Where Learning Closes

The ELLIS Research Programmes Read Through the Vacuum.Net Theory

J. Konstapel
Constable Research, Leiden · November 2026

Abstract

ELLIS is Europe's network of excellence in machine learning. It runs sixteen research programmes. This essay maps all sixteen and reads them through the Vacuum.Net theory. The theory starts from a single closed path. Its key notions are closure, ratios instead of absolute units, balanced ternary values and the observer as part of the path.

Three findings stand out. First, three programmes ask in parallel why learning generalises. Their answers converge on one idea: what fits, stays. Second, three programmes take symmetry groups as given. The theory derives symmetry from the loop. Third, three programmes move from the single model to structure and collective. That is a move from point to path.

The Quantum and Physics Based Machine Learning programme connects most concretely. Its members show that discrete learning works in dense regions, that control is a sum over paths, and that physical learning requires reversibility. Each motif has a direct counterpart in the theory. The essay ends with gaps in the portfolio and possible entry points.

1. Introduction

ELLIS stands for the European Laboratory for Learning and Intelligent Systems. It is a pan-European network. Machine learning is its engine of modern AI. Its research programmes bring academics together at regular meetings across Europe. The aim is free exchange, keeping talent in Europe and international visibility.

A programme is not a funded project. It is a network of top researchers around workshops. Two or three directors lead it. Ten to fifteen Fellows and Scholars join it. Each programme should hold at least two workshops a year and pays for them itself. The research itself happens in the members' own groups. The programme area is currently under review. There is no open call.

This essay has two aims. The first is to describe what the programmes do. The second is to read that work through the Vacuum.Net theory. The reading is a correspondence in the language of the theory. ELLIS does not make these claims itself.

2. Method

All sixteen programme pages on ellis.eu were read in November 2026. So were the proposal guidelines and recent workshop reports. For each programme one or more key publications by directors or members were traced. For the physics-based programme each leading member was studied separately.

The lens uses these notions from the theory:

- Everything follows from the trinity. Values are balanced ternary: −1, 0, +1.
- The loop is a real line. Light moves along it as a point. Complexity is winding.
- Only ratios matter. Length and time are human reference points.
- The observer is part of the path.

- Closure: what fits, stays.
- The closure register: keep only what is still open and what remained as a result.
- Start from the first distinction.

3. The sixteen programmes

Twelve programmes started on 9 January 2019, when ELLIS began. Four followed: Semantic and NLP in 2020, Multimodal in 2021, Graphics and Vision and Molecule Discovery in 2023. The table gives an overview.

Programme	Est.	Directors	Core of the mission
Geometric Deep Learning	2019	Bekkers, Welling (UvA); Bronstein (Oxford)	Structure, symmetry and invariance as the foundational language of deep learning
Quantum and Physics Based ML	2019	Kappen (Radboud); Zecchina (Bocconi)	Concepts from quantum and statistical physics; energy-efficient physical hardware
Theory, Algorithms and Computations	2019	Bach (ENS/Inria); Rosasco (IIT/Genoa); Hennig (Tübingen)	Close the gap between empirical performance and theory
Natural Intelligence	2019	Mathis (EPFL); Bethge (Tübingen); Dayan (MPI)	From task-centric to agent-centric intelligence
Interactive Learning and Interventional Representations	2019	Krause (ETH); Schölkopf (MPI-IS); Cesa-Bianchi (Milan)	Causal modelling between observation and intervention
Semantic, Symbolic and Interpretable ML	2020	Kersting (Darmstadt); Frasconi (Florence); Tresp (LMU)	Learning at the human level of entities, concepts and relations
Robust Machine Learning	2019	Holmes (Oxford); Kaski (Aalto); Teh (Oxford)	Principles, quantification and verification of robust ML
Human-centric ML	2019	Weller (Cambridge); Oliver (Alicante); Angelov (Lancaster)	Transparency, accountability, fairness, verifiability
ELLIS Health	2019	Rätsch (ETH); van der Schaar (Cambridge); Stegle (EMBL)	AI for biomedicine and health care
Robot Learning: Closing the Reality Gap	2019	Billard (EPFL); Peters (Darmstadt); Asfour (KIT)	How a robot should move, act and interact
Learning for Graphics and Vision	2023	Boyer (Inria); Pons-Moll (Tübingen); Agapito (UCL)	Graphics, geometry processing and 3D vision inside deep learning
ML and Computer Vision	2019	Schiele (MPI-INF); Schmid (Inria); Weiss (Jerusalem)	Forum for European computer vision
ML for Earth and Climate Sciences	2019	Camps-Valls (Valencia); Reichstein (MPI-BGC)	Earth system modelling; extremes; hybrid physics and ML
ML for Molecule Discovery	2023	Grisoni (TU/e); Hernández-Lobato (Cambridge); Schneider (Novartis)	Faster molecular discovery
Multimodal Learning Systems	2021	Snoek (UvA); Lavrač (JSI); Sebe (Trento)	Multimodal systems and foundation models
Natural Language Processing	2020	Martins (Lisbon); Gurevych (Darmstadt); Titov (UvA/ Edinburgh)	Cooperation of European NLP labs

For this essay the programmes fall into five clusters:

- **Foundations:** Geometric Deep Learning, Quantum and Physics Based ML, Theory, Semantic.
- **Agency and causality:** Natural Intelligence, Interactive Learning, Robot Learning.
- **Trust:** Robust ML, Human-centric ML.
- **Domains:** Health, Earth and Climate, Molecule Discovery.
- **Perception and language:** Computer Vision, Graphics and Vision, Multimodal, NLP.

4. Three questions ELLIS asks in parallel

4.1 Why does learning generalise?

Three programmes work on this question, each in its own way.

The Theory programme says openly that theory lags behind practice. Its 2025 Oberwolfach workshop, held with the Interactive Learning programme, focused on overparameterisation, regularisation, identifiability and uncertainty. Forty-eight researchers took part.

The physics-based programme gives an answer from statistical physics. Loss landscapes contain a few very wide, flat minima next to many narrow ones. Good algorithms find the wide ones. Solutions there sit in regions of high local entropy. This holds for discrete weights too.

The Robust ML programme gives a third answer. Grünwald and colleagues developed e-values. Their error guarantee holds no matter when data collection stops.

Read through the theory, these are three views of one thing. A wide flat minimum is a basin of closure. Many neighbouring solutions also fit. That makes the solution stable. An e-value is a state that is complete at every moment. That is what the closure register asks for.

4.2 Where does symmetry come from?

Geometric Deep Learning wants symmetry, structure and invariance to be the foundational language of deep learning. It derives common architectures from symmetry groups. A well-known example is the E(n)-equivariant graph network of Welling's group. It respects rotations, translations, reflections and permutations. Graphics and Vision and Molecule Discovery build on the same tools.

In all three programmes the group is given in advance. The theory turns this around. The loop comes first. Symmetry is what stays the same after a full turn of the path. A gauge is the local frame of an observer on the strand. One open question follows: why E(n), and not the group of the loop itself?

4.3 From point to path

Three programmes are moving from the single model to structure.

- Natural Intelligence: the Tolman-Eichenbaum Machine separates the structure of a map from what lies on it. One relational basis serves space and abstract knowledge.
- Interactive Learning: causal representation learning assumes that only a few modules shift when conditions change. The rest stays stable.
- Human-centric ML: the 2026 Alicante workshop concluded that "human-in-the-loop" is no longer enough. Human-AI systems are now treated as sociotechnical infrastructure. Accountability sits at the level of collectives and institutions.

In the theory's terms each case separates the path from what lies on it. Stable modules are closed windings. Only open parts move. The shift from model to network is the shift from point to path.

5. The physics-based programme in depth

Quantum and Physics Based Machine Learning aims to use quantum and statistical physics for new algorithms. Its stated goal is new, energy-efficient physical hardware. It is the only programme with that goal. Its members show three recurring motifs.

5.1 Discrete learning works in dense regions

Baldassi, Zecchina and colleagues studied a network with weights of only −1 and +1. Standard analysis says solutions are isolated and almost impossible to find. In practice simple algorithms find them easily. The reason is that rare, very dense clusters of solutions exist. Those solutions are robust and generalise better than typical ones.

A follow-up extended this to several discrete levels. The picture stays the same. Capacity saturates quickly as the number of levels grows. Biological synapses store only one to five bits.

This is the most direct evidence in the ML literature that discrete computing works. It supports ternary values. It matches BitNet b1.58, a language model outside ELLIS whose weights are all −1, 0 or +1. In the language of the theory a dense cluster is a region where many neighbours also close. It is "what fits, stays" expressed in numbers.

5.2 Control is a sum over paths

Kappen (Radboud) showed that for a class of control problems the optimal choice is a path integral. All trajectories count, weighted by their action. At low noise one path dominates. At high noise the best choice is a mixture of paths. At a critical noise level the symmetry breaks. Control becomes inference.

This is the closest conceptual link found in ELLIS. The path is the basic object. The single "deterministic" path is a limiting case. Symmetry breaking is the moment one winding closes.

5.3 Physical learning requires reversibility

Marquardt (MPI for the Science of Light) and López-Pastor designed machines that learn through their own physical dynamics. No external digital feedback is needed. The process must be time-reversible and non-linear. The method runs the dynamics backwards like an echo and adds a small error signal. The motion itself then delivers the learning step. The team works with an experimental group on an optical machine.

Here light as carrier and walking the path back meet. In the theory, rewinding is a basic operation. This line is closest to the new computing device proposed in Computing on the Net.

5.4 Further members

- Carleo uses a neural network as a compressed wave function. Learning means finding which correlations close.
- Briegel lets an agent walk through its own memory and project itself into possible futures before acting. The observer walks the path before taking it.
- Noé uses invertible networks to find equilibrium states directly instead of simulating step by step. Reversibility appears again.

The three motifs map onto the theory one to one. Discrete dense regions match the trit and closure. Sums over paths match the path. Reversibility matches rewinding.

6. Readings of the other programmes

Robot Learning

The reality gap comes from omitted physics, wrong parameter estimates and discrete numerical integration. There is consensus that more accurate simulators alone do not close it. The working method is domain randomization. It lets go of fixed parameters. This supports a reading from the theory: simulation computes in absolute units, the world works in ratios.

Earth and Climate

The programme argues for hybrid models that join physics and data. Its 2024 workshop noted that foundation models trained on satellite and weather data lack interpretability, consistency and robustness. Hybrid means two layers. That touches the author's two-layer reading of El Niño.

Multimodal Learning

This programme joins modalities after the fact. In the theory there is one carrier. Modalities are aspects of one path. Integration then becomes the starting point, not the last step.

Natural Language Processing

The directors list their own open problems. Language models hallucinate too often. They give no reliable uncertainty. They are weak at reasoning, common sense, grounding and physical intuition. They are costly and too English-centred. Each item is a point where the author's claim-bounded system maze.ai formulates an answer: a gate against invention, an explicit "I do not know", one carrier, and ternary, register-based computing.

Semantic and Symbolic ML

De Raedt argues that neuro-symbolic AI is neural plus logical plus probabilistic. The integration must keep the pure forms as special cases. That is a triad. Balanced ternary offers a natural logic: true, false, open.

Vision, Graphics and Health

Vision workshops call vision a uniquely good testing ground because of its ties to perception and the physical world. In the theory, seeing is itself a motion of light. Gaussian splatting builds scenes only where something exists. Nothing is computed in empty space. Health studies cell perturbation: intervene on a path and observe what closes.

7. Gaps in the portfolio

Seen from the theory, five topics have no programme:

1. No programme works from ratios or ternary representation. BitNet shows this works outside ELLIS.
2. No programme studies computation that keeps only what is unresolved and the result.
3. No programme derives modalities from one carrier.
4. No programme treats the observer and perspective as part of the system.
5. No programme asks what learning minimally is: the first distinction.

8. Entry points

The most promising contacts follow from the analysis:

- Kappen (Radboud): path integrals and control. A Dutch entry.
- Zecchina and Baldassi (Bocconi): discrete weights and dense clusters.

- Marquardt (Erlangen): self-learning physical machines with light.
- Welling and Bekkers (UvA): equivariance and symmetry.
- Grünwald (Leiden/CWI) and Koolen (Twente/CWI): e-values and safe testing.
- Hennig (Tübingen): contact for the Theory programme.

9. Scope

This essay is based on public pages, workshop reports and selected key publications. It does not cover every member's work. The reading through the theory is a correspondence. It shows how the programmes can be described in the language of the theory. It does not report experiments.

10. Conclusion

ELLIS is a network, not a single research agenda. Yet its programmes circle a few shared questions. Why does learning generalise? Where does symmetry come from? How does a system move from point to structure?

The Vacuum.Net theory gives one answer to all three: closure on a path. The physics-based programme comes closest. Its members already show that discrete learning works, that control is a sum over paths, and that physical learning needs reversibility. The gaps in the portfolio mark where the theory could add something new: ratios, ternary values, a closure register, one carrier and the observer.

Annotated references

ELLIS (2026). Research Programs. https://ellis.eu/research/programs

The primary source. Lists all sixteen programmes. Each programme page gives mission, founding date, directors and members. Start here for any follow-up.

ELLIS (2026). Guidelines for Writing Proposals for ELLIS Programs. https://ellis.eu/research/programs/guidelines-proposals-ellis-programs

Explains how programmes are formed: two or three directors, ten to fifteen Fellows and Scholars, two workshops a year, own funding. Shows that programmes are networks, not funded projects.

ELLIS News (2022–2026). Workshop reports and interviews. https://ellis.eu/news

Reports on the Oberwolfach workshop (2025), the HCML workshop in Alicante (2026), the vision workshops in Bad Teinach (2025) and Sofia (2026), the Multimodal annual meeting (2026) and the NLP directors interview (2024). Best source for current open questions.

Bronstein, M., Bruna, J., Cohen, T., Veličković, P. (2021). Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. arXiv:2104.13478

The blueprint of the GDL programme. Derives common architectures from symmetry groups. Read it to see which symmetries are assumed rather than derived.

Satorras, V. G., Hoogeboom, E., Welling, M. (2021). E(n) Equivariant Graph Neural Networks. ICML. arXiv:2102.09844

A network that respects rotations, translations, reflections and permutations. Widely used for molecules. A concrete example of symmetry built into the architecture.

Baldassi, C., Ingrosso, A., Lucibello, C., Saglietti, L., Zecchina, R. (2015). Subdominant Dense Clusters Allow for Simple Learning and High Computational Performance in Neural Networks with Discrete Synapses. Phys. Rev. Lett. 115:128101. arXiv:1509.05753

Key paper for this essay. Weights of only −1 and +1. Simple algorithms find rare, dense clusters of solutions that are robust and generalise well.

Baldassi, C. et al. (2016). Learning may need only few bits of synaptic precision. arXiv:1602.04129

Extends the dense-cluster result to several discrete levels. Capacity saturates quickly with the number of states. Relevant to ternary computing.

Baldassi, C., Pittorino, F., Zecchina, R. (2020). Shaping the learning landscape in neural networks around wide flat minima. PNAS 117(1):161–170

Shows that a few wide flat minima coexist with many narrow ones. Good learning algorithms are drawn to the wide ones. Covers both discrete and continuous weights.

Zdeborová, L., Krzakala, F. (2016). Statistical physics of inference: thresholds and algorithms. Advances in Physics. arXiv:1511.02476

A long pedagogical review. Treats learnability as a phase transition. Useful background on how physicists read learning problems.

Kappen, H. J. (2005). Path integrals and symmetry breaking for optimal control theory. J. Stat. Mech. P11011. arXiv:physics/0505066

Optimal control as a weighted sum over all paths. One path wins at low noise; a mixture of paths at high noise; symmetry breaks at a critical value. The closest conceptual link found in ELLIS.

López-Pastor, V., Marquardt, F. (2023). Self-Learning Machines Based on Hamiltonian Echo Backpropagation. Phys. Rev. X 13:031020

Physical systems that learn through their own dynamics. Conditions: time-reversible and non-linear. Uses a time-reversal echo. Aimed at optical hardware.

Carleo, G., Troyer, M. (2017). Solving the quantum many-body problem with artificial neural networks. Science 355:602–606

A neural network as a compressed wave function. Founding paper of neural quantum states.

Briegel, H. J., De las Cuevas, G. (2012). Projective simulation for artificial intelligence. Sci. Rep. 2:400

An agent that walks through its own episodic memory and projects itself into possible futures before acting. Relevant to the question of the observer.

Noé, F., Olsson, S., Köhler, J., Wu, H. (2019). Boltzmann generators. Science 365:eaaw1147

An invertible network maps a simple distribution onto equilibrium states. Finds stable states directly instead of simulating step by step.

Hennig, P., Osborne, M. A., Kersting, H. P. (2022). Probabilistic Numerics: Computation as Machine Learning. Cambridge University Press

Treats numerical routines as inference. A useful contrast with a computing model that keeps only what is unresolved and the result.

Whittington, J. C. R. … Behrens, T. E. J. (2020). The Tolman-Eichenbaum Machine. Cell

Separates structural codes from sensory codes in the hippocampus. One relational basis for space and non-space.

Schölkopf, B., Locatello, F. et al. (2021). Toward Causal Representation Learning. Proc. IEEE 109(5):612–634

States the central problem of discovering causal variables from raw observations. Introduces the sparse mechanism shift hypothesis.

Manhaeve, R. … De Raedt, L. (2018). DeepProbLog: Neural Probabilistic Logic Programming. NeurIPS. arXiv:1805.10872

Combines neural, logical and probabilistic reasoning. Requires that the pure forms remain special cases.

Grünwald, P., de Heide, R., Koolen, W. (2024). Safe Testing. J. R. Stat. Soc. B 86(5):1091–1128

E-values: error guarantees that hold whenever data collection stops. A Leiden/CWI contribution within the Robust ML programme.

Muratore, F. … Peters, J. (2022). Robot Learning from Randomized Simulations: A Review. Frontiers in Robotics and AI

Causes of the reality gap and the case for domain randomization. Notes the consensus that better simulators alone do not close the gap.

Aljalbout, E. et al. (2026). The Reality Gap in Robotics: Challenges, Solutions, and Best Practices. arXiv:2510.20808

Recent survey. Treats the reality gap as many sub-gaps.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G. (2023). 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM Trans. Graph. 42(4)

Scenes built from local elements that exist only where there is something. No computation in empty space.

Reichstein, M., Camps-Valls, G. et al. (2019). Deep learning and process understanding for data-driven Earth system science. Nature 566:195–204

The founding argument for hybrid physics-aware ML in Earth science.

Ma, S. et al. (2024). The Era of 1-bit LLMs: All Large Language Models are in 1.58 Bits. arXiv:2402.17764

Outside ELLIS. Every weight is ternary {−1, 0, 1}. Matches full precision at much lower cost. Opens the way to ternary hardware.

Konstapel, J. (2026). Vacuum.Net Theory, Ninth Edition; Computing on the Net; Provable AI and The Interpreter Behind the Gate. https://constable.blog

The author’s own work used as the lens. The Ninth Edition is the current statement of the theory.