Geometric Deep Learning for Glassy Materials

1 Internship context

Research lab: LISN, team A&O, Inria team TAU (joint team between Inria, CNRS and LISN, Université Paris-Saclay)
Location: LISN (building 660 "Digiteo", Université Paris-Saclay)
Supervision: François Landes, Associate Professor (MCF) (francois.landes@universite-paris-saclay.fr);
Luca Teodorescu, 2nd-year PhD student (luca.teodorescu@inria.fr)

2 Problem statement

Context: geometric deep learning Graph Neural Networks (GNNs) are a well-known kind of Deep Networks, allowing to enforce node-permutation equivariance (re-indexing), a symmetry present in all graphs. Likewise, **rotation-equivariant networks** allow to enforce rotation equivariance at the level of the very architecture (it constrains the choice in aggregation or update steps). This is a game-changer that many are comparing to the introduction of CNNs and then GNNs, and is rapidly taking over the other methods for physics/chemistry-derived problems, where a **learned representation of the input should not depend on the viewpoint** (a rotation of the system of coordinates). The paradigm of (rotation) equivariant GNNs now dominates the field of molecular properties prediction (or ab initio replacements with machine learned energies and forces), but also other fields (using other symmetry groups), and is known as **geometric deep learning** [BBCV21]. Formally, a map $\phi : X \to Y$ (e.g. a network) is equivariant with respect to a group G if

$$\forall g \in G : \qquad \rho_Y(g) \circ \phi = \phi \circ \rho_X(g),$$

where ρ_X, ρ_Y are group representations of G in the spaces X, Y, respectively.

Context: glassy materials We propose a test case that offers ideal conditions for **studying and designing new equivariant GNNs**: glassy materials. The task, summarized in Fig. 1, is a node-wise regression on 3D point clouds, so **no prior knowledge of glass physics is needed**. This test case matters in its own right, since understanding glassy dynamics is a long-standing open problem in physics. It is also a hard problem[1]: solving it satisfactorily requires the design of extremely expressive and sensitive GNNs. It also differs from the more common benchmarks, which usually involve **sparse assemblies of atoms** (molecules, or at most an adsorbate on a surface): **here, space is densely packed with particles**. Predicting glassy dynamics has attracted a surge of interest recently [JAB+23], and several approaches now compete for the state of the art.

[1] 2021 Physics Nobel laureate Giorgio Parisi devoted part of his career to this notoriously hard problem.

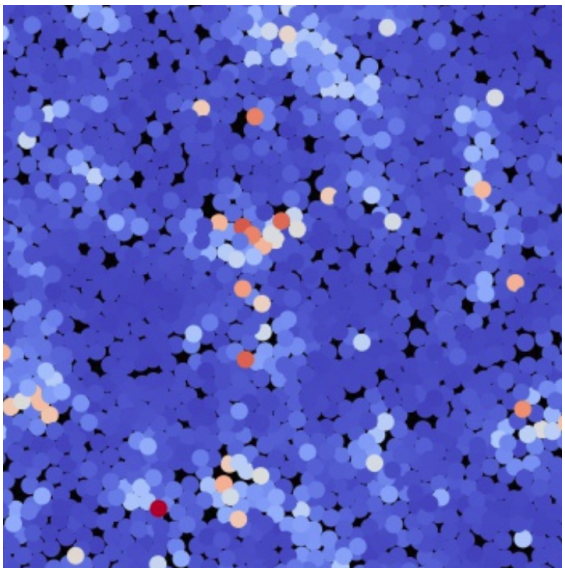


Figure 1: **Problem statement**. The input is the 3D location of all particles (here shown as disks) in the simulation box. Here only a 2D slice is shown, black areas mean no particle was present at this location of the slice. The target value is the mobility (indicated by color) of each particle. Notice the heterogeneous nature of the target field. Current SOTA is a correlation coefficient of $\rho \sim 0.8$. Maximum reachable is about 0.95. To the naked eye, it is impossible to identify a structure-dynamics relationship.

3 State of the art

Work done in our team Francesco Saverio Pezzicoli (PhD defended in December 2024) developed a very competitive equivariant GNN model [PCL24, JAB+23]. Inspired by the NequIP architecture [BMS+22], it combines local node features with the relative positions of nodes using Clebsch–Gordan tensor products, following the Tensor Field Networks approach [TSK+18]. We refer to [PCL24] and to Francesco's PhD thesis [Pez24] for a full description.
Luca Teodorescu then vastly improved on these results by combining MACE blocks [BKS+22] with more standard equivariant convolutions [BMS+22] (publication in preparation). Geometric deep learning for physics and chemistry is a very active area, from wich MACE emerged. We are interested in other original approaches coming from that field too (see internship topics).

Other methods for glassy dynamics prediction For the specific problem of predicting the dynamics of glassy materials, the state of the art was redefined in late 2022 by our SciPost paper [PCL24] and a few other independent works [SHSS23, ASF23, JBB23]. Notably, these four contributions largely surpassed **DeepMind's 2020 state of the art** [BKGB+20], with **ours leading on a range of sub-tasks** (and being the only equivariant approach). A review of the interplay between glassy physics and machine learning, including dynamics prediction and a benchmark of the various methods, can be found in [JAB+23], to which we contributed.

4 Internship Topics

We have many ideas of separate, mostly independent internship topics. They come in three flavors: computer science, physics or mathematics. Below are these well-defined subjects, ready to be investigated, but your own ideas are welcome too. **Each bullet point is essentially an internship subject**. If one of them interests you, we can discuss the details together.

4.1 CS oriented: architecture and training design

- **Equivariant transformers**: combine equivariant attention layers [FWFW20, LS22, LWDS23, LHS+26] with our current MACE+NequIP codebase. This may raise stability issues, so the integration may not be straightforward. Beyond attention, these recent works contain other interesting ideas worth importing, as their way to use gated activation, and choice of batch/layer-norms.

- **Multi-task decoding**: mobility can be measured, and predicted, at any time[2]. We propose a single timescale-aware decoder, conditioned on an embedding of the timescale as in FiLM [PSDV+18] (see also [GB22]), and we propose to **improve over FiLM using Laguerre polynomials**, an orthogonal family. Similarly, the backbone can be trained on several temperatures at once, with a temperature-aware decoder. Ultimately, combining both would **yield a decoder that is both timescale-aware and temperature-aware**. Alternatively, the decoder could be replaced by symbolic regression models, yielding a very interpretable model.

- **Improved training**: data augmentation (e.g. using short-time future positions as additional input features when predicting long-time mobility), pre-training on a pretext task (e.g. predicting the current potential energy of each particle), or **self-supervised learning**, e.g. **denoising** (adding non-physical noise to the input positions and training the network to remove it). This would be especially meaningful when using Inherent Structures' positions as input, since there is less randomness in these arrangements.

- **Beyond energies and forces in molecular simulations.** The major recent successes of equivariant models (MACE, AIMNet) consist in replacing the first principles computations of forces with Machine Learned approximations, to speed up molecular dynamics simulations, but still integrate the dynamics one timestep at a time. Some other recent methods (FlashMD [BCKC26]) attempt to directly learn the transition density over relatively long timescales, using diffusion or flow models, obtaining orders of magnitude accelerations with respect to molecular dynamics simulations. We propose to contribute to this effort on the methodological side. After a more careful literature review, one could use one of these methods (such as FlashMD) on Glass data, and possibly our model on their data. Then the question would be, what guarantees can one have when using such methods, and/or what classes of model families best approximate the flow in configuration space?

4.2 Physics oriented: task design and modeling

- **Vector targets**: *(this actually involves CS, physics and maths)* fully leverage the equivariance of the network to predict the displacement vector (three components in 3D) rather than only its magnitude. At fixed backbone, the question is **whether diversifying the task helps learning**. This is quite an **original task**: we are not aware of any work predicting vector quantities as a pretext task and studying how this choice impacts generalization. There are interesting questions to address in how to average over the configurations of the iso-configurational ensemble.

- **A new system**: in collaboration with Sylvain Patinet (PMMH, Paris), apply our networks to 2D amorphous solids under athermal quasi-static (AQS) shear [PVF16],

[2]Each timescale is de facto a new task, but these tasks are not independent.

where the targets are inherently vectorial and multivariate.

- **Interpretability**: revisit how well very shallow (1- or 2-layer) networks with large body order and ℓ_{max} can perform. If they do reasonably well, interpret them as far as possible by unfolding their explicit equations. This would combine naturally with the Multi-task decoding subject.
- **Reduced-order model of the dynamics**: *(very ambitious project)* use the network's learned features as a local state variable, in the spirit of what was done with the "softness" field [ZRL21, ZXY$^+$22, RTL23], in particular, focusing on the direction explored in [ZRL21]. This involves running short molecular dynamics simulations, extracting local features with our best network, and modeling the evolution of this feature field.

4.3 Maths oriented: learning theory

- **Lazy vs. rich learning**: determine in which regime our networks are trained, why, and possibly how to switch from one regime to the other [COB19, KRD$^+$24].
- **Learning stability**: empirically, (equivariant) batch normalization is crucial for stable training, but why remains poorly understood. A meticulous study of the learning dynamics and of the role of this specific equivariant batch normalization could yield valuable insights. This is partly experimental, partly theoretical.

The internship could lead to a PhD thesis, provided there is mutual interest.

5 Expected skills

We are looking for motivated students, from computer science, physics or applied mathematics. Required:

- Proficiency in Python and PyTorch
- Scientific rigor

A plus, depending on the chosen direction:

- Some knowledge of GNNs (and possibly of PyTorch Geometric)
- Some knowledge of geometric deep learning (and possibly of e3nn or a similar package)
- A taste for theory (e.g. the maths underlying equivariant networks, or learning theory)

For MVA students, the related topics and courses are listed below. Students from other Master's programs are equally welcome to apply.

- **Representation learning**: Representation Learning for Computer Vision and Medical Imaging
- **Deep learning**: Deep Learning; Deep Learning in Practice
- **Geometric deep learning**: Geometric Data Analysis (esp. session 4)

- **GNNs**: Advanced Learning for Text and Graph Data (ALTEGRAD), esp. the graph part (not NLP)

Somewhat related but not necessary:

- Nuages de Points et Modélisation 3D (NPM3D)
- Geometry Processing and Geometric Deep Learning (not required; actually not that relevant despite the title)
- Graphs in Machine Learning (although our graphs are actually quite regular)

References

[ASF23] Rinske M. Alkemade, Frank Smallenburg, and Laura Filion. Improving the prediction of glassy dynamics by pinpointing the local cage. January 2023. arXiv:2301.13106 [cond-mat].

[BBCV21] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

[BCKC26] Filippo Bigi, Sanggyu Chong, Agustinus Kristiadi, and Michele Ceriotti. Flashmd: long-stride, universal prediction of molecular dynamics. *Advances in Neural Information Processing Systems*, 38:88537–88574, 2026.

[BKGB+20] V. Bapst, T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. S. Schoenholz, A. Obika, A. W. R. Nelson, T. Back, D. Hassabis, and P. Kohli. Unveiling the predictive power of static structure in glassy systems. *Nature Physics*, 16(4):448–454, April 2020.

[BKS+22] Ilyes Batatia, David P Kovacs, Gregor Simm, Christoph Ortner, and Gábor Csányi. Mace: Higher order equivariant message passing neural networks for fast and accurate force fields. *Advances in Neural Information Processing Systems*, 35:11423–11436, 2022.

[BMS+22] Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1):2453, 2022.

[COB19] Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

[FWFW20] Fabian B Fuchs, Daniel E Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-translation equivariant attention networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 1970–1981, 2020.

[GB22] Jayesh K Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized pde modeling. *arXiv preprint arXiv:2209.15616*, 2022.

[JAB+23] Gerhard Jung, Rinske M Alkemade, Victor Bapst, Daniele Coslovich, Laura Filion, François P Landes, Andrea Liu, Francesco Saverio Pezzicoli, Hayato Shiba, Giovanni Volpe, et al. Roadmap on machine learning glassy liquids. *arXiv preprint arXiv:2311.14752*, 2023.

[JBB23] Gerhard Jung, Giulio Biroli, and Ludovic Berthier. Predicting dynamic heterogeneity in glass-forming liquids by physics-inspired machine learning. *Physical Review Letters*, 130(23):238202, 2023.

[KRD+24] Daniel Kunin, Allan Raventós, Clémentine Dominé, Feng Chen, David Klindt, Andrew Saxe, and Surya Ganguli. Get rich quick: exact solutions reveal how unbalanced initializations promote rapid feature learning. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

[LHS+26] Yi-Lun Liao, Alexander J Hoffman, Sabrina C Shen, Alexandre Duval, Sam Walton Norwood, and Tess Smidt. Equiformerv3: Scaling efficient, expressive, and general se (3)-equivariant graph attention transformers. *arXiv preprint arXiv:2604.09130*, 2026.

[LS22] Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d atomistic graphs. In *The Eleventh International Conference on Learning Representations*, 2022.

[LWDS23] Yi-Lun Liao, Brandon Wood, Abhishek Das, and Tess Smidt. Equiformerv2: Improved equivariant transformer for scaling to higher-degree representations. *arXiv preprint arXiv:2306.12059*, 2023.

[PCL24] Francesco Saverio Pezzicoli, Guillaume Charpiat, and François Pascal Landes. Rotation-equivariant graph neural networks for learning glassy liquids representations. *SciPost Physics*, 16(5):136, 2024.

[Pez24] Francesco Pezzicoli. *Statistical Physics-Machine Learning Interplay: from Addressing Class Imbalance with Replica Theory to Predicting Dynamical Heterogeneities with SE (3)-equivariant Graph Neural Networks*. PhD thesis, Université Paris-Saclay, 2024.

[PSDV+18] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

[PVF16] Sylvain Patinet, Damien Vandembroucq, and Michael L. Falk. Connecting local yield stresses with plastic activity in amorphous solids. *Physical Review Letters*, 117(4):045501, 2016.

[RTL23] Sean A. Ridout, Indrajit Tah, and Andrea J. Liu. Building a "trap model" of glassy dynamics from a local structural predictor of rearrangements. *Europhysics Letters*, 144(4):47001, November 2023. arXiv:2308.07913 [cond-mat.soft].

[SHSS23] Hayato Shiba, Masatoshi Hanai, Toyotaro Suzumura, and Takashi Shimokawabe. BOTAN: BOnd TArgeting Network for prediction of slow glassy dynamics by machine learning relative motion. *The Journal of Chemical Physics*, 158(8):084503, 02 2023.

[TSK+18] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. *arXiv (unpublished)*, May 2018. arXiv: 1802.08219.

[ZRL21] Ge Zhang, Sean A. Ridout, and Andrea J. Liu. Interplay of Rearrangements, Strain, and Local Structure during Avalanche Propagation. *Physical Review X*, 11(4):041019, October 2021.

[ZXY+22] Ge Zhang, Hongyi Xiao, Entao Yang, Robert J. S. Ivancic, Sean A. Ridout, Robert A. Riggleman, Douglas J. Durian, and Andrea J. Liu. Structuro-elasto-plasticity model for large deformation of disordered solids. *Physical Review Research*, 4(4):043026, October 2022.