedicated QCS routes
- 3 Reduce QUBO translation and orchestration overhead
- 4 Package WNTR-Quantum as a reusable experimentation asset for water operators

Energy, mobility and critical infrastructure

Operational problems with hard constraints, large search spaces and demanding classical baselines.

TNO + ALLIANDER

N-1 security in power grids

Rapid reconfiguration after a single-link failure. Quantum annealing has shown valid solutions on small networks; the gate-based route remains exploratory and needs more realistic constraint checks.

SURF + TNO + ALLIANDER

PyPSA-Quantum

A plugin that lets energy-network researchers and operators experiment with quantum routines inside the familiar PyPSA environment. Current position: very early proof of concept.

UDC

Prosumer Coalition Formation

Smart-grid coalition optimisation. Current work is building algorithm primitives for benchmarking quantum and classical approaches on an NP-hard partitioning problem.

Railroad conflict resolution under disruptions

Problem and approach

Internal QEX only.



RAILWAY TIMETABLING & DISRUPTION MANAGEMENT

TNO + NS

The use case separates long-horizon timetable construction from operational conflict resolution under disruption. Candidate quantum routes include QUBO/Ising formulations, Grover-style search, variational methods and graph-native analogue platforms.

Current position: conceptual / exploratory. Near-term value is hybrid experimentation and benchmarking on small-to-moderate instances - not full-scale deployment.

Earth observation and maritime applications

Two use cases show how quantum kernels can enter real data-science workflows.

QSVM4EO on an analog neutral atom quantum computer

February 26, 2026

1 Summary

QSVM4EO aims to perform land cover classification of Earth Observation data (EO) using quantum-enhanced support vector machine (SVM) models. Specifically, the approach analyses datasets of spectral band reflectance acquired from Sentinel-2 [1] satellites. The objective is to predict land cover classes (vegetation, bare soil, water, etc.) of each pixel within a tile based on these reflectance measures.

To implement the quantum-enhanced SVM, a quantum evolution (QE) kernel [2] is used. This kernel allows us to measure the similarity between graph-structured data based on the time evolution of a quantum system. Once the kernel is computed, we can use it to train the SVM and make predictions on land cover classes.

2 Mathematical formulation

The equation for a kernelised binary SVM is the following:

$$f(x) = \sum_i \alpha_i y_i K(x_i, x) + b \quad (1)$$

where the kernel $K(x, x')$ is a function that measures the similarity between two data points and corresponds to an inner product in a feature space:

$$K(x, x') = \langle \phi(x), \phi(x') \rangle \quad (2)$$

Different kernel methods can be used depending on the data and type of problem that is being tackled. The kernel used in this study is the quantum evolution kernel, which is suitable to be implemented in a neutral atom quantum computer.

This kernel associates takes graphs $\mathcal{G}$ and associates them with a probability distribution $\mathcal{P}_{\mathcal{G}}$. The quantum evolution kernel between two graphs $\mathcal{G}$ and $\mathcal{G}'$

ICHEC

QSVM4EO

Land-cover classification from Sentinel-2 Earth-observation data using a quantum-evolution kernel on neutral-atom systems. Classical spectral features are encoded into graph-like atom geometries, and the resulting quantum kernel feeds a classical SVM.

QEX focus

- accuracy vs classical SVM
- scaling with feature count / qubits
- execution on relevant neutral-atom backends
- reproducibility and benchmarking

TNO + S&T

Quantum Bathymetry



A hybrid quantum–classical machine-learning workflow for estimating shallow-water depth from satellite imagery. The current library is a proof of concept with explicit limitations around noise, transferability and the open question of quantum advantage.

QEX focus

- new coastal datasets
- classical vs quantum-enhanced comparison
- model robustness across regions
- operational relevance for maritime and mapping stakeholders

Earth observation and maritime applications

Two use cases show how quantum kernels can enter real data-science workflows.



ICHEC

QSVM4EO

Land-cover classification from Sentinel-2 Earth-observation data using a quantum-evolution kernel on neutral-atom systems. Classical spectral features are encoded into graph-like atom geometries, and the resulting quantum kernel feeds a classical SVM.

QEX focus

- accuracy vs classical SVM
- scaling with feature count / qubits
- execution on relevant neutral-atom backends
- reproducibility and benchmarking

TNO + S&T

Quantum Bathymetry


A hybrid quantum–classical machine-learning workflow for estimating shallow-water depth from satellite imagery. The current library is a proof of concept with explicit limitations around noise, transferability and the open question of quantum advantage.

QEX focus

- new coastal datasets
- classical vs quantum-enhanced comparison
- model robustness across regions
- operational relevance for maritime and mapping stakeholders

Finance: reusable software and real optimisation challenges

A mature library can accelerate experimentation; a domain proof of concept tests where the real bottlenecks are.



CESGA / UDC

QQuantLib

An operational Python library covering quantum algorithms and finance-oriented modules such as amplitude estimation, quantum machine learning for Value-at-Risk, and QPU integration.

Current position

Operational and available on GitHub. QEX can adapt it to the project software framework and benchmark components across different hardware.

Strategic value

Reuse strong European R&D assets instead of rebuilding them.

TNO + RABOBANK + BRIGHTLANDS

Multi-objective portfolio optimisation

A quantum-enhanced optimisation module that reformulates portfolio allocation into QUBO form, balancing competing objectives such as return, concentration risk and carbon-footprint constraints.

Current position

Proof of concept. Main issues include discretisation, classical-quantum orchestration overhead and scaling.

QEX opportunity

Benchmark meaningful business scenarios and identify which parts of the workflow actually benefit from quantum or hybrid solvers.

Quantum simulation and protocols

These cases stress-test the software stack, hardware capabilities and the limits of current quantum methods.

QSE: Quantum Simulation Environment

ICHEC

March 2026

1 QSE package

QSE is an open source [\[1\]](#) [\[2\]](#) package developed at ICHEC by the quantum-HPC team, written purely in python. It is adapted from the Atomic Simulation Environment (ASE) [\[3\]](#), which is a flexible, customisable, and easy to use python package for setting up, steering, and analysing atomistic simulation.

In analogue quantum computing, instead of quantum gates one has a set of quantum objects, each identified as qubits, which interact with each other, and instead of individually manipulating them, one evolves this quantum system in time. Several quantum computing problems can be mapped into this evolution and thus it serves as an alternative form of computation. Since the interaction among the qubits depends on their positions (unlike gate-based systems, where position has no meaning), there is a correspondence with ASE, where one manipulates the positions of the atoms, particularly systems with periodic structure.

1.1 Package summary

QSE has two main python classes, `Qbits` and `Calculator`. The `Qbits` class encapsulates the details of the system, while the `Calculator` class connects to different backends which perform the type of calculations the user would want. Initial implementation supports MyQLM and Pulser backends, which allows us to seamlessly run our calculations of interest on Qaptiva appliances, and the code design aims to enable multiple calculator backends in the future.

QSE serves as high level wrapper to facilitate analogue computation, particularly for setting up and analysing lattice-based systems. We are also adding tools within QSE that aren't readily available in different backends to bridge the gap of capability while staying backend-agnostic.

2 Analogue quantum computing

2.1 Neutral atoms

Analogue quantum computers, specifically quantum annealers let a quantum system naturally evolve under a well controlled, programmable, often time de-

1

ICHEC

QSE library

Open-source Python tooling for analogue quantum simulation of lattice-based systems. Backend-agnostic design, with MyQLM and Pulser currently supported and room for additional backends.

CESGA + INL

Spinons

Numerically tested in silico. A benchmark for demanding subroutines such as high-fidelity ground-state preparation and LCU-based non-unitary operations before hardware execution.

SDU + QPURPOSE

Holographic error correction

A quantum-gravity toy model used to test holography and error-correction ideas. Theoretical checks are defined; computational validation still lies ahead.

QML / ANALOGUE

Analogue QC regression

Quantum-kernel ridge regression to infer quantities that are hard to measure directly on neutral-atom devices. Train on smaller simulated systems; test extrapolation towards larger systems.


QML / PHYSICS

Unsupervised phase discovery

Use quantum kernels and clustering to discover structure in many-body phase space - a candidate setting where the underlying quantum state space is expensive to represent classically.

Quantum simulation and protocols

These cases stress-test the software stack, hardware capabilities and the limits of current quantum methods.



ICHEC

QSE library

Open-source Python tooling for analogue quantum simulation of lattice-based systems. Backend-agnostic design, with MyQLM and Pulser currently supported and room for additional backends.

CESGA + INL

Preparation of Spinons

Numerically tested in silico. A benchmark for demanding subroutines such as high-fidelity ground-state preparation and LCU-based non-unitary operations before hardware execution.

SDU + QPURPOSE

Holographic error correction

A quantum-gravity toy model used to test holography and error-correction ideas. Theoretical checks are defined; computational validation still lies ahead.

QML / ANALOGUE

Analogue QC regression

Quantum-kernel ridge regression to infer quantities that are hard to measure directly on neutral-atom devices. Train on smaller simulated systems; test extrapolation towards larger systems.

QML / PHYSICS

Unsupervised phase discovery

Use quantum kernels and clustering to discover structure in many-body phase space - a candidate setting where the underlying quantum state space is expensive to represent classically.

From proof of concept to EuroHPC evidence

Hardware time is most valuable when the use case, baseline and success criteria are already clear.

Creating a client journey through Hybrid Quantum Computing

<https://qdnl.nl/qex>

This project has received funding from the European High Performance Computing Joint Undertaking under grant agreement No: 101194491

VISION & POSITIONING

QEX establishes a virtual European hub to accelerate the adoption of hybrid quantum computing.


By connecting users, technology providers, research institutions and HPC infrastructures, QEX creates a seamless client journey from awareness to application deployment.

- QEX Vision and principles for engagement
- QEX as a European quantum one-stop-shop
- End-user centricity as a co-creation philosophy
- Alignment with EuroHPC, Quantum-EOSC and the European quantum ecosystem
- Guiding principles for partnerships and engagement

USE CASES

- PyPSA-Quantum
- QC2
- QSE Library
- Testing discrete holography
- WNTR-Quantum
- Hyperfine Coupling Constants
- Linear Response with HHL
- Potential Energy Surfaces

Consortium partners



Access is not automatic: eligibility, availability and the applicable EuroHPC / QEX procedures still govern execution.

What we are learning - and what we invite from you

QEX is building a European capability by turning diverse quantum experiments into comparable evidence.

- 1 Start with the problem** A material end-user challenge and a strong classical baseline are more important than starting with a preferred quantum technology.
- 2 Be honest about maturity** PoC, reusable library and benchmark-ready application are different states - and require different next steps.
- 3 Benchmark the full workflow** Algorithmic speed-up alone is not enough; encoding, orchestration, sampling, data and post-processing matter.
- 4 Reuse before rebuilding** European libraries such as QSE and QQuantLib can shorten the path from research to experimentation.
- 5 Use hardware deliberately** Match the architecture to the problem, and use scarce QCS time to answer a defined question.

A special thank you to the QEX “Use cases and applications” team (WP5) for their contributions and support. WP5 Contact: belliard@imada.sdu.dk, giaccap@imada.sdu.dk

Thank you

Questions & discussion

Bring us a challenge

A computational problem with domain relevance, committed users and a willingness to test what quantum can - and cannot - add.

ENERGY

WATER

MOBILITY

EARTH / MARITIME

FINANCE

SCIENCE

QEX can help formulate the problem, connect expertise, identify a credible route, prepare for access, benchmark honestly, and make the learning reusable.



EuroHPC
User Days

Dublin Ireland
22→24
September 2026



EuroHPC
Joint Undertaking

ICHEC
National Institute for Digital
Convergence





ÉIRE2026.EU
IRELAND2026.EU

Europe's Quantum Playground:

Experimenting, Learning, and Co-Designing
for the future







→ Through experimental use, users became participants in determining the technology's future form and applications

What is this?

- A first toy to learn, discover, dream, fail
- A knowledge, expertise and skills development machine
- A chance to be active designers in the technology
- **A use case discovery machine!**



The power is not in the machine itself (yet) is
in what we learn from it and with it.

Exploring quantum computing today for advantage tomorrow



Experimentation reveals
possibilities

Possibilities influence
technology

Technology enables
new experimentation

EuroSSQ-HPC

Consortium of NL, FR, and BE

netherlands
eScience center




Nikhef

 GENCI
HPC at the service of knowledge TU DelftUniversiteit
LeidenUniversiteit
Antwerpen

16+ qubits

Semiconductor Spin Qubit (SSQ)
Co-located with NL national HPC

HPC, quantum algorithms,
software, integration, co-design



Systematic Polymorph Crystal Structure Prediction to De-Risk Drug Development

Nested sampling for crystal structure prediction

Dr Catherine Watkinson

Lead Data Scientist, PolyChord Ltd

PhD in Astrophysics, Imperial College London


Crystal structure prediction

The problem, the costs, and why it needs a supercomputer.

Same molecule, different solid

The same molecule can settle into more than one packing. Each one is a different material.

FORM I



Packs one way. Dissolves faster, but can convert to the other form on the shelf.

FORM II



Packs another way. More stable, but slower to dissolve and harder to process.

When unexpected form appears, it can break a product.

THE ASSET AT RISK

\$2.67bn

Average R&D cost of taking one drug from discovery to launch, top-20 biopharma cohort, 2025.

Deloitte, Measuring the Return from Pharmaceutical Innovation

RITONAVIR (NORVIR)

\$250m+

Two years after launch, a more stable and much less soluble Form II appeared and forced withdrawal and reformulation. Estimated loss to Abbott.

Bučar et al., Angew. Chem. 2015, 54, 6972

ROTIGOTINE (NEUPRO)

4 years

Form II crystallised inside the patch in 2007. Off the US market from 2008 until July 2012.

UCB, US8232414B2; J. Pharm. Sci. 2015

The problem is prevalent

85%

of marketed drugs have more than one crystalline form

Gui, Pharm. Res. 2023; Karpinski, Chem. Eng. Technol. 2006

5.5

crystal forms found on average per free form, and 3.7 per salt, across 425 polymorph screens

Li et al., J. Pharm. Sci. 2025, 114, 586

2,102

crystal forms identified from 476 new chemical entities screened between 2016 and 2023, across 250 companies

Li et al., J. Pharm. Sci. 2025, 114, 586

Ritonavir Form III was found 24 years after Form II, by accident, during a study of something else. *Yao, Henry & Zhang, J. Pharm. Sci. 2023, 112, 237*

Crystal structure prediction (CSP)

Start from the molecule on its own and work out which packings are possible, and which are most likely.



One molecule

The chemical structure, and nothing else. No experimental data needed.



Generate packings

Search over space groups, cell shape, and where each molecule sits and points.



Rank by energy

Compute the lattice energy of every candidate and put them in order.



Shortlist of forms

The structures most likely to appear in a real experiment.

In 2019, computational crystal structure prediction recovered the late-appearing rotigotine form using no prior information about it. *Nature Commun. Chem.* 2019, 2, 71

CSP is challenging

Two problems pull against each other: the space to search is enormous and challenging, and the differences that decide the answer are tiny.

A very large space to search

- Which of 230 crystallographic space groups
- The shape and size of the unit cell
- Where each molecule sits and points
- The molecule's own twists about its rotatable bonds

These multiply. A single search routinely produces many thousands of candidates.

Very small energy differences

energy



Competing packings often sit within a few kJ/mol of each other, so the final ranking has to be more accurate than that gap. The cheap models that make the search affordable usually are not.

The search has to be broad and the ranking has to be accurate. No single method is good at both.

Why this is an HPC problem

The cost depends on how many structures have to be considered and how expensive each one is to evaluate.

Broad search

Cheap energy model, run across the whole candidate space

Refinement

More accurate model, applied to a shortlist

Final ranking

First-principles methods, a few structures

Every candidate structure can be evaluated independently of the others, so the work parallelises naturally. Large-scale HPC is the right place to run it.


Building this workflow at scale is what the project is for.

Our solution, built with EuroHPC

The sampler, and the evidence that it works on real chemistry.

PolyChord nested sampling

The sampler does not hunt for the best answer. It maps the whole function, so every basin is found and weighted.




PolyChord is a novel implementation of John Skilling's nested sampling.

Where conventional methods stall, it delivers full-coverage sampling of functions with multiple peaks, valleys and narrow ridges, in hundreds of dimensions.

This is the natural tool for mapping the energy function of a molecular system.

The sampler in action

Melatonin, sampled over its conformational space.



Live points are scattered across the full conformational space, with no prior knowledge of where the minima are.

As the run proceeds they contract onto the basins that carry the probability mass, and converge on the lowest-energy structures.

Nothing is seeded...
The landscape is discovered.

Validated on real chemistry problems

Before the solid state, the same sampler was put to work on a live industrial problem in solution.

The problem

- The dominant solvent in metal extraction is cyanide
- Highly toxic and environmentally damaging
- The industry collaborator needs greener molecular alternatives that still bind metal effectively


The approach

- Map the full energy landscape of the metal plus solution molecules
- Identify binding configurations and their **relative statistical weights** in solution
- Compare binding energetics across candidate molecules
- **Query at any process temperature without re-running**

The problem is different, but the technology extends naturally to CSP.

Sampling a metal-ligand system

Ammonia coordinating to a copper centre in solution: which configurations bind, and how much each matters.




A real liquid-phase problem, solved by the pipeline.

The sampler explores multiple coordination numbers and locates the octahedral arrangement as a minimum.

Progress, structures explored, and active live points are tracked throughout the run.

Blind validation against manual scans

Sampling started from a random distribution over the full configuration space, with no prior knowledge of the answer.



Minima found by the sampler for the ammonia- Cu^+ complex, overlaid on minima the collaborator's computational scientists had located by manual scan.


The pipeline recovered all expected conformations, with the correct energy ranking.

**Correct chemistry, predicted from a blind start...
Evidence the sampler works on real systems.**

Filled markers are structures found by the sampler. Outlined markers are the collaborator's reference structures.

Blind validation against manual scans

Sampling started from a random distribution over the full configuration space, with no prior knowledge of the answer.



The pipeline recovered all expected conformations, with the correct energy ranking.

System	ΔE kcal/mol
$\text{Fe}(\text{NH}_3)_6(\text{OCT})$	0.00
$\text{Fe}(\text{NH}_3)_5(\text{SPY}) + \text{NH}_3$	+61.61
$\text{Fe}(\text{NH}_3)_4(\text{SAW}) + 2\text{NH}_3$	+125.59
$\text{Fe}(\text{NH}_3)_4(\text{SQP}) + 2\text{NH}_3$	+148.82
$\text{Fe}(\text{NH}_3)_3(\text{TPY}) + 3\text{NH}_3$	+213.22
$\text{Fe}(\text{NH}_3)_3(\text{TSH}) + 3\text{NH}_3$	+225.62
$\text{Fe}(\text{NH}_3)_2(\text{BNT}) + 4\text{NH}_3$	+309.49

Stability spans over 300 kcal/mol
across the series.

Temperature response comes
from the same run, **re-queried**
rather than **recomputed**.

The next 22 months

Where the project sits today, and what the remaining programme delivers.

Two-year research project

MolyticsCSP, a Innovate-UK funded research project, commenced July 2026.

EuroHPC allocation will support vital optimisation for GPU to ensure tool is commercially viable.

VALIDATED

The sampler

Nested sampling maps the energy surface of a molecular system instead of searching it heuristically. Validated blind on molecules in solution, recovering known conformations with the correct energy ranking.



IN PROGRESS

Crystal packing

Extending the same parameterisation to space group, unit cell, molecular position and orientation, and integrating energy functionals suited to CSP.



THE TARGET

CSP at scale

Ranked crystal structure landscapes for real molecules, at a resolution useful to the people who have to choose a form.

Where else the same approach applies

The pipeline was built to find alternatives to cyanide for metal extraction. The underlying question generalises.

Catalyst recovery

Lithium, cobalt and nickel from
recycled batteries

Greener descaling agents for industrial
plant

Heavy metal removal from industrial
wastewater

Ligand-metal enzyme binding for
herbicide optimisation

More biodegradable fertiliser chelation
agents

All of them reduce to the same question: **which combination of molecules binds most strongly and stably with a given target**

Questions

Dr Catherine Watkinson

Lead Data Scientist, PolyChord Ltd

Catherine.watkinson@polychord.co.uk

Domyn | Edge

Domyn Edge: coming soon



The Why

Built for constrained environments and efficient domain specialization, benchmarking our stack against SOTA agentic models.

- Version 1.0
- Expected to be released in October 2026, alongside a detailed Tech Report.
- The weights will be fully accessible for research use on Hugging Face.
- MIT license.



Technical Challenges

- Maximizing agentic capabilities on SLMs
- Maximizing context length
- Optimizing architecture and throughput for trillion-token pretraining on European HPC



Domyn Edge

- 4 billion parameters
- Dense model with interleaved sliding window and different head size
- On/off thinking mode
- Long context: +256k
- Trained on ~13 trillion tokens

Infrastructure

Leonardo HPC

- 3,456 booster nodes, 4x A100 64GB SXM each, ~13,800 GPUs, EuroHPC Tier-0
- Login nodes are the only path to the internet
- SLURM batch only, multiple QoS with different node/wall-time envelopes
- Storage tiering: \$HOME / \$WORK / \$SCRATCH



Challenges

- Max nodes per run 256, throughput has to come from per-GPU efficiency, not from adding nodes.
- Wall time per run is 24 hours, every run is a chain of queued jobs
- A100 means BF16, not FP8.
- Limited storage (100TB) implies aggressive lifecycle management, and scratch purge discipline

Budget

- 4 Million GPU-Hours in 12 months
- Approximated Budget Breakdown:
 - pretraining run: 40%
 - architecture and tokenizer ablations: 10%
 - post-training: 25%
 - evaluation: 10%
 - data processing and generation: 15%

Domyn Edge: Tokenizer



A 115k vocabulary tuned for European languages

We extended the GPT-4 tokenizer, by training it on 5M samples for five target languages (EN, IT, ES, FR, DE).

- European average parity drops from 1.60 (GPT-4) to 1.53.
- Vocabulary size matters for small models

Tokenizer	Size	Parity ↓					Fertility ↓			
		ITA	SPA	FRA	DEU	Avg.	Code	ENG	Multi.	Avg.
Baguettotron	65k	1.48	1.45	<u>1.36</u>	1.38	1.42	3.67	1.46	2.06	2.40
GPT-4	100k	1.64	1.54	1.61	1.59	1.60	<u>2.88</u>	<u>1.40</u>	2.32	2.20
DomynEdge	115k	1.58	1.51	1.56	1.48	1.53	<u>2.88</u>	<u>1.40</u>	2.25	2.18
Llama-3.2-1B	128k	1.62	1.52	1.58	1.56	1.57	2.87	<u>1.40</u>	2.11	<u>2.13</u>
Mistral-3.1	131k	1.38	1.30	1.31	1.29	1.32	3.27	1.47	<u>1.93</u>	2.22
Qwen3	152k	1.63	1.52	1.58	1.56	1.57	2.99	1.44	2.29	2.24
GPT-4o	200k	<u>1.46</u>	<u>1.32</u>	1.37	<u>1.31</u>	<u>1.37</u>	<u>2.88</u>	1.38	1.84	2.03

Table 1: Comparison of tokenizer parity across languages and fertility across types of data between DomynEdge and a selection of popular LLMs tokenizer. Lower is better for both metrics; the best value in each column is in bold and the runner-up is underlined. Parity and fertility are averaged separately.

Domyn Edge: Architecture



Cutting-edge design, built for the edge

A dense 36-layer backbone, reshaped for on-device inference.

- Interleaved attention: only 10 of 36 layers carry a full KV cache.
- Attention capacity where it counts: RoPE on windowed layers, NoPE and head size 256 on the full-attention ones.
- Asymmetric QK-norm: removing a rescaling symmetry that hinders optimization.

Hyperparameter	value
Layers	36
Hidden size	2560
Feedforward size	9728
Attention heads	32
Attention KV heads	8
RoPE base	10,000
RMSNorm epsilon	1e-6
Attention head size	$[256, 256] + [64, 64, 64, 256] \times 8 + [256, 256]$
Attention Window size	$[F, F] + [1024, 1024, 1024, F] \times 8 + [F, F]$
Positional embeddings	$[N, N] + [R, R, R, N] \times 8 + [N, N]$

Table 2: DomynEdge model hyperparameters. R stands for “RoPE”, N for “NoPE”, and F stands for “full attention”.

Domyn Edge: Data



8.55T tokens, built on open corpora

Five data families driven primarily by web content, synthetic data, and a multilingual data skewed toward European languages.

- CC-2025-47 processed FineWeb-style for post-2023 coverage.
- SDG at scale: 600B tokens of reasoning traces over Stack-Edu C++/Python/Java and Nemotron-CC-Math.
- Open licenses throughout, PII masked, quality-gated, robots.txt, blacklisting

GENERAL WEB & NATURAL LANGUAGE

4.72T · 55%

Dolma 3 CC + FineWeb CC-2025-47

SYNTHETIC

1.68T · 20%

Nemotron pretraining sets + TPT traces + FinePhrase

MULTILINGUAL

0.88T · 10%

12 languages · FineWeb2-HQ, HPLT v3, FineWiki

CODE & TECHNICAL

0.65T · 8%

Stack Edu + GitHub Issues & Kaggle

STEM


0.62T · 7%

OLMoCR science PDFs, FineMath, ProofPile arXiv

Total 8.55T tokens · bars in trillions of tokens

Synthetic Data Generation – Agentic Trajectories

An agentic trajectory factory that runs on any agentic environment, from HPC sandboxes to live Domyn Platform tenants.



Domyn Edge: how it was trained



Training methodology



TorchTitan


NeMo RL

Torch (FSDP2)

Leonardo HPC (4xA100)

Domyn Edge on BFCL v4

Preliminary checkpoint, overall accuracy against Qwen 4B-class references



Domyn's open-source tools

Tools built for our own training pipeline, now open to the community.

Domyn Swarm Open Source

A CLI and Python library that manages vLLM inference replicas, load balancing, and job execution, built around Slurm and NVIDIA DGX Cloud Lepton

Domyn Evals Open Sourcing Soon


A unified evaluation framework for running multiple LLM benchmark suites, with Slurm integration and Domyn Swarm deployment support.

FinerFlows Open Sourcing Soon

An SDAG-based Python library for scaling HPC data workloads, covering pretraining, post-training, validation and statistics, with integrated deduplication and sequence packing.

SDAG Open Sourcing Soon


A Python library for running DAGs of jobs on Slurm clusters. It handles interdependent jobs, multiple entry points, failure recovery, and caching of completed steps.



```
src > finer_flows > workflows > demo > example.py > compute_stats
29     download_lid_176,
30 )
31
32
33 @pipeline
34 def example() -> None:
35     """Example pipeline.
36
37     Made to process a small chunk of domyn/FinSafeGuard.
38
39     Stages:
40     - Download the Domyn-Edge tokenizer
41     - Download the lid model
42     - Process the training split
43     - Compute stats
44
45     Raw data must be downloaded in `data/raw/<dataset-name>`
46     while processed artifacts should be placed in
47     `data/processed/<dataset-name>`.
48     """
49     tokenizer_task = download_domyn_edge_tokenizer()
50     lid_task = download_lid_176()
51     download_task = download_hf_dataset(
52         repo_id="domyn/FinSafeGuard", local_dir="data/raw/example"
53     )
54
55     process_task = process(
56         input_path=download_task.artifacts["local_dir"],
57         lid_path=lid_task.artifacts["lid"],
58         tokenizer_path=tokenizer_task.artifacts["tokenizer"],
59         output_path="data/processed/example/data/train.jsonl",
60     )
61
62     compute_stats(
63         input_path=process_task.artifacts["output_path"],
64         output_path="data/processed/example/stats/train.json",
65     )
66
67
68 @example.task("launch_scripts/submit_small.sh", cache_local=True)
69 def process(
70     input_path: Path,
71     lid_path: str,
72     tokenizer_path: str,
73     output_path: Artifact[Path],
74 ) -> None:
75     """Process the fin_safe_guard training split.
```

```
> Chat
Terminal
(finier-flows) [arognoni@login07 finer-flows]$ ff view example
[2026-09-01T15:50:58Z INFO sdag.commands] Compiling pipeline
[2026-09-01T15:50:58Z INFO sdag.discovery] Pipelines are search
[2026-09-01T15:50:58Z INFO sdag.discovery] Importing pipeline
```

Domyn Edge Demo





Thank you

Explore more

Contact us at info@domyn.com

Visit our website domyn.com

Follow us on [LinkedIn](#)




A Frontier-Scale Protein Language Foundation Model for Conditional De Novo Enzyme Design

Taha Koçyiğit, PhD, Head of AI

A reusable system for enzyme engineering


Discover

Build enzyme-family
training datasets



Train and design

Adapt PLMs and
select new candidates




Test and evolve

Measure function and
learn from experiments



12 enzyme families modelled. PETase candidates validated in the wet lab

EnzymeOS: protein engineering workbench



Discover
Score
Evolve
Generate

One workflow for
iterative engineering

Discovery Kits bring AI designs to the bench

Submit a sequence and
screen an AI -designed
96-well variant library.

<https://design.curecraft.co>

Enzyme programmes

PETase and nylonase
KREDs and detergent proteases



PETase: experimentally
validated activity that matches
wild type, with %50 seqID

What should the next generation of PLMs do?

Text language model

Predict the next token

Protein language model

Predict the next amino acid

Scaling shows emergent biological capabilities

Our next step: reliable design from functional requirements.

Our fourth-generation PLM programme

Model capacity

273B total-parameter target with sparse MoE

Dataset

26 billion protein samples

Learning methods

FIM, MTP and Muon optimizer

Research direction

Broader capability with richer post-training

A much larger protein sequence resource

Sequence resource	Full corpus	90% identity
ESM3	3.15B incl. synthetic	Not compared
ProGen3 / PPA-1	3.4B	1.2B clusters
CureCraft Atlas v3	26.0B	~11.5B representatives

~9.6× as many 90% identity representatives as PPA -1

Post-training for controllable enzyme design

Richer conditioning

Enzyme family, temperature and environment

Instruction fine -tuning

Connect design requirements with protein sequences

RL based alignment

Evaluate designs with simulators and predictors

Experimental feedback

Check whether improved rewards translate to activity

What EuroHPC enabled, and how we used it

**Why compute was
needed**

Broad pretraining and experiments across
model scales

How we used it

Family-model workflows and staged
foundation-model runs

A mature software base

Keeping training supplied

Parallel dataset reading and dataset
sharding

Reuse LLM advances that fit the biology

Adopt useful training methods and mature software (NVIDIA containers with NeMo / BioNeMo).

Validate them at small scale before committing compute. Feed the training properly

Our use cases do not need million-token context. Standard multi-head attention meets our current needs.

The impact we seek: more experimentally useful enzymes.

Q&A: training objectives and reward signals

Causal modelling

Predict the next residue from preceding residues

Fill-in-the-middle

Add masks in the middle of training with tags and predict them with Decoder transformer

Multi -token prediction

Add targets for multiple future residues

RLVR-style rewards

Simulator and predictor feedback, checked in the lab

Q&A: scale, evidence and next milestones

Demonstrated

12 enzyme-family models and PETase wet-lab activity

Validated model scales

1.5B, 3B, 10B, 45B experimentation

Scale-up programme

90B and 273B total-parameter milestones

Planned capability

Richer conditioning, instruction tuning and RLVR