Novelty Assessment Report

Paper: Reducing Symmetry Increase in Equivariant Neural Networks

PDF URL: https://openreview.net/pdf?id=dUgq4bLY4X

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

Equivariant Neural Networks (ENNs) have empowered numerous applications in scientific fields. Despite their remarkable capacity for representing geometric structures, ENNs suffer from degraded expressivity when processing symmetric inputs: the output representations are invariant to transformations that extend beyond the input's symmetries. The mathematical essence of this phenomenon is that a symmetric input, after being processed by an equivariant map, experiences an increase in symmetry. While prior research has documented symmetry increase in specific cases, a rigorous understanding of its underlying causes and general reduction strategies remains lacking. In this paper, we provide a detailed and in-depth characterization of symmetry increase together with a principled framework for its reduction: (i) For any given feature space and input symmetry group, we prove that the increased symmetry admits an infimum determined by the structure of the feature space; (ii) Building on this foundation, we develop a computable algorithm to derive this infimum, and propose practical guidelines for feature design to prevent harmful symmetry increases. (iii) Under standard regularity assumptions, we demonstrate that for most equivariant maps, our guidelines effectively reduce symmetry increase. To complement our theoretical findings, we provide visualizations and experiments on both synthetic datasets and the real-world QM9 dataset. The results validate our theoretical predictions.

Core Task Landscape

This paper addresses: **Reducing Symmetry Increase in Equivariant Neural Networks**

A total of **13 papers** were analyzed and organized into a taxonomy with **7 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Theoretical Foundations and Characterization of Symmetry Phenomena**
- **Symmetry Breaking Methods and Frameworks**
- **Computational Efficiency and Architectural Optimization**
- **Specialized Equivariant Architectures and Applications**

Complete Taxonomy Tree

- Reducing Symmetry Increase in Equivariant Neural Networks Survey Taxonomy
- Theoretical Foundations and Characterization of Symmetry Phenomena
  - Symmetry Increase Analysis and Reduction Frameworks ★ (1 papers)
  - [0] Reducing Symmetry Increase in Equivariant Neural Networks (Anon et al., 2026) View paper
  - Equivariance Relaxation and Subgroup Constraints (2 papers)
  - [7] Relaxed Equivariant Graph Neural Networks (Hofgard, 2024) View paper
  - [12] Equivariance-aware Architectural Optimization of Neural Networks (Maile, 2022) View paper
- Symmetry Breaking Methods and Frameworks
  - Explicit Symmetry Breaking Mechanisms (3 papers)
  - [3] Equivariant symmetry breaking sets (Xie Yu-qing, 2024) View paper
  - [9] Any-Subgroup Equivariant Networks via Symmetry Breaking (A Goel, 2025) View paper
  - [10] Breaking the Symmetry: Resolving Symmetry Ambiguities in Equivariant Neural Networks (Balachandar, 2022) View paper
  - Probabilistic and Training-Based Symmetry Breaking (2 papers)
  - [5] Improving equivariant networks with probabilistic symmetry breaking (Lawrence, 2025) View paper
  - [6] Relaxed Equivariance via Multitask Learning (Elhag, 2024) View paper
- Computational Efficiency and Architectural Optimization
  - Dimensional Reduction and Convolution Optimization (3 papers)
  - [1] Reducing SO (3) convolutions to SO (2) for efficient equivariant GNNs (Passaro, 2023) View paper
  - [2] Efficient equivariant network (Lingshen He, 2021) View paper
  - [11] Exploiting Learned Symmetries in Group Equivariant Convolutions (Attila Lengyel, 2021) View paper
  - Attention and Learned Symmetry Exploitation (1 papers)
  - [13] Co-Attentive Equivariant Neural Networks: Focusing Equivariance On Transformations Co-Occurring In Data (David W. Romero GuzmÃ¡n, 2019) View paper
- Specialized Equivariant Architectures and Applications
  - Wavelet-Based and Scale-Invariant Architectures (2 papers)
  - [4] Global Control for Local SO(3)-Equivariant Scale-Invariant Vessel Segmentation (Patryk Rygiel, 2024) View paper
  - [8] Equivariant Wavelets: Fast Rotation and Translation Invariant Wavelet Scattering Transforms (Andrew K Saydjari, 2021) View paper

Narrative

Core task: Reducing symmetry increase in equivariant neural networks. Equivariant architectures are designed to respect known symmetries in data, but a subtle challenge arises when network layers inadvertently increase symmetry beyond what the problem requires, leading to reduced expressiveness or ambiguous representations. The field has organized around several complementary directions. Theoretical foundations examine how and why symmetry increase occurs, characterizing the mathematical conditions under which equivariant operations preserve or expand symmetry groups. Symmetry breaking methods develop explicit techniques—ranging from learnable breaking mechanisms to probabilistic approaches—that allow networks to selectively reduce unwanted symmetries while maintaining beneficial equivariances. Computational efficiency and architectural optimization focus on designing layers and operations that balance equivariance constraints with practical scalability, often leveraging sparse representations or efficient group convolutions. Specialized equivariant architectures tailor these principles to specific domains such as molecular modeling, image analysis, or graph learning, where particular symmetry groups (e.g., SO(3), SE(3)) dominate.

Recent work has explored diverse strategies for managing symmetry. Some studies introduce relaxed equivariance frameworks that soften strict constraints to improve flexibility, as seen in Relaxed Equivariant GNNs[7] and Relaxed Equivariance Multitask[6], while others propose explicit symmetry breaking sets or probabilistic mechanisms (Symmetry Breaking Sets[3], Probabilistic Symmetry Breaking[5]) to disambiguate representations. Domain-specific innovations like SO3 to SO2[1] and SO3 Vessel Segmentation[4] demonstrate how reducing from higher to lower symmetry groups can enhance task performance. The original paper, Reducing Symmetry Increase[0], sits within the theoretical and analytical branch, focusing on formal frameworks for understanding and mitigating symmetry increase. Its emphasis on rigorous characterization contrasts with more application-driven works like SO3 Vessel Segmentation[4], yet complements method-oriented papers such as Symmetry Breaking Sets[3] by providing foundational insights that guide when and how symmetry reduction should be applied.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

Both subtopics address the challenge of overly restrictive symmetry constraints in equivariant neural networks, but from different angles. The original leaf focuses on analyzing and characterizing the symmetry increase phenomenon itself—understanding when and why networks gain unwanted symmetries—and developing formal frameworks to reduce this increase. The sibling subtopic takes a more architectural approach by relaxing equivariance constraints to appropriate subgroups or optimizing which symmetries to enforce, rather than analyzing symmetry increase per se.

Similarities: - Both aim to improve equivariant network expressiveness by addressing overly restrictive symmetry constraints - Both require understanding the relationship between architectural choices and resulting symmetries - Both develop principled mathematical approaches rather than ad-hoc solutions - Both distinguish themselves from simple symmetry breaking methods that don't address the underlying constraint structure

Differences: - Original leaf focuses on characterizing and reducing symmetry increase (when networks gain more symmetries than intended), while sibling focuses on relaxing initial equivariance constraints to subgroups - Original leaf emphasizes analysis and formal characterization of phenomena, while sibling emphasizes architectural design and constraint optimization - Original leaf addresses unintended symmetry emergence, while sibling addresses intentional modification of symmetry requirements - Sibling subtopic involves algorithmic selection of appropriate symmetry groups, while original leaf develops frameworks to prevent symmetry amplification

Suggested Search Directions: - Papers analyzing the relationship between subgroup relaxation and symmetry increase—whether relaxing to subgroups can inadvertently cause symmetry increase - Methods that combine symmetry increase analysis with adaptive subgroup selection - Theoretical characterizations of when equivariance relaxation preserves desired symmetries without introducing unwanted ones

Sibling Subtopics

- **Equivariance Relaxation and Subgroup Constraints** (leaves: 1, papers: 2)
- Scope: Papers developing methods to relax equivariance constraints to subgroups or optimize architectural symmetry constraints algorithmically.
- Exclude: Excludes papers on breaking symmetry in outputs or training procedures; those belong in Symmetry Breaking Methods.

Contributions Analysis

Overall novelty summary. The paper contributes a rigorous mathematical framework for understanding and reducing symmetry increase in equivariant neural networks, proving the existence of a symmetry infimum determined by feature space structure and developing a computable algorithm to derive it. Within the taxonomy, it occupies the 'Symmetry Increase Analysis and Reduction Frameworks' leaf under 'Theoretical Foundations and Characterization of Symmetry Phenomena'. Notably, this leaf contains only the original paper itself—no sibling papers—indicating this is a relatively sparse and underexplored research direction focused specifically on formal characterization of symmetry increase phenomena.

The taxonomy reveals neighboring work in 'Equivariance Relaxation and Subgroup Constraints' (two papers on relaxing equivariance to subgroups) and 'Symmetry Breaking Methods' (five papers on explicit or probabilistic breaking mechanisms). While these adjacent directions address related challenges—managing unwanted symmetries or enabling lower-symmetry outputs—they differ fundamentally in approach. The original paper provides theoretical foundations for understanding why symmetry increases occur, whereas neighboring leaves focus on architectural mechanisms or training procedures to break symmetries. The taxonomy's scope notes clarify that formal characterization of increase belongs here, while breaking mechanisms without such characterization belong elsewhere.

Among seventeen candidates examined across three contributions, none were found to clearly refute any claimed novelty. The characterization of symmetry infimum examined three candidates with zero refutations; the computable algorithm examined four candidates with zero refutations; the theoretical guarantee examined ten candidates with zero refutations. This suggests that within the limited search scope of top-K semantic matches and citation expansion, no prior work appears to provide the same combination of formal symmetry increase characterization, infimum derivation, and reduction framework. The contributions addressing theoretical guarantees received the most scrutiny but still showed no overlapping prior work among examined candidates.

Based on the limited literature search of seventeen candidates, the work appears to occupy a novel position providing formal mathematical foundations for a phenomenon documented but not rigorously characterized in prior work. The absence of sibling papers in its taxonomy leaf and zero refutations across all contributions suggest substantive theoretical novelty, though this assessment is constrained by the search scope and does not constitute an exhaustive field survey.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Characterization of symmetry increase with a unique symmetry infimum

Description: The authors establish a mathematical foundation showing that for any feature space and input symmetry group, the symmetry increase phenomenon has a lower bound called the symmetry infimum. This infimum is uniquely determined by the algebraic structure of the feature space itself.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Tensor Algebra Toolkit for Folded Mixture Models: Symmetry-Aware Moments, Orbit-Space Estimation, and Poly-LAN Rates

URL: View paper

Brief Assessment

Folded Mixture Models[26] addresses symmetry in statistical mixture models with folding group actions on parameter spaces, not symmetry increase in equivariant neural network feature spaces or input-output transformations.

2. On the conformal walk dimension: quasisymmetric uniformization for symmetric diffusions

URL: View paper

Brief Assessment

Conformal Walk Dimension[25] addresses conformal walk dimension in symmetric diffusions on metric spaces, not symmetry increase in equivariant neural networks or feature spaces with input symmetry groups.

3. Symmetry-induced disentanglement on graphs

URL: View paper

Brief Assessment

Symmetry-induced Disentanglement[24] focuses on disentangled representations in graphs using Lie groups and symmetry decomposition for generative models. It does not address symmetry increase in equivariant neural networks or establish infimum bounds for feature spaces with input symmetry groups.

Contribution 2: Computable algorithm for deriving the symmetry infimum

Description: The authors propose an algorithm that computes the symmetry infimum through orbit type analysis. This algorithm enables practical guidelines for feature design that can predict and prevent harmful symmetry increases in equivariant neural networks.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Fourier-Orbit Construction of GKZ-Type Systems for Commutative Linear Algebraic Groups

URL: View paper

Brief Assessment

Fourier-Orbit GKZ[28] addresses GKZ-type D-modules for commutative algebraic groups and holonomicity analysis, not symmetry infimum computation for equivariant neural networks through orbit type analysis.

2. Tensor Algebra Toolkit for Folded Mixture Models: Symmetry-Aware Moments, Orbit-Space Estimation, and Poly-LAN Rates

URL: View paper

Brief Assessment

Folded Mixture Models[26] develops algorithms for orbit-space estimation in mixture models, not for computing symmetry infimum through orbit type analysis in equivariant neural networks.

3. The equivariant LS-category of polar actions

URL: View paper

Brief Assessment

Equivariant LS-category[29] focuses on topological invariants for polar actions on symmetric spaces, not on computing symmetry infimum through orbit type analysis for equivariant neural networks.

4. Symmetry and generalisation in machine learning

URL: View paper

Brief Assessment

Symmetry and Generalisation[27] focuses on generalisation theory in machine learning using averaging operators and orbit representatives, not on computing symmetry infimum through orbit type analysis for equivariant networks.

Contribution 3: Theoretical guarantee for symmetry reduction under regularity conditions

Description: The authors prove that under standard regularity assumptions like the manifold hypothesis, their framework effectively reduces symmetry increase for most equivariant maps. The output symmetry becomes exactly the predicted infimum, preventing loss of orientational information.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. General e (2)-equivariant steerable cnns

URL: View paper

Brief Assessment

E2 Steerable CNNs[16] focuses on constructing equivariant convolutions through kernel constraints for the Euclidean group E(2), not on analyzing symmetry increase phenomena or providing theoretical guarantees about symmetry reduction under manifold hypothesis conditions.

2. Equivariant motion manifold primitives

URL: View paper

Brief Assessment

Motion Manifold Primitives[17] focuses on equivariant motion primitives for robot trajectory learning, not on theoretical guarantees for symmetry reduction in equivariant neural networks under manifold hypothesis conditions.

3. Group equivariant subsampling

URL: View paper

Brief Assessment

Equivariant Subsampling[22] focuses on subsampling operations in neural networks to preserve group equivariance, not on proving theoretical guarantees about symmetry reduction under manifold hypothesis regularity conditions for general equivariant maps.

4. Dimensionality reduction in deep learning through group actions

URL: View paper

Brief Assessment

Group Actions Dimensionality[18] focuses on dimensionality reduction through group actions in deep learning, not on symmetry increase in equivariant neural networks or manifold hypothesis conditions for equivariant maps.

5. -Equivariant Dimensionality Reduction on Stiefel Manifolds

URL: View paper

Brief Assessment

Equivariant Dimensionality Reduction[14] focuses on dimensionality reduction for Stiefel/Grassmannian manifolds with O(k)-equivariance, not on symmetry increase in equivariant neural networks or manifold hypothesis regularity conditions for general equivariant maps.

6. Equivariant Verlinde formula from fivebranes and vortices

URL: View paper

Brief Assessment

Equivariant Verlinde Formula[23] focuses on complex Chern-Simons theory and fivebranes on Riemann surfaces, not on equivariant neural networks or manifold hypothesis regularity conditions for symmetry reduction in machine learning contexts.

7. G-invariant diffusion maps

URL: View paper

Brief Assessment

G-invariant Diffusion Maps[21] focuses on diffusion maps for group-invariant manifolds in data analysis, not on symmetry reduction in equivariant neural networks under manifold hypothesis conditions.

8. Structuring representations using group invariants

URL: View paper

Brief Assessment

Group Invariants Representations[20] focuses on learning equivariant representations through symmetry regularization using group invariants, not on reducing symmetry increase in equivariant neural networks under manifold hypothesis conditions. The paper addresses a different problem domain.

9. Shaping manifolds in equivariant recurrent neural networks

URL: View paper

Brief Assessment

Shaping Manifolds[15] focuses on equivariant recurrent neural networks for continuous attractor models in neuroscience, analyzing fixed-point manifolds through group representation theory. This is a fundamentally different domain from the original paper's work on equivariant neural networks for scientific applications and symmetry increase reduction.

10. Dimensionless machine learning: Imposing exact units equivariance

URL: View paper

Brief Assessment

Dimensionless Machine Learning[19] focuses on units equivariance through dimensional analysis and dimensionless features, not on symmetry reduction in equivariant maps under manifold hypothesis regularity conditions as described in the original contribution.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Reducing Symmetry Increase in Equivariant Neural Networks View paper
- [1] Reducing SO (3) convolutions to SO (2) for efficient equivariant GNNs View paper
- [2] Efficient equivariant network View paper
- [3] Equivariant symmetry breaking sets View paper
- [4] Global Control for Local SO(3)-Equivariant Scale-Invariant Vessel Segmentation View paper
- [5] Improving equivariant networks with probabilistic symmetry breaking View paper
- [6] Relaxed Equivariance via Multitask Learning View paper
- [7] Relaxed Equivariant Graph Neural Networks View paper
- [8] Equivariant Wavelets: Fast Rotation and Translation Invariant Wavelet Scattering Transforms View paper
- [9] Any-Subgroup Equivariant Networks via Symmetry Breaking View paper
- [10] Breaking the Symmetry: Resolving Symmetry Ambiguities in Equivariant Neural Networks View paper
- [11] Exploiting Learned Symmetries in Group Equivariant Convolutions View paper
- [12] Equivariance-aware Architectural Optimization of Neural Networks View paper
- [13] Co-Attentive Equivariant Neural Networks: Focusing Equivariance On Transformations Co-Occurring In Data View paper
- [14] -Equivariant Dimensionality Reduction on Stiefel Manifolds View paper
- [15] Shaping manifolds in equivariant recurrent neural networks View paper
- [16] General e (2)-equivariant steerable cnns View paper
- [17] Equivariant motion manifold primitives View paper
- [18] Dimensionality reduction in deep learning through group actions View paper
- [19] Dimensionless machine learning: Imposing exact units equivariance View paper
- [20] Structuring representations using group invariants View paper
- [21] G-invariant diffusion maps View paper
- [22] Group equivariant subsampling View paper

- [23] Equivariant Verlinde formula from fivebranes and vortices View paper
- [24] Symmetry-induced disentanglement on graphs View paper
- [25] On the conformal walk dimension: quasisymmetric uniformization for symmetric diffusions View paper
- [26] Tensor Algebra Toolkit for Folded Mixture Models: Symmetry-Aware Moments, Orbit-Space Estimation, and Poly-LAN Rates View paper
- [27] Symmetry and generalisation in machine learning View paper
- [28] Fourier-Orbit Construction of GKZ-Type Systems for Commutative Linear Algebraic Groups View paper
- [29] The equivariant LS-category of polar actions View paper