

Geometric Deep Learning for Glassy Materials



1 Internship context

Research lab: INRIA TAU team (joint team between INRIA, CNRS and LISN of Université Paris-Saclay)

Location: LISN (building 660 “Digiteo”, at Université Paris-Saclay)

Supervision: François Landes (francois.landes@universite-paris-saclay.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, equivariance of the map $\phi : X \rightarrow Y$ (e.g. a network) can be written, for the group G :

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

Where ρ_X, ρ_Y are group representations in the space X, Y (respectively).

Context: Glassy materials We propose to work on a test which represents ideal conditions for **studying and designing new equivariant GNNs**: the case of glassy materials. The task is summarized in Fig. 1. This test case has its own scientific significance for theoretical physics, but also represents a hard problem¹, so that solving it satisfactorily necessitates the design of extremely expressive, sensitive GNNs. It is somehow different from the more common benchmarks, which usually involve **sparse assemblies of atoms** (as molecules, or at most an adsorbant), **here instead space is densely packed with particles**. Predicting glassy dynamics has known a surge of interest recently [JAB⁺23], and the state of the art is now disputed by different approaches.

¹Physics Nobel Prize 2021 recipient Giorgio Parisi spent a part of his career on 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, a PhD student soon to defend his PhD, has developed a very competitive equivariant GNN model [PCL24, JAB⁺23]. It is a network inspired by the NequIP [BMS⁺22] architecture, where we combine local node features with the relative positions of nodes, using Clebsh-Gordan tensor products, following the Tensor Field Networks approach [Thomas et al 2018]. We do not fully describe our model here, and refer to [PCL24] or the soon-to-be-released PhD thesis of Saverio for a full description.

Other methods for glassy dynamics prediction For the specific problem of predicting glassy materials’ dynamics, the state of the art has been redefined in late 2022: our own SciPost paper [PCL24], along with a few other independent works [SHSS23, ASF23, JBB23]. It’s worth noting that **Deepmind’s redefinition of the SOTA** (from 2020, [BKGB⁺20]) has been largely surpassed by these 4 independent contributions, with **ours in leading position over a range of sub-tasks** (and the only truly equivariant approach). A good review about the interplay between glassy physics questions and machine learning, including predicting dynamics, and a benchmark of the various methods, can be found in [JAB⁺23] (soon to be published as a Nature Physics Review).

Geometric Deep Learning applied to physics/chemistry in general This is a very active area, and we are interested in the newer original approaches such as MACE [BKS⁺22] or the equivariant attention schemes that are being developed, such as those in [LS22, HLLZ⁺21] or [FWFW20]. Part of the intern’s work would be to get up to date with the most recent/promising developments.

4 Scientific proposals

There are multiple avenues to explore that look promising. Here are some of the possible directions that are well-defined and ready to be studied, but this is not exclusive of your own ideas:

- Fully leveraging the equivariant properties of the network to predict directly the displacement vector (comprising three components in 3D space), rather than just its scalar magnitude. At equal backbone, one is interested to **understand whether diversifying the task can help learn better**. This is quite an **original task**: we are not aware of a paper actually making predictions for vector quantities and studying how this choice impacts generalization.
- As in [JBB23], adding non-local quantities as additional target labels (i.e., adding terms in the loss function), such as global correlation functions evaluated at specific lengths (computed for the entire sample, resulting in a graph-wide target), or the **local variance of the target value** (variance within a node’s neighborhood). This could improve prediction quality, particularly in terms of spatio-temporal correlations, and help **address over-smoothing, a known issue in GNNs**. Studying carefully the impact of these auxiliary loss terms on the main task could help **draft generic strategies to fight over-smoothing** (for any node-wise prediction task).
- Decoding various timescales² using a single timescale-aware decoder, similar to FiLM [PSDV⁺18] (conditioning the decoder with an embedding of the timescale, as proposed in [GB22], allowing for a single final decoder). Here **to improve over FiLM, we propose to use Laguerre polynomials**, an orthogonal family. Also, training the backbone on several temperatures simultaneously, with a single temperature-aware decoder. Ultimately, one would combine this with the previous idea to **create a decoder that is both timescale-aware and temperature-aware**.
- Simply use more expressive equivariant architectures, such as those recently introduced in [BKS⁺22] (**MACE**), or **equivariant transformer**. This could be combined with previous ideas. The **interplay** between Spherical Harmonics order (ℓ_{max}), network depth (L) and MACE-style self interaction order needs to be studied.
- Propose strategies to produce **interpretable networks**, or produce **interpretations of the network’s output**. In terms of interpretability, *equivariant* GNNs offer a better outlook than regular GNNs, and we already have some ideas about how to start thinking about this.

Since the learning curve is not saturated, **Self-supervised learning** is also an interesting playground. We can think of at least two directions for pre-training:

- **Denoising**: Adding nonphysical noise to the input positions and asking the network to denoise the input.
- **Predicting known quantities** (defined at time 0) such as E_{pot} , or quenched/thermal positions from thermal/quenched inputs.

²Mobility can be measured, and predicted, at any chosen time. Each time constitutes de-facto a new task, but these tasks are not independent.

5 Expected results

Using the testbed of glassy materials and **starting from our existing codebase**, the intern would **study equivariant networks**, then **develop new architectures** or **training curricula** to improve generalizability.

Following one or some of the scientific proposition above:

- Study rigorously how it impacts generalization power w.r.t. temperature (performance at the training temperature *and* when generalizing to unseen temperatures).
- Study rigorously how it impacts generalization power w.r.t. training data (compute the learning curve).
- When applicable, draft conclusions that apply more generally than the case of glassy materials (*i.e.* apply to geometric deep learning for science). For instance, adding non-local terms to the loss may be helpful to fight over-smoothing for a large range of GNN applications.
- Along the research, draft guidelines for building efficient equivariant GNNs. In CNNs, there are some well-known rule of thumb (like, one often combines blocks of two (3×3) kernels, with batch norm and skip connection in between such blocks). Such rules of thumb do not exist for equivariant networks. Drafting such rules could help build geometrical deep learning foundation models for some families of tasks.

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

6 Expected skills

We aim to recruit very motivated and talented students. The skills required are:

- Proficiency in python and pytorch
- Scientific rigor
- Some knowledge on GNNs (and possibly of pytorch-geometric)
- Some knowledge on Geometric Deep Learning (and possibly of e3nn or other package)
- Appeal for theory is appreciated (to understand the maths underlying equivariant networks)

The **topics** related and classes of the MVA are:

- **Representation learning:** Representation Learning for Computer Vision and Medical Imaging
- **Deep Learning:** Deep Learning, Deep learning in practice
- **Geometric Deep:** 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.
- [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.
- [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.
- [HLLZ⁺21] Michael J Hutchinson, Charline Le Lan, Sheheryar Zaidi, Emilien Dupont, Yee Whye Teh, and Hyunjik Kim. Lietransformer: Equivariant self-attention for lie groups. In *International Conference on Machine Learning*, pages 4533–4543. PMLR, 2021.
- [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.
- [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.
- [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.
- [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.
- [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.