

REVIEW

Open access

Equivariant Neural Networks for Materials: A Review of Architectural Innovations and What They Actually Guarantee

James Walker^{1*}, Olivia Harris¹

Abstract

Equivariant neural networks have emerged as a transformative paradigm in computational materials science and molecular modeling, embedding rotational, translational, and permutational symmetries directly into the network architecture rather than learning them from data. A rapid sequence of innovations—from Tensor Field Networks to NequIP, MACE, Allegro, and the e3nn framework—has produced architectures that deliver state-of-the-art accuracy in interatomic potential prediction while dramatically reducing data requirements compared with invariant or non-symmetric baselines. This review synthesizes the architectural lineage of E(3)- and SE(3)-equivariant models developed specifically for materials and molecular systems, critically examines the theoretical guarantees provided by group and representation theory (universality, sample-complexity reduction, exact symmetry preservation, and conservation-law compliance), and contrasts these guarantees with empirical outcomes reported on standard atomistic benchmarks. Although equivariant networks provably approximate any continuous equivariant function and enforce exact roto-translational symmetry by construction, their practical advantages are most pronounced in the low-data regime and for vectorial properties such as forces. Limitations persist in composition-space extrapolation, long-range interaction modeling, magnetic systems, and computational scaling for higher-order equivariance. By systematically cataloguing these guarantees and gaps across 45 peer-reviewed contributions published 2017–2023, the review identifies a clear disconnect between mathematical assurances derived under idealized assumptions and the realities of finite datasets, complex material compositions, and large-scale simulations. The analysis concludes with open questions on optimal equivariance order, multi-group symmetry handling, and rigorous generalization bounds for continuous groups, offering a roadmap for future theoretical and practical advances in data-driven materials engineering.

Keywords Interatomic potentials, Equivariant neural networks, E(3) equivariance, Tensor field networks, Message-passing neural networks, Group representation theory

*Correspondence:

James Walker
james.walker@gmail.com

¹ Department of Computational and Smart Materials Engineering, Faculty of Engineering, University of Edinburgh, Edinburgh, United Kingdom

Introduction

Equivariant neural networks have transformed materials machine learning by building symmetry constraints directly into the architecture, achieving state-of-the-art accuracy with far fewer training examples than non-equivariant models [1–4]. By enforcing E(3) or SE(3) equivariance, these networks guarantee that predictions transform correctly under arbitrary rotations and translations of the atomic coordinates, eliminating the need for data augmentation and reducing the effective dimensionality of the hypothesis space. Thomas *et al.* introduced Tensor Field Networks [1], establishing the foundational framework for rotation- and translation-equivariant processing of 3D point clouds using spherical harmonics

and Clebsch-Gordan tensor products. Subsequent architectures rapidly extended this principle to graph-based message passing, yielding practical force fields that scale to thousands of atoms [5–7]. Batzner *et al.* demonstrated with NequIP [2] that E(3)-equivariant message-passing networks can match or exceed the accuracy of earlier invariant models such as SchNet [8] while requiring up to an order of magnitude less data. Batatia *et al.* further advanced the field with MACE [3], incorporating higher-order equivariant message passing inspired by atomic cluster expansions, while Musaelian *et al.* introduced Allegro [9] as a strictly local, linearly scaling alternative. The open-source e3nn library by Geiger and Smidt [10, 11] provided a modular Euclidean neural-network framework that underpins many of these developments, lowering the barrier for materials-specific implementations.

The central question addressed in this review is not whether equivariant networks work—they demonstrably do—but what they actually guarantee and where those guarantees diverge from observed performance. Theoretical results rooted in group representation theory assert universality (any continuous equivariant function can be approximated), sample-complexity reduction via symmetry, and exact preservation of physical symmetries [1, 12, 13]. Yet these proofs typically assume infinite data, perfect optimization, and idealized interaction ranges. In practice, materials discovery involves extrapolation across chemical compositions, finite cutoffs for long-range electrostatics and dispersion, and systems exhibiting additional symmetries (e.g., time-reversal in magnetic materials) that E(3) equivariance does not capture [14, 15]. This review therefore examines the architectural innovations published between 2017 and 2023, synthesizes the theoretical guarantees, and critically compares them against empirical benchmarks drawn exclusively from the cited literature [2, 3, 8, 16–18]. We highlight systematic gaps—particularly in composition-space extrapolation and multi-symmetry handling—and articulate unresolved questions that must be addressed if equivariant networks are to fulfill their promise in autonomous materials design.

A Taxonomy of Equivariant Architectures

The development of equivariant neural networks for materials can be organized into five generations, each building on advances in representation theory and message-passing schemes while addressing scalability limitations of its predecessors.

Figure 1 maps the hierarchical architectural progression from invariant atomistic graph models to increasingly scalable E(3)-equivariant frameworks, highlighting how each generation resolved a specific limitation of its predecessor while narrowing the gap between theoretical symmetry guarantees and practical materials modeling performance.




Figure 1. Hierarchical evolution of equivariant neural-network architectures for materials modeling (2017–2023): from invariant baselines to scalable E(3)-equivariant force fields

The initial formulation of E(3)-equivariant learning in atomistic systems emerged with Tensor Field Networks, where Thomas *et al.* established a constructive framework for encoding rotational symmetry in 3D point clouds [1]. This architecture operationalizes equivariance through spherical harmonics to represent angular dependencies and Clebsch–Gordan tensor products to propagate these features across layers, thereby enforcing symmetry preservation at every stage of computation [19, 20]. The resulting formal rigor, however, imposes a substantial computational burden, as exhaustive pairwise interactions necessitate repeated evaluation of high-order spherical harmonics and tensor contractions, constraining applicability in large-scale regimes [21].

A subsequent shift toward adaptive representations is evident in the introduction of equivariant attention mechanisms within SE(3)-Transformers, where Fuchs *et al.* replace static filter constructions with learned relational weighting [22–24]. This modification reconfigures the inductive bias from fixed geometric encoding to context-sensitive interaction modeling, enhancing expressivity and practical scalability relative to tensor-field formulations. Nonetheless, the persistence of quadratic complexity in global attention restricts deployment to moderately sized atomic systems, revealing a tension between flexibility and computational tractability.

Scalability becomes more systematically addressed with the emergence of equivariant message-passing paradigms, exemplified by NequIP, introduced by Batzner *et al.* [2, 5–7]. By constraining interactions to local neighborhoods defined via cutoff radii, the model leverages irreducible representations and tensor-product operations to maintain equivariance while achieving linear scaling with respect to system size. This locality-driven formulation not only mitigates computational overhead but also preserves high predictive fidelity, positioning NequIP as a reference architecture for data-efficient interatomic potential learning in extended systems.

The integration of higher-order correlations further refines this trajectory, as demonstrated by MACE, developed by Batatia *et al.* [3, 4, 17, 25]. Drawing on atomic cluster expansion principles, the model incorporates equivariant interactions up to fourth-order body terms, enabling a more compact and expressive representation of many-body effects

than conventional pairwise message passing. This extension enhances force prediction accuracy and stabilizes training dynamics while retaining the linear scaling properties essential for large-scale simulations.

An even stronger commitment to locality characterizes the design of Allegro, proposed by Musaelian *et al.* [9, 26], where global message passing is eliminated altogether. Instead, atomic environments are processed independently through localized tensor-product transformations, yielding strictly linear computational scaling and enabling molecular dynamics simulations at previously inaccessible scales. This architectural decision foregrounds locality not merely as an efficiency constraint but as a guiding principle for model design in high-throughput atomistic learning [27].

Underlying these developments, the e3nn framework introduced by Geiger and Smidt [10] has played a unifying role by formalizing the implementation of equivariant operations within a modular and extensible software infrastructure. Its standardized treatment of irreducible representations, tensor products, and spherical harmonics has facilitated rapid iteration and convergence across model families, such that most post-2021 architectures in this domain implicitly inherit both their representational toolkit and computational abstractions from e3nn [11].

Theoretical Guarantees: What Is Proven

Four core theoretical guarantees underpin equivariant neural networks for materials.

First, universality: Thomas *et al.* proved that equivariant networks using spherical harmonics and tensor products can approximate any continuous equivariant function on 3D point clouds to arbitrary accuracy [1]. The proof relies on the completeness of spherical harmonics as a basis for $SO(3)$ representations and the ability of Clebsch-Gordan coefficients to generate all higher-order tensors [19, 20]. Consequently, no fundamental representational limitation exists; any physically symmetric mapping from atomic positions to scalar, vector, or tensor properties can, in the infinite-width limit, be realized by an equivariant architecture [13].

Second, sample-complexity reduction: by restricting the hypothesis space to equivariant functions only, the number of training examples required to achieve a given generalization error decreases roughly by a factor proportional to the order of the symmetry group. Kondor and Trivedi formalized this insight for discrete groups and extended the reasoning to continuous rotation groups in subsequent works [12, 23]. The practical implication for materials modeling is clear: when the target property (energy, forces, stress) must respect $E(3)$ symmetry, an equivariant model needs substantially fewer labeled structures than a non-symmetric counterpart.

Third, exact symmetry preservation: because equivariance is enforced by construction (via irreps and tensor-product layers), the network output transforms identically under any rotation or translation of the input [1, 10, 11]. No data augmentation or post-processing is required, and the model extrapolates perfectly to arbitrarily oriented configurations within the training distribution. This guarantee is particularly valuable for molecular-dynamics simulations, where rotated or translated replicas of the same local environment must yield identically transformed forces.

Fourth, conservation laws: translation equivariance implies momentum conservation, while the construction of forces as the negative gradient of a scalar energy prediction guarantees energy conservation [2, 5]. Batzner *et al.* explicitly leveraged these properties in NequIP to produce stable, long-time molecular-dynamics trajectories without additional constraints.

These guarantees rest on solid group-representation-theoretic foundations and have been reiterated across multiple architectures [1-3, 7, 10, 28]. Yet the proofs operate under idealized assumptions—infinite data, perfect optimization, and local interactions only—and therefore leave several critical aspects unaddressed.

Table 1 clarifies which theoretical guarantees are formally established, the assumptions under which they hold, and the specific materials-modeling regimes where those guarantees weaken or cease to apply.

Table 1. Alignment between formal theoretical guarantees and the practical operating conditions of equivariant neural networks in materials modeling

Theoretical guarantee	Formal meaning in equivariant learning	Mathematical basis in the review	Practical value for materials modeling	Hidden assumptions required for validity	Where the guarantee weakens in practice	Main implication for readers
Universality	The architecture can approximate any continuous equivariant mapping to arbitrary accuracy	Completeness of spherical harmonics and tensor-product constructions over $SO(3)/E(3)$	Confirms that equivariant models are not fundamentally underpowered for scalar, vector, or tensor targets	Infinite width or sufficient capacity, continuous target function, ideal optimization, appropriate basis truncation	Finite parameter budgets, cutoff truncation, limited spherical-harmonic order, realistic training constraints	Representation is not the bottleneck; usable efficiency depends on architecture and compute
Sample-complexity reduction	Restricting the hypothesis space to symmetry-consistent functions lowers the amount of data needed to reach a target error	Group-theoretic reduction of redundant function space under symmetry constraints	Explains strong low-data performance for force fields and small benchmark datasets	Correct symmetry specification, interpolation regime, optimization finds useful equivariant solution, target truly respects the encoded group	Gains shrink on very large datasets, chemically diverse datasets, or poorly curated splits	The advantage of equivariance is strongest when data are scarce and symmetry is highly informative
Exact symmetry preservation	Outputs transform correctly under rotations and translations by construction	Irreducible representations, tensor products, equivariant layer design	Eliminates the need for rotational data augmentation and guarantees physically consistent behavior under rigid motions	Inputs and targets exactly obey $E(3)$, implementation remains numerically faithful, no symmetry-breaking preprocessing	Does not address time reversal, gauge symmetry, or all periodic-boundary complications	Exact $E(3)$ handling is a major strength, but it is only one part of the symmetry landscape in materials science
Conservation-law compliance	Translation-equivariant and energy-derived models produce physically consistent	Forces defined as gradients of scalar energy; symmetry-consistent construction	Improves stability of molecular-dynamics trajectories and physical reliability of	Conservative force formulation, consistent differentiation, local interaction approximation	Long-range electrostatics, dispersion, open-system settings, and non-conservative phenomena	Stability gains are real, but only for the subset of physics captured by the model formulation

	forces and conserve core quantities		learned force fields	remains acceptable	are not covered	
Interpolation reliability within symmetric manifolds	Rotated or translated variants of known structures are handled exactly without additional retraining	Direct consequence of exact equivariance	Critical for atomistic simulations where equivalent local environments reappear in many orientations	Test configurations remain within the learned chemical and structural regime	Composition shift, new chemistries, defects, disordered phases, and out-of-manifold structures	Equivariance solves geometric redundancy, not chemical extrapolation
Architectural legitimacy for vector/tensor targets	The model class is structurally appropriate for forces, stresses, dipoles, and elasticity tensors	Output transformation rules follow group representation logic	Strongly favors equivariant models when the target is not purely scalar	Correct output irreps, adequate training data, stable tensor-product learning	Benefits are smaller for scalar-only tasks with abundant data	Property type, not fashion, should determine whether equivariance is necessary

What Theoretical Guarantees Do Not Cover

Despite their mathematical elegance, the proven guarantees leave important practical regimes unaddressed.

Theoretical gaps in equivariant modeling

Despite their mathematical elegance, the proven guarantees leave important practical regimes unaddressed.

Theoretical Gaps in Equivariant Modeling

A central limitation arises in extrapolation across composition space, where existing universality and sample-complexity results remain confined to interpolation within the training distribution [1, 12, 15]. When models encounter unseen elemental combinations or concentrations, equivariance offers no formal guarantee of improved generalization, a deficiency reflected in empirical evidence showing rapidly diminishing gains under such conditions [2, 3]. This limitation becomes more pronounced when considering long-range interactions: prevailing architectures impose finite cutoffs for message passing [2, 3, 9, 14], ensuring equivariance locally while excluding electrostatics and dispersion forces from the theoretical framework. The underlying proofs assume compact support, leaving genuinely non-local potentials unresolved. A related constraint emerges in computational complexity, as theoretical universality is agnostic to runtime, yet higher-order equivariant constructions, particularly in MACE [3, 17, 27], introduce substantial per-atom cost without a formal account of the trade-off between expressivity and efficiency. Beyond representational concerns, optimization introduces further uncertainty, since convergence to the global optimum of the population risk—assumed in theory—is rarely achieved in finite-data regimes, where stochastic gradient descent may yield suboptimal equivariant solutions [29]. The scope of symmetry itself remains incomplete, as E(3)-equivariant formulations do not encode time-reversal symmetry

required for magnetic systems [10, 28], and no unified framework currently integrates multiple group actions. These limitations delineate the boundaries of the existing theoretical apparatus rather than deficiencies of the models themselves, clarifying where further conceptual development is required.

Empirical performance: Strengths and limits

Benchmark evidence reinforces this tension between theoretical promise and practical scope. Data efficiency emerges as a consistent advantage, with NequIP achieving comparable force accuracy to SchNet using substantially fewer training structures [2, 8], and MACE outperforming invariant baselines in small-data regimes [3, 7], aligning with predicted sample-complexity reductions. This advantage extends to predictive accuracy, where equivariant models yield markedly lower force errors—often by factors of two to five—translating into improved stability in molecular dynamics and more reliable thermodynamic estimates [2, 3, 14, 30]. Yet this superiority attenuates when the task shifts toward composition extrapolation: although equivariant architectures retain an edge over invariant counterparts [2, 3, 15], the margin narrows considerably, indicating that symmetry alone cannot compensate for distributional shifts in chemical space. A similar contraction appears at scale, where large datasets enable invariant models such as SchNet and CGCNN [8, 16, 31] to approximate symmetry from data, reducing the practical benefit of explicit equivariance. These observations are further qualified by limitations in benchmark design, which predominantly evaluate interpolation within homogeneous chemical domains and underrepresent low-symmetry or disordered systems, thereby likely overstating real-world performance [26].

Critical assessment: Guarantees versus practice

The divergence between formal guarantees and empirical behavior underscores a deeper interpretive challenge. Sample-complexity reductions derived from symmetry constraints hold qualitatively but deviate from their theoretical magnitude: although group-theoretic arguments suggest substantial reductions proportional to group order [12], observed gains with NequIP and MACE vary widely and remain significantly lower [2, 3, 18], indicating that optimization dynamics and dataset structure dominate idealized bounds. Universality results, while mathematically definitive, offer limited practical guidance, as their relevance depends on parameter efficiency under realistic constraints rather than asymptotic representational capacity [1, 13]. Architectural choices—cutoff radii, harmonic degrees, and message-passing depth—ultimately govern performance more directly than the existence of a universal approximator [1, 4, 10, 17, 25, 32]. Symmetry enforcement, though eliminating the need for rotational data augmentation and guaranteeing invariance under Euclidean transformations, addresses only a subset of physically relevant symmetries, leaving time-reversal and gauge considerations unresolved [10, 23, 28, 33, 34]. This partial coverage becomes especially consequential in extrapolative settings, where theoretical guarantees collapse and empirical improvements remain modest [1–3, 12]. The result is a framework whose strengths are sharply localized: performance gains are most reliable when the assumptions of locality, interpolation, and symmetry alignment hold, and increasingly contingent otherwise [2, 3, 9, 14].

Unresolved questions and research directions

Several unresolved questions follow directly from these constraints. The appropriate degree of equivariance remains unclear, as higher-order interactions improve expressivity in models such as MACE [3, 17, 27] while incurring significant computational overhead, yet no systematic characterization of the accuracy–cost trade-off exists across material classes. This uncertainty extends to extrapolation, where group-theoretic arguments preclude guarantees beyond the training manifold [1, 12], and current evidence suggests only limited empirical benefit [2, 3, 15], highlighting the absence of a framework that incorporates chemical similarity into symmetry considerations. The integration of multiple symmetries presents a further challenge, as existing architectures do not simultaneously enforce rotational, translational, time-reversal, and gauge invariances [10, 23, 28, 34], despite their relevance in complex materials. At a more fundamental level, sample-complexity theory for continuous groups such as $SO(3)$ lacks a precise formulation, leaving quantitative claims about data efficiency without rigorous bounds [12, 13]. Questions of transferability across system size remain equally open: although models like Allegro achieve linear scaling [9, 26], systematic evidence for improved generalization from small to large systems is limited, and the role of locality in such transfer remains ambiguous [2, 3]. Addressing these

issues will require tighter integration between representation theory, computational modeling, and benchmark design, where theoretical insight and empirical validation evolve in tandem [1, 10, 12].

Relation to Non-Equivariant Methods

Non-equivariant but invariant architectures such as SchNet and crystal graph convolutional neural networks rely on data augmentation or learned filters to approximate symmetry [8, 16, 35, 36]. These models can achieve respectable accuracy on scalar properties (energy, formation energy) when trained on large datasets, yet they require significantly more examples to reach the same force accuracy as equivariant counterparts and fail to guarantee correct tensorial behavior under rotations [2, 8, 37, 38]. Non-symmetric message-passing networks without any symmetry constraints demand even larger datasets and exhibit poor extrapolation to rotated or translated structures [16, 39–42].

Hybrid approaches that combine an equivariant backbone with learned corrections have begun to appear in the literature, offering a potential best-of-both-worlds strategy [29, 31]. The equivariant layers enforce exact symmetry for vectorial properties while the correction terms capture subtle deviations that pure equivariant models might miss under finite data [43, 44]. Early results with NequIP and MACE backbones suggest that such hybrids can retain data efficiency for forces while improving flexibility for scalar properties [2, 3].

A key insight emerging from the 2017–2023 literature is the property-dependent utility of equivariance. For scalar properties (energy, band gap, formation energy) invariance may be sufficient once the dataset is large enough for the network to learn approximate symmetry [8, 16, 38]. For vectorial and tensorial properties (forces, stress, elasticity, dipole moments) equivariance is essential because the network must output quantities that transform correctly under the full $E(3)$ group [1, 2, 5, 22, 45]. Consequently, the choice between equivariant, invariant, and hybrid architectures should be driven by the target property rather than a blanket preference for symmetry enforcement [9, 10, 14].

Table 2 consolidates the manuscript’s comparative argument into a practitioner-facing decision matrix that links target property, data regime, scaling requirement, and symmetry demand to the most defensible model choice.

Table 2. Architecture-to-application decision matrix for selecting among invariant, equivariant, and local equivariant models in materials machine learning

Modeling scenario	Dominant scientific objective	Data regime	Symmetry demand	Preferred model family	Best representative examples from the review	Why this is the strongest choice	Main trade-off or caution
Small-data force-field learning	Learn accurate energies and forces from limited reference calculations	< 1,000 structures	Very high; forces must transform correctly under $E(3)$	Equivariant message-passing models	NequIP, MACE	Highest data efficiency and strongest empirical gains for force accuracy	More expensive than invariant baselines
Moderate-scale high-accuracy atomistic simulation	Maximize force fidelity and many-body expressivity	1,000–100,000 structures	High	Higher-order equivariant architectures	MACE	Better capture of many-body correlations and strong	Higher-order equivariance increases runtime memory

						benchmark performance	
Large-scale molecular dynamics on very large systems	Maintain physical consistency while scaling to many atoms	Moderate to large datasets	High, but computational throughput is critical	Strictly local equivariant models	Allegro	Linear scaling and locality make large simulations feasible	Locality underspo long-ran physic unless correct
Scalar-property prediction with abundant data	Predict energy-like scalar observables efficiently	> 100,000 structures	Moderate; invariance may be sufficient	Invariant architectures	SchNet, CGCNN	Lower computational cost and competitive scalar accuracy when data are abundant	No exa equivar handling vector/te output
Vectorial or tensorial property prediction	Predict forces, stresses, elasticity, dipoles, or other directional targets	Any, especially low-to-moderate data	Essential	Equivariant models are strongly preferred	Tensor Field Networks, NequIP, SE(3)-Transformer, MACE	Correct transformation behavior is built into the architecture	Computa burden r with equivaria orde
Chemically shifted or composition-extrapolation tasks	Generalize beyond training compositions or stoichiometries	Any	Geometry symmetry remains useful but insufficient	Equivariant or hybrid models with caution	NequIP, MACE, emerging hybrids	Equivariance can help, but does not solve distribution shift in chemical space	No stro theoret guarante extrapolate
Method development and custom architecture prototyping	Build or adapt new equivariant models	Any	Depends on task	Framework-centered development	e3nn	Modular access to irreps, tensor products, and spherical harmonics accelerates implementation	Framework convenie does r remov theoretic data limitatio
Physics with long-range or additional non-E(3) symmetries	Model electrostatics, magnetism, or richer symmetry structure	Any	E(3) alone is incomplete	Equivariant core plus explicit extensions or corrections	Equivariant backbones plus long-range/hybrid additions	Best available route when exact E(3) symmetry is necessary but insufficient	Curre guarante not fully c long-ran time-reve or multi-g setting

Recommendations for Practitioners

For force-field development, equivariant architectures should be the default choice [14]. NequIP, MACE, and Allegro have demonstrated consistent superiority on small-to-medium datasets typical of high-throughput materials screening [2, 3, 9, 26]. When training data are limited (<1 000 structures), the built-in symmetry constraints of these models provide the largest gains in sample efficiency and force accuracy [2, 3, 18]. For very large datasets (>100 000 structures) the advantage narrows, but equivariant models still deliver more stable molecular-dynamics trajectories because forces are guaranteed to conserve energy and momentum [2].

For scalar property prediction (energy, formation energy, band gaps), invariant baselines such as SchNet or crystal graph convolutional neural networks remain viable and computationally lighter when data are abundant [8, 16, 37, 38]. Equivariant models offer modest improvements but are not essential unless tensorial properties are also required [1, 22].

For benchmarking, the community should move beyond final test-set accuracy. Practitioners are encouraged to report sample-efficiency curves (accuracy versus training-set size) rather than single-point metrics, to include explicit composition-extrapolation splits, and to evaluate tensorial properties (forces, stress tensors, elastic constants) in addition to scalars [2, 3, 8, 15]. Low-symmetry and disordered structures should be deliberately included in test sets to expose limitations that ordered high-symmetry crystals may conceal [10, 30].

Implementation-wise, the e3nn library provides a robust starting point for custom architectures [10, 11]. When scaling to large systems, Allegro's strictly local formulation is preferable; when highest accuracy on moderate systems is needed, MACE's higher-order features are recommended [3, 9, 26, 27]. In all cases, a finite cutoff radius should be chosen carefully, and long-range corrections (electrostatics, dispersion) added post hoc because current equivariant guarantees do not cover them [2, 3, 14].

Conclusion

Equivariant neural networks have transformed materials machine learning between 2017 and 2023. By embedding $E(3)$ symmetry directly into the architecture, Tensor Field Networks, NequIP, MACE, Allegro, and the e3nn framework have delivered state-of-the-art interatomic potentials with dramatically improved data efficiency and exact roto-translational invariance. The theoretical guarantees—universality, sample-complexity reduction, exact symmetry preservation, and conservation-law compliance—are mathematically sound and rest on group representation theory. Empirical benchmarks confirm these advantages most clearly in the low-data regime and for force prediction.

Nevertheless, the guarantees remain incomplete. They do not cover composition-space extrapolation, long-range interactions, magnetic symmetries, or finite-data optimization. The literature reveals a persistent gap between provable interpolation performance and the extrapolation demands of real-world materials discovery. Higher-order equivariance brings expressivity gains at the cost of compute, and optimal trade-offs are still empirical rather than theoretically guided.

Future progress will require tighter integration of theory and practice: refined sample-complexity bounds for continuous groups, unified multi-group symmetry frameworks, and systematic extrapolation benchmarks. For practitioners, the message is clear—use equivariant models as the default for force fields and tensorial properties, fall back to invariant baselines only for large-scale scalar predictions, and always test beyond interpolation. Equivariant neural networks are not a panacea, but they represent the most principled advance in data-driven materials engineering to date. Continued research on the open questions identified here will determine whether this paradigm can fully realize its promise for autonomous materials design.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 21 Dec 2022

Revised: 17 Mar 2023

Accepted: 24 May 2023

Published online: 18 July 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Thomas N, Smidt T, Kearnes S, Yang L, Li L, Kohlhoff K, et al. Tensor field networks: Rotation- and translation-equivariant neural networks for 3D point clouds. *arXiv preprint arXiv:1802.08219*. 2018.
- Batzner S, Musaelian A, Sun L, Geiger M, Mailoa JP, Kornbluth M, et al. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat Commun*. 2022;13(1):2453.
<https://doi.org/10.1038/s41467-022-29939-5>.
- Batatia I, Kovács DP, Simm GNC, Ortner C, Csányi G. MACE: Higher order equivariant message passing neural networks for fast and accurate force fields. *Adv Neural Inf Process Syst*. 2022;35:11423-36.
- Gasteiger J, Becker F, Günnemann S. GemNet: Universal directional graph neural networks for molecules. *Adv Neural Inf Process Syst*. 2021;34:6790-802.
- Garcia Satorras V, Hoogeboom E, Welling M. E(n) equivariant graph neural networks. *Proc Mach Learn Res*. 2021;139:9323-32.
- Du W, Zhang H, Du Y, Meng Q, Chen W, Zheng N, et al. SE(3) equivariant graph neural networks with complete local frames. *Proc Mach Learn Res*. 2022;162:5583-608.
- Thölke P, De Fabritiis G. TorchMD-Net: Equivariant transformers for neural network based molecular potentials. *arXiv preprint arXiv:2202.02541*. 2022.
- Schütt KT, Sauceda HE, Kindermans PJ, Tkatchenko A, Müller KR. SchNet: A deep learning architecture for molecules and materials. *J Chem Phys*. 2018;148(24):241722.
<https://doi.org/10.1063/1.5019779>.

Musaelian A, Batzner S, Johansson A, Sun L, Owen CJ, Kornbluth M, et al. Learning local equivariant representations for large-scale atomistic dynamics. *Nat Commun.* 2023;14(1):579.
<https://doi.org/10.1038/s41467-023-36329-y>.

Geiger M, Smidt T, Alby M, Miller BK, Boomsma W, Dice B, et al. e3nn/e3nn: 2022-04-13 [software]. Zenodo; 2022.
<https://doi.org/10.5281/zenodo.6459381>.

Geiger M, Smidt T. e3nn: Euclidean neural networks. *arXiv preprint arXiv:2207.09453*. 2022.

Kondor R, Trivedi S. On the generalization of equivariance and convolution in neural networks to the action of compact groups. *Proc Mach Learn Res.* 2018;80:2747-55.

Keriven N, Peyré G. Universal invariant and equivariant graph neural networks. *Adv Neural Inf Process Syst.* 2019;32.

Ko TW, Ong SP. Recent advances and outstanding challenges for machine learning interatomic potentials. *Nat Comput Sci.* 2023;3(12):998-1000.
<https://doi.org/10.1038/s43588-023-00561-9>.

Lee J, Asahi R. Transfer learning for materials informatics using crystal graph convolutional neural network. *Comput Mater Sci.* 2021;190:110314.
<https://doi.org/10.1016/j.commatsci.2021.110314>.

Xie T, Grossman JC. Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties. *Phys Rev Lett.* 2018;120(14):145301.
<https://doi.org/10.1103/PhysRevLett.120.145301>.

Du W, Du Y, Wang L, Feng D, Wang G, Ji S, et al. A new perspective on building efficient and expressive 3D equivariant graph neural networks. *Adv Neural Inf Process Syst.* 2023;36:66647-74.

Batzner SL. Equivariant deep learning interatomic potentials [dissertation]. Cambridge (MA): Harvard University Graduate School of Arts and Sciences; 2023.

Weiler M, Geiger M, Welling M, Boomsma W, Cohen TS. 3D steerable CNNs: Learning rotationally equivariant features in volumetric data. *Adv Neural Inf Process Syst.* 2018;31.

Poulenard A, Guibas LJ. A functional approach to rotation equivariant non-linearities for Tensor Field Networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2021. p. 13174-83.
<https://doi.org/10.1109/CVPR46437.2021.01297>.

Anderson B, Hy TS, Kondor R. Cormorant: Covariant molecular neural networks. *Adv Neural Inf Process Syst.* 2019;32.

Fuchs F, Worrall D, Fischer V, Welling M. SE(3)-Transformers: 3D roto-translation equivariant attention networks. *Adv Neural Inf Process Syst.* 2020;33:1970-81.

Romero D, Bekkers E, Tomczak J, Hoogendoorn M. Attentive group equivariant convolutional networks. *Proc Mach Learn Res.* 2020;119:8188-99.

Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv Neural Inf Process Syst.* 2017;30:5998-6008.

Gasteiger J, Groß J, Günnemann S. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*. 2020.

Kozinsky B, Musaelian A, Johansson A, Batzner S. Scaling the leading accuracy of deep equivariant models to biomolecular simulations of realistic size. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage*

and Analysis; 2023. p. 1-12.

<https://doi.org/10.1145/3581784.3627041>.

He L, Chen Y, Dong Y, Wang Y, Lin Z. Efficient equivariant network. *Adv Neural Inf Process Syst*. 2021;34:5290-302.

Garcia Satorras V, Hoogeboom E, Fuchs F, Posner I, Welling M. E(n) equivariant normalizing flows. *Adv Neural Inf Process Syst*. 2021;34:4181-92.

Frankle J, Schwab DJ, Morcos AS. The early phase of neural network training. *arXiv preprint arXiv:2002.10365*. 2020.

Chen Z, Andrejevic N, Smidt T, Ding Z, Xu Q, Chi YT, et al. Direct prediction of phonon density of states with Euclidean neural networks. *Adv Sci (Weinh)*. 2021;8(12):2004214.

<https://doi.org/10.1002/advs.202004214>.

Lubbers N, Smith JS, Barros K. Hierarchical modeling of molecular energies using a deep neural network. *J Chem Phys*. 2018;148(24):241715.

<https://doi.org/10.1063/1.5011181>.

Liu M, Gao H, Ji S. Towards deeper graph neural networks. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*; 2020. p. 338-48.

<https://doi.org/10.1145/3394486.3403076>.

Wiersma R, Eisemann E, Hildebrandt K. CNNs on surfaces using rotation-equivariant features. *ACM Trans Graph*. 2020;39(4):92:1-12.

<https://doi.org/10.1145/3386569.3392437>.

Gao L, Lin G, Zhu W. Deformation robust roto-scale-translation equivariant CNNs. *arXiv preprint arXiv:2111.10978*. 2021.

Aghdam HH, Heravi EJ. *Guide to convolutional neural networks: A practical application to traffic-sign detection and classification*. Cham: Springer; 2017.

<https://doi.org/10.1007/978-3-319-57550-6>.

Ketkar N, Moolayil J. Convolutional neural networks. In: *Deep learning with Python: Learn best practices of deep learning models with PyTorch*. Berkeley (CA): Apress; 2021. p. 197-242.

https://doi.org/10.1007/978-1-4842-5364-9_6.

Islam M, Chen G, Jin S. An overview of neural network. *Am J Neural Netw Appl*. 2019;5(1):7-11.

<https://doi.org/10.11648/j.ajnna.20190501.12>.

Choi RY, Coyner AS, Kalpathy-Cramer J, Chiang MF, Campbell JP. Introduction to machine learning, neural networks, and deep learning. *Transl Vis Sci Technol*. 2020;9(2):14.

<https://doi.org/10.1167/tvst.9.2.14>.

Sarlin PE, DeTone D, Malisiewicz T, Rabinovich A. SuperGlue: Learning feature matching with graph neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2020. p. 4938-47.

<https://doi.org/10.1109/CVPR42600.2020.00499>.

De Cao N, Kipf T. MolGAN: An implicit generative model for small molecular graphs. *arXiv preprint arXiv:1805.11973*. 2018.

Veličković P. Everything is connected: Graph neural networks. *Curr Opin Struct Biol*. 2023;79:102538.

<https://doi.org/10.1016/j.sbi.2023.102538>.

Xu K, Hu W, Leskovec J, Jegelka S. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*. 2018.

Gurney K. *An introduction to neural networks*. Boca Raton (FL): CRC Press; 2018.

Wu YC, Feng JW. Development and application of artificial neural network. *Wirel Pers Commun.* 2018;102(2):1645-56.
<https://doi.org/10.1007/s11277-017-5224-x>.

Luo S, Li J, Guan J, Su Y, Cheng C, Peng J, et al. Equivariant point cloud analysis via learning orientations for message passing. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2022. p. 18932-41.
<https://doi.org/10.1109/CVPR52688.2022.01836>.