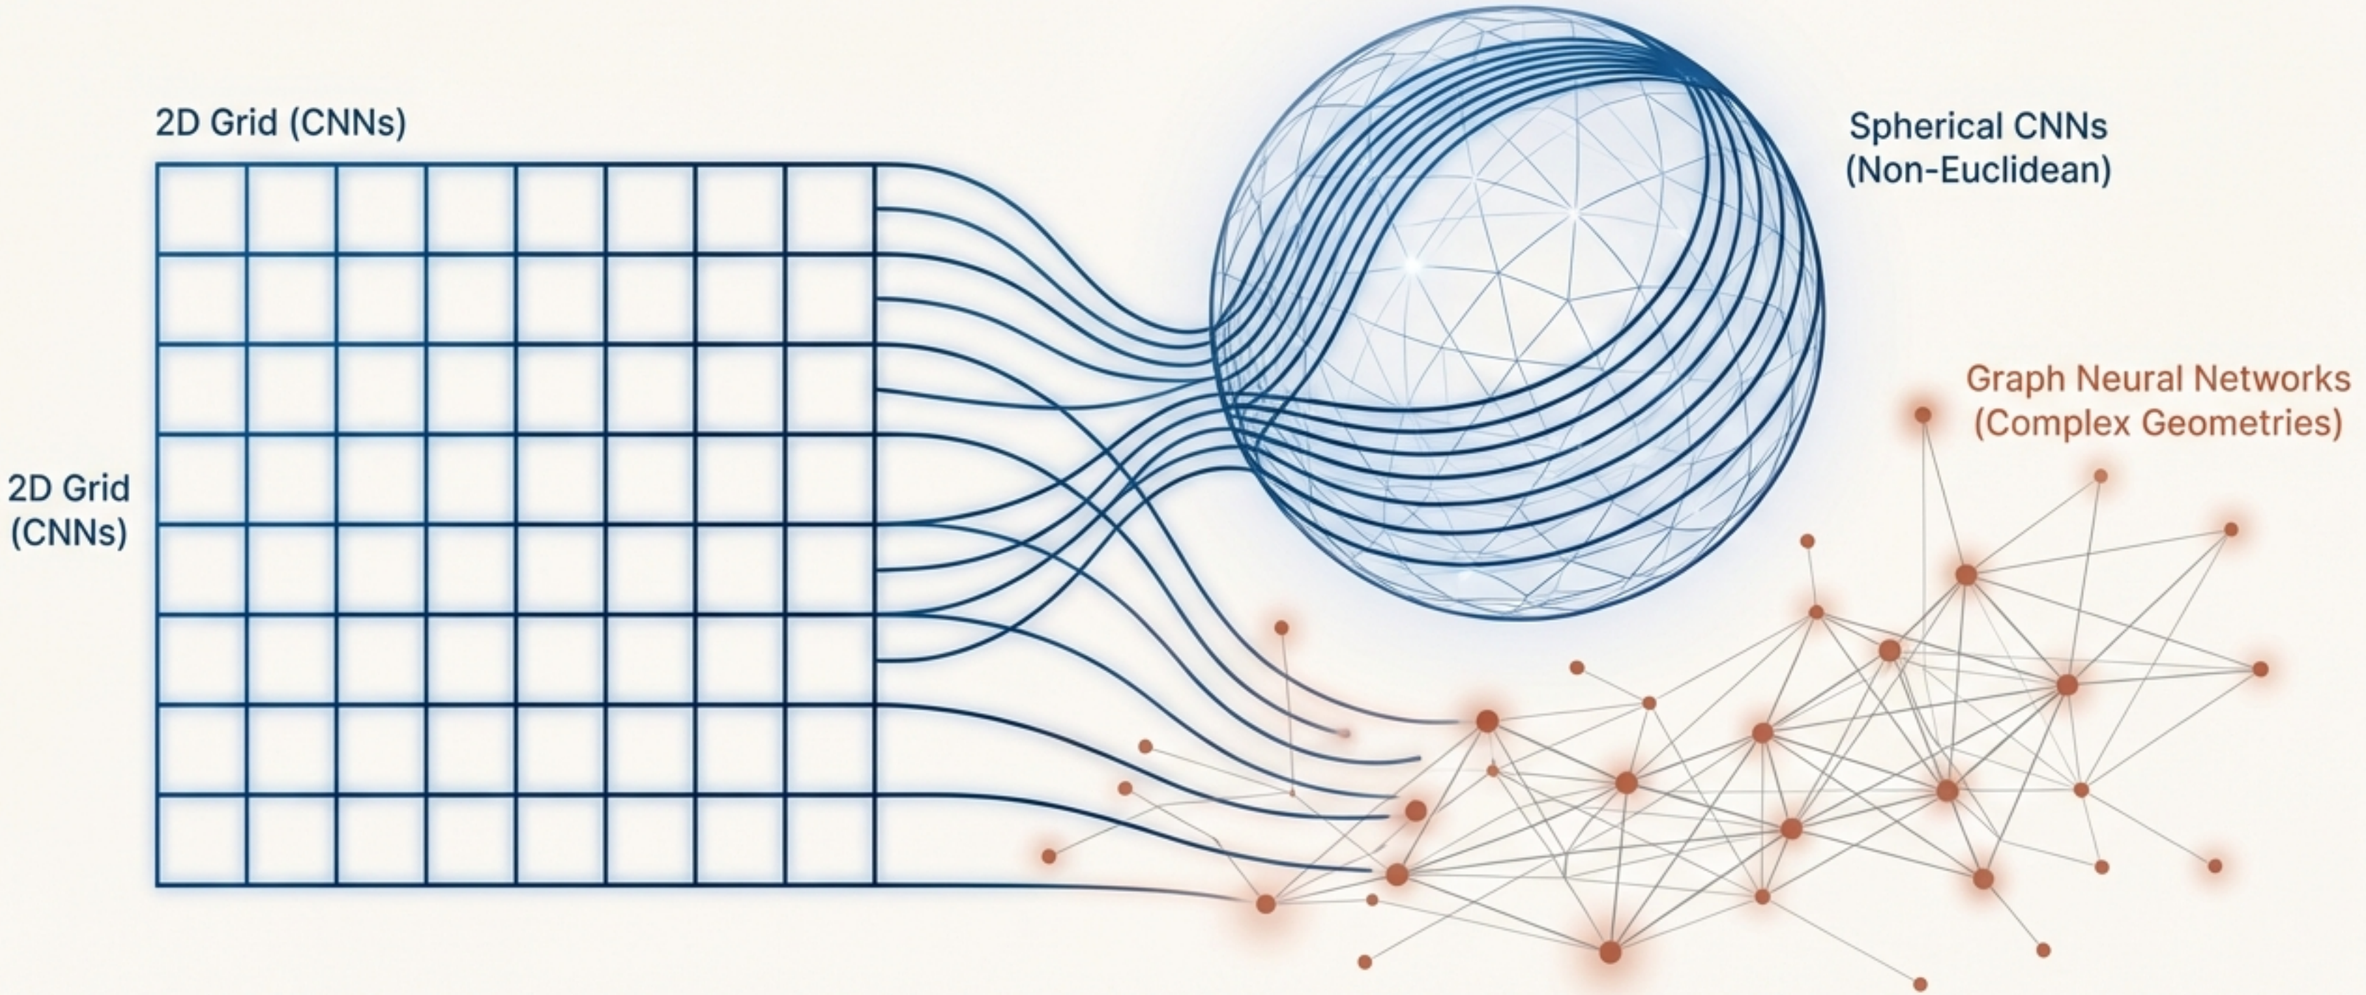
Breaking the Grid: From CNNs to Geometric Deep Learning

Understanding Spherical CNNs, Graph Neural Networks, and Their Place in the Modern AI Landscape



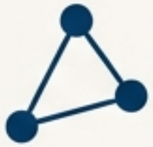
Our Journey of Generalization



1. **The Classic: What Makes CNNs Work?**
   Subtitle: The power and assumptions of the grid.



2. **Beyond the Grid I: Data on a Sphere**
   Subtitle: Introducing Spherical CNNs for rotational symmetry.



3. **Beyond the Grid II: Data as a Network**
   Subtitle: Introducing Graph Neural Networks for arbitrary structures.



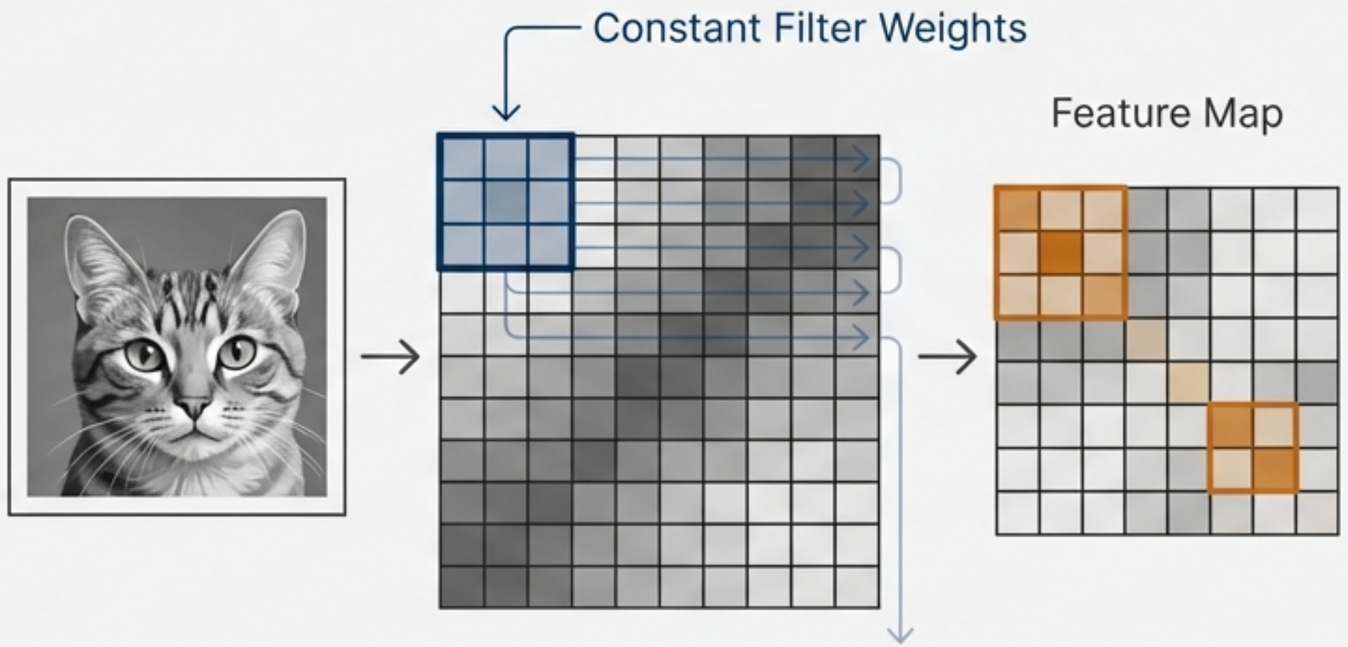
4. **The Modern Landscape: GNNs, CNNs, and Transformers**
   Subtitle: A unified view and a principled comparison.



5. **Synthesis & Key Takeaways**
   Subtitle: The principle of matching architecture to data geometry.

CNNs: Masters of the Flat World

Convolutional Neural Networks have revolutionized machine vision, achieving unprecedented success on 2D grid-like data such as images. Their efficiency stems from a core assumption: **translational weight sharing**. A single filter, or kernel, is applied across the entire image, assuming that the features it detects (like edges or textures) are independent of their location.



The CNN: A Master of the Grid Built on Translation Equivariance

Operation

A kernel (filter) slides across an image, sharing weights to detect local patterns efficiently.

The Key Symmetry: Translation Equivariance

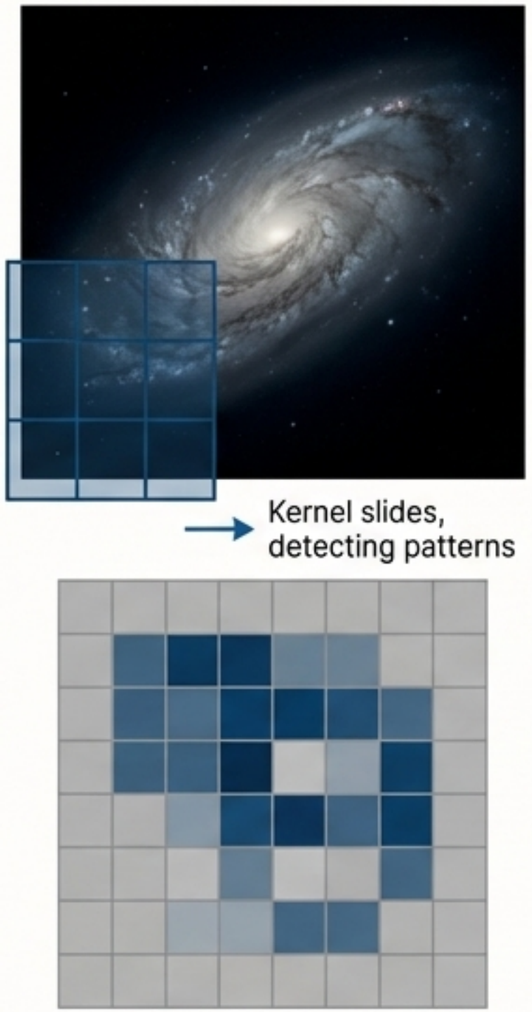
This is the crucial property. If an object moves in the input, its feature representation moves by the same amount in the output.

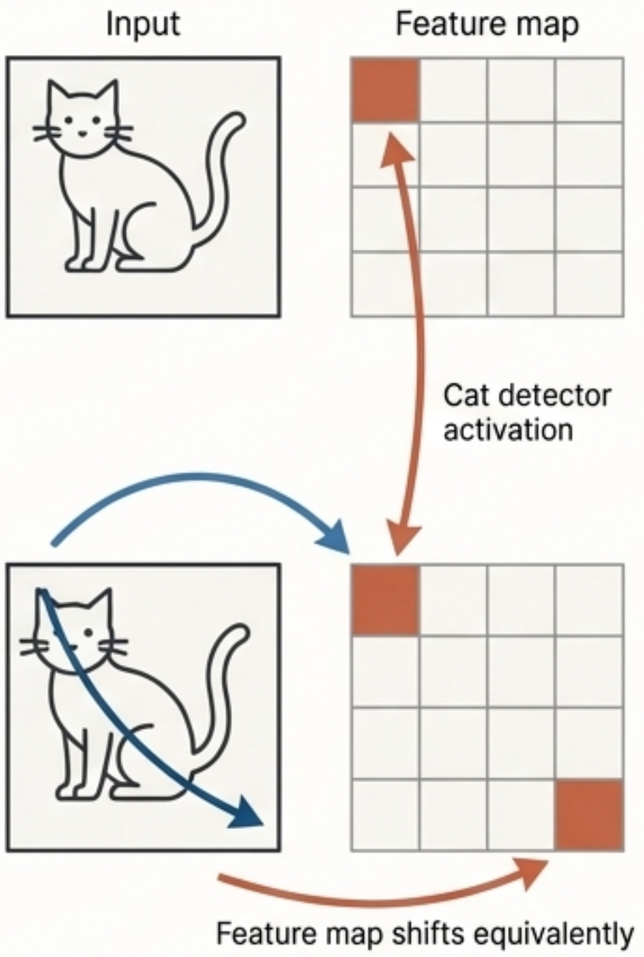
Mathematical Representation

$$f(\text{translate}(x)) = \text{translate}(f(x))$$

Key Takeaway

This property is incredibly powerful for images but is fundamentally tied to the assumption that data lives on a regular grid where 'translation' is a globally consistent and meaningful operation.





The Flat-World Assumption Breaks

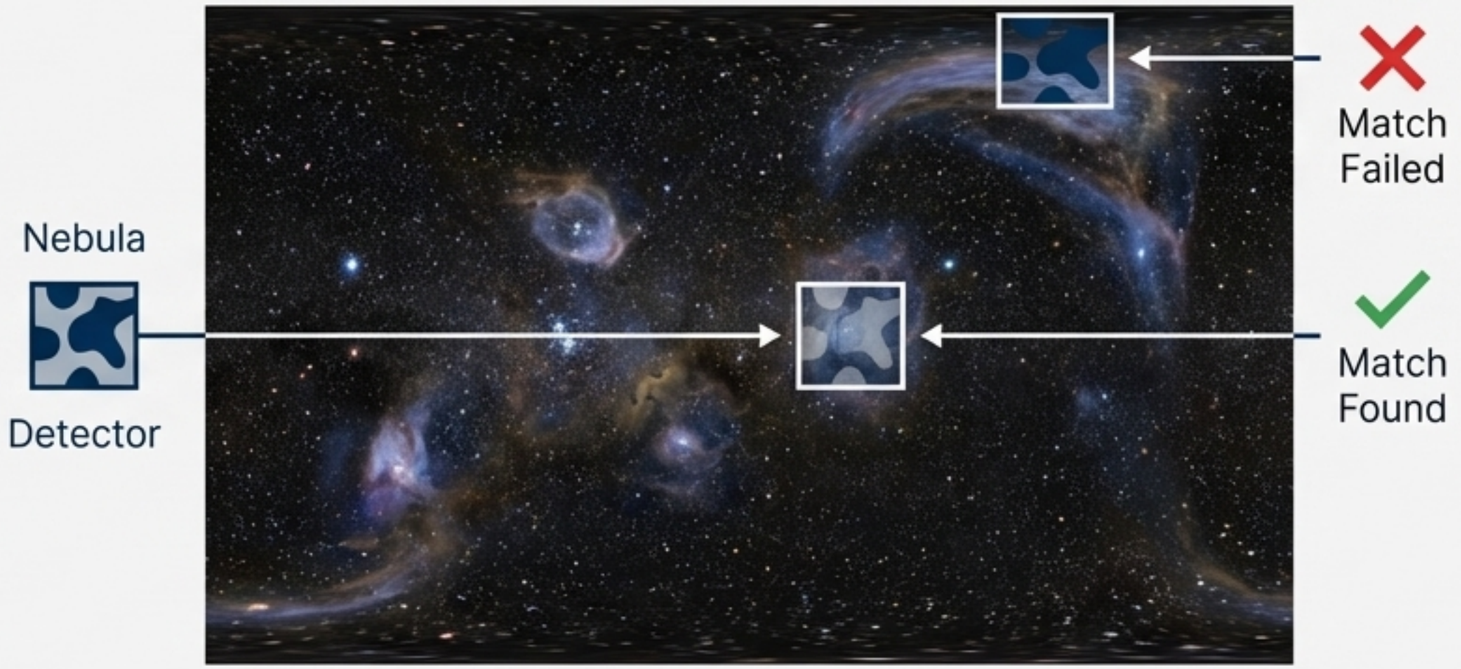
What happens when data isn't a flat grid? Consider omnidirectional (360°) vision, global climate maps, or molecular surfaces.

Projecting spherical data onto a flat plane introduces severe, space-varying distortions. This makes the core principle of CNNs ineffective.

"A naive application of convolutional networks to a planar projection of the spherical signal is destined to fail, because the space-varying distortions introduced by such a projection will make translational weight sharing ineffective." — (Attribution)

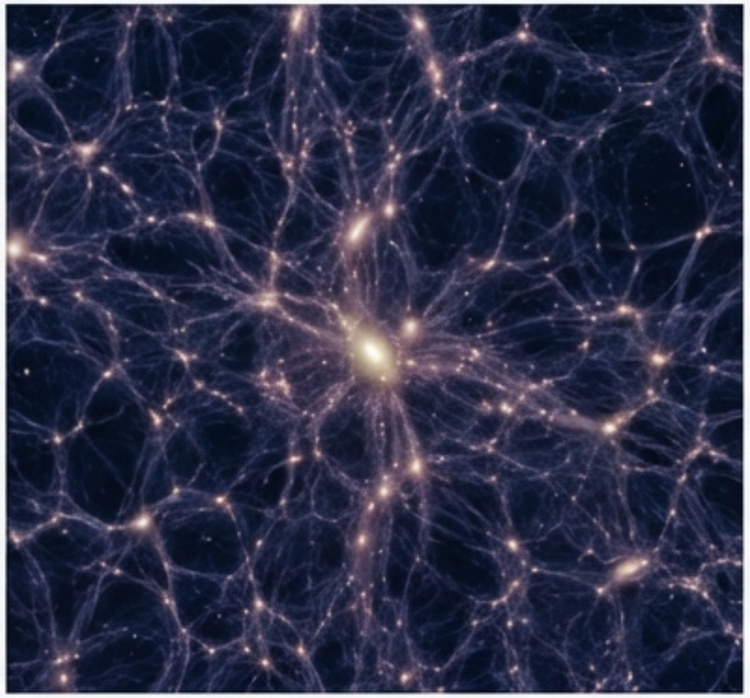


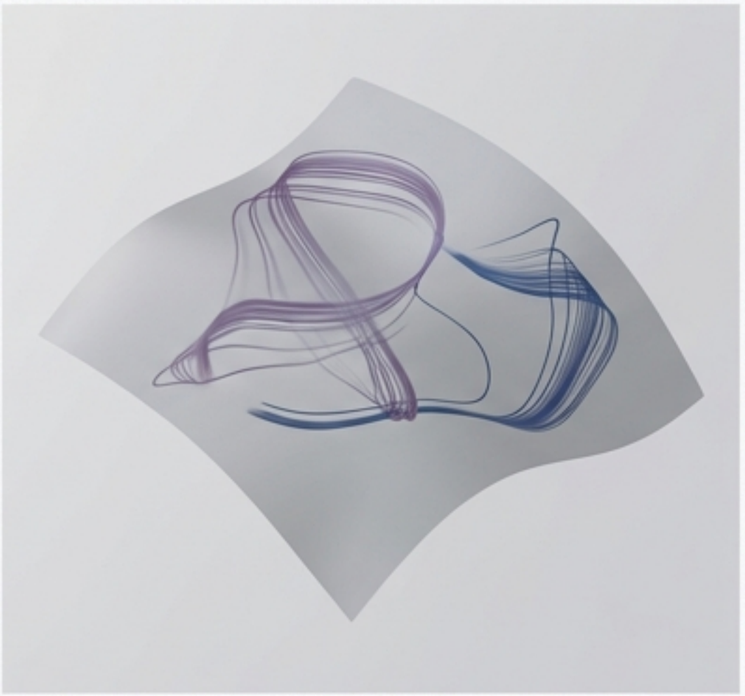
Native Spherical Data

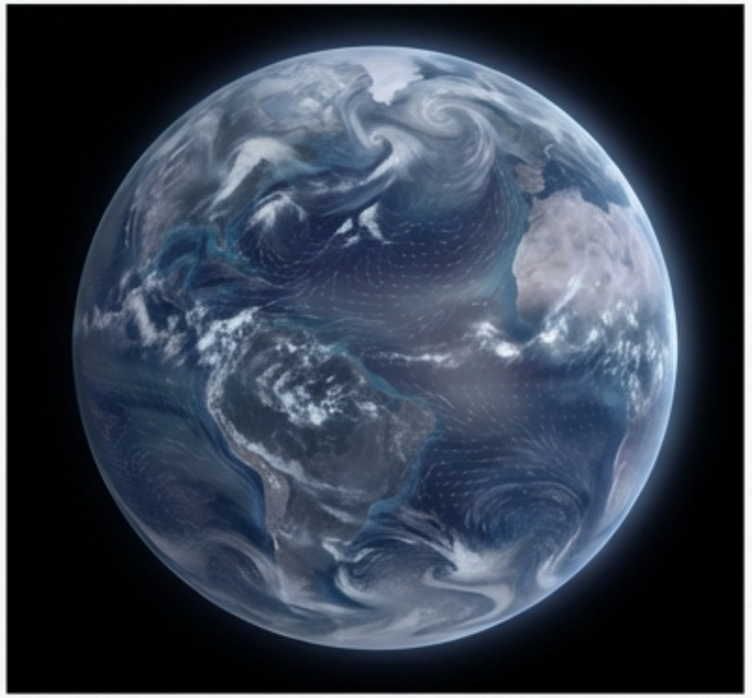


Planar Projection

Science is Not Flat

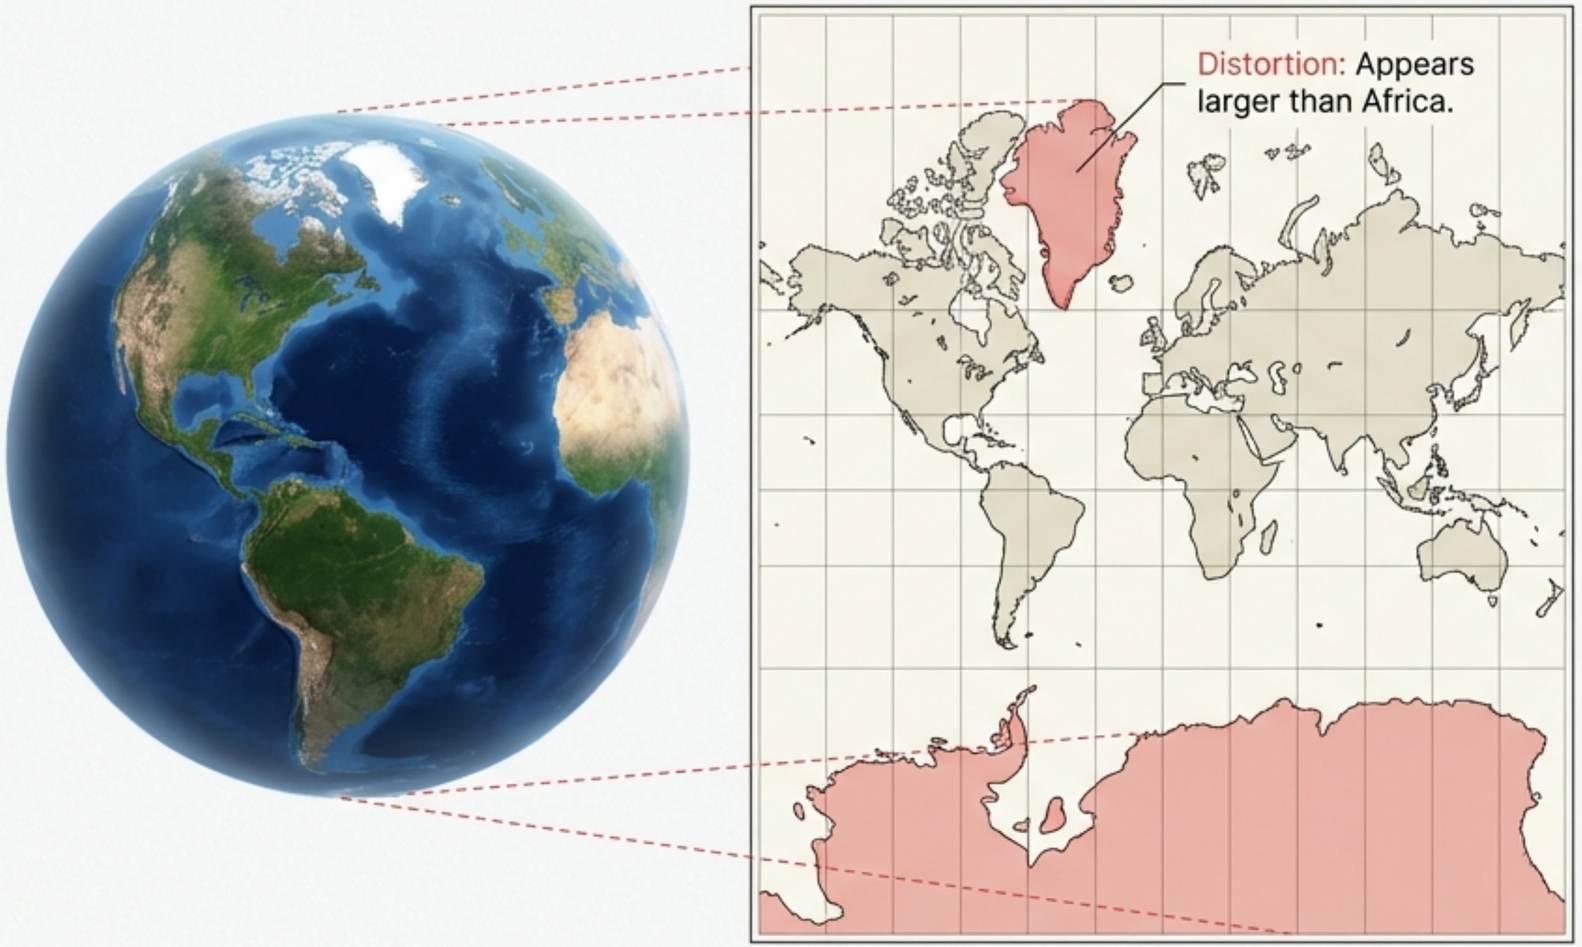






For decades, our most powerful computational models have been built on a powerful but flawed assumption: that data lives - **lives on a simple grid**. But the fundamental structures of the **universe**—from the gravitational dance of galaxies to the neural pathways of thought—**possess an intrinsic geometry**. This deck explores a new paradigm that respects this underlying shape, unlocking a new class of scientific discovery.

The Mercator Problem in Scientific AI



Applying standard CNNs to geometric data is like navigating with a distorted map.

The Mercator projection sacrifices geometric accuracy for the convenience of a flat grid. This creates massive, space-varying distortions.

Similarly, forcing non-Euclidean scientific data onto a grid-based model (like a standard CNN) breaks its fundamental symmetries and relationships. Translational weight sharing becomes ineffective. We are teaching our models a distorted version of reality.

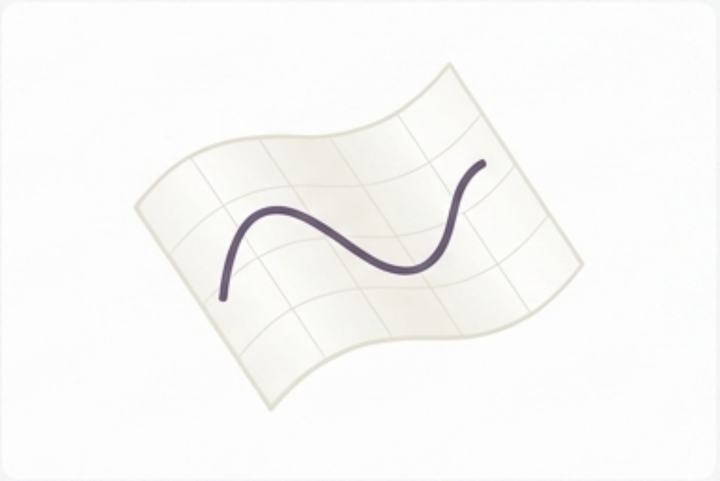
A Blind Spot for Structure, Symmetry, and Relationships

This “flat-world” assumption creates three fundamental challenges that limit scientific progress:



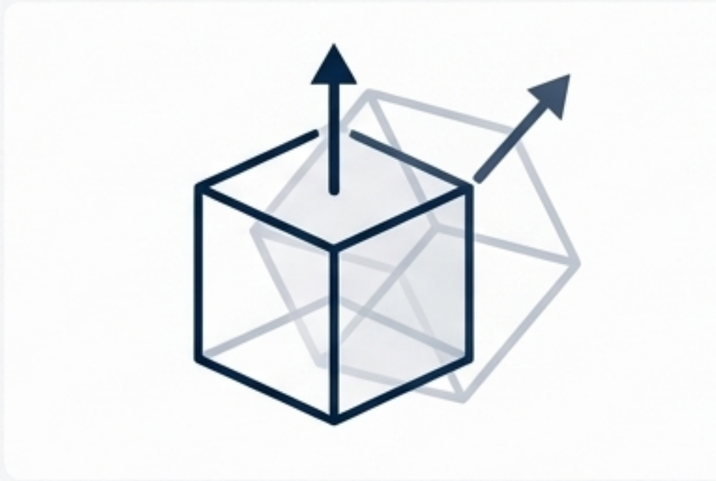
1. Ignoring Relationships

How do we model the intricate, non-grid connections between galaxies in a dark matter halo? GNNs treat the cosmic web as a graph, where nodes are galaxies and edges encode interactions. Standard models see pixels, not a network.



2. Ignoring Intrinsic Structure

How do we analyze neural activity that unfolds on a complex, low-dimensional manifold? Neural states may be embedded differently across animals, making direct comparison impossible without understanding the underlying geometric space.



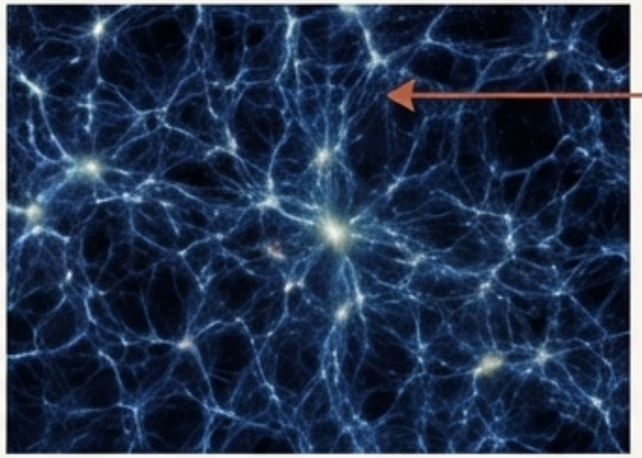
3. Ignoring Physical Symmetries

How do we create models that respect fundamental physical laws? An AI model for cosmology or climate should understand that rotating a system doesn't change the physics. This requires building in rotational and translational equivariance (E(3) symmetry).

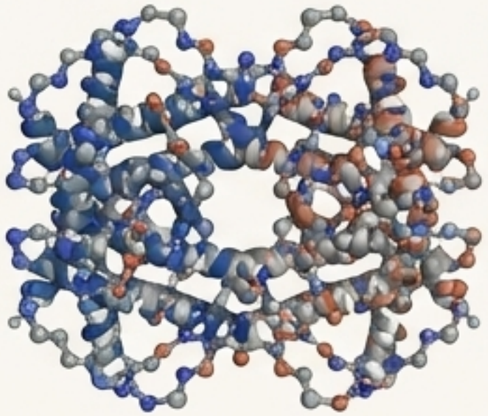
The Problem: The Universe is Not a Grid

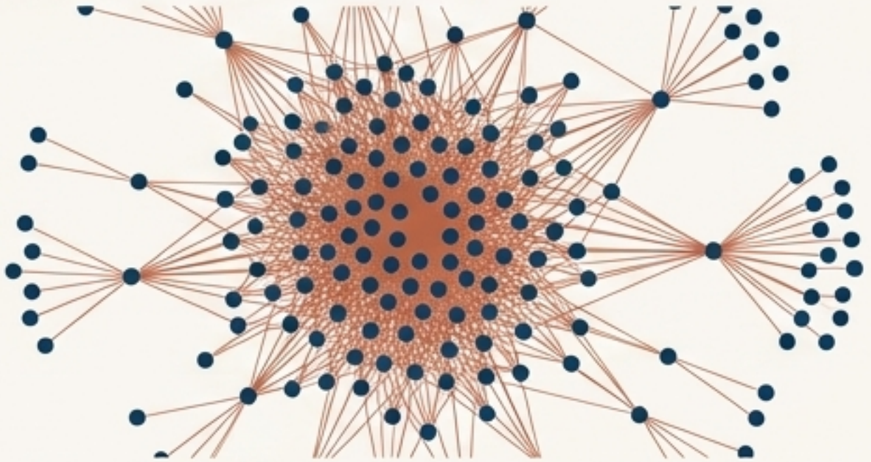
The CNN's assumption of a regular grid structure breaks down for many high-value datasets. Applying standard CNNs to these domains leads to ineffective learning due to distortions and a mismatch between the model's architecture and the data's inherent geometry.





"Galaxies are not randomly distributed, but rather are organized into clusters, filaments, and voids, forming a rich network of connections."





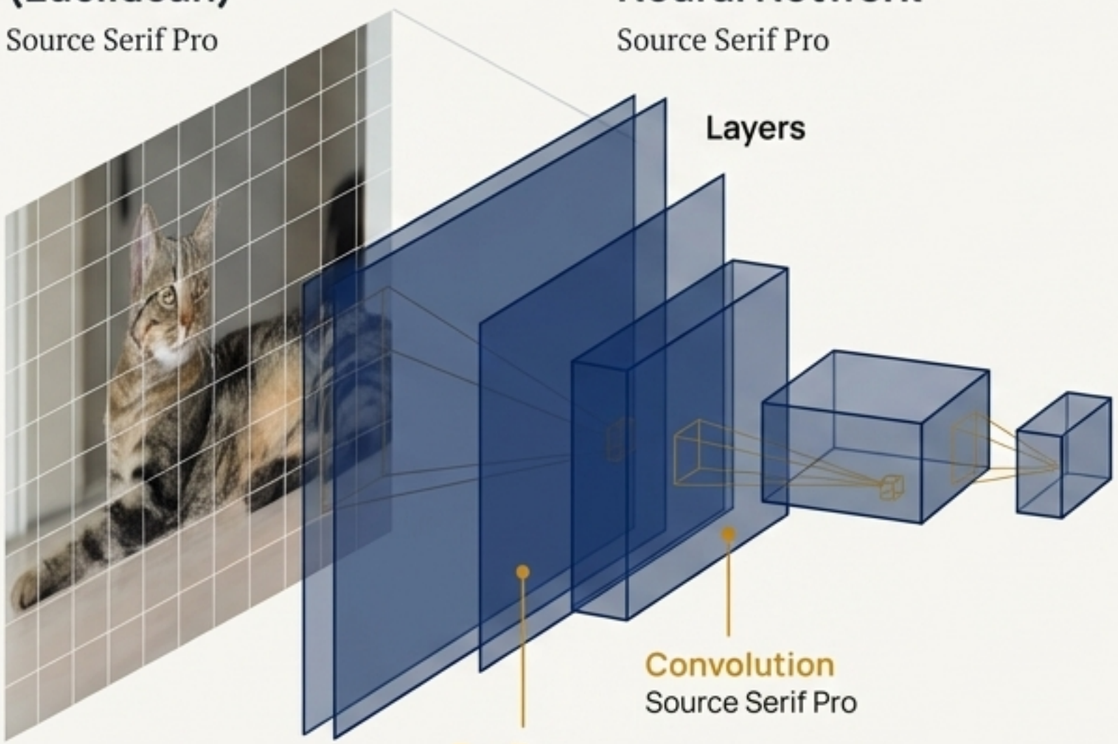
Beyond the Grid: Geometric Deep Learning and the Principle of Symmetry

Grid-Based Data (Euclidean)

Source Serif Pro

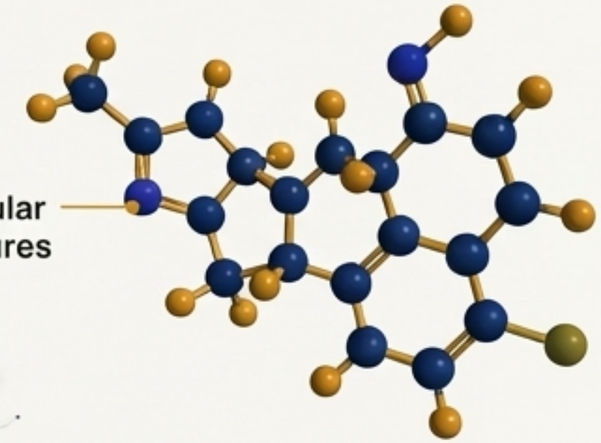
Convolutional Neural Network

Source Serif Pro

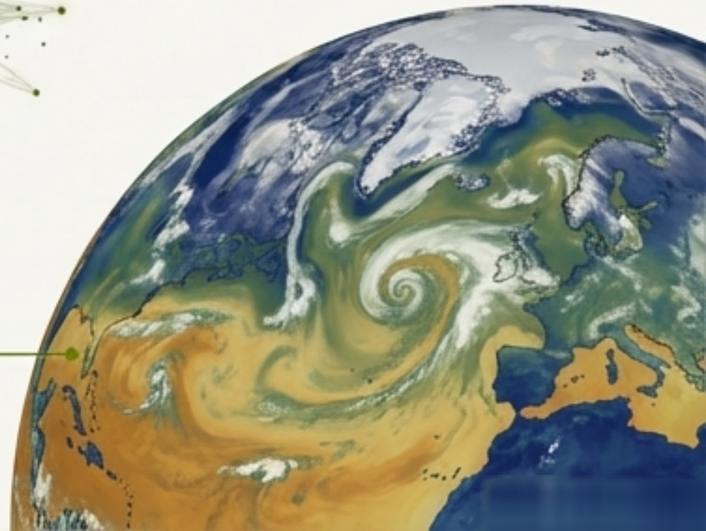


Non-Euclidean Data (Geometric)

Source Serif Pro





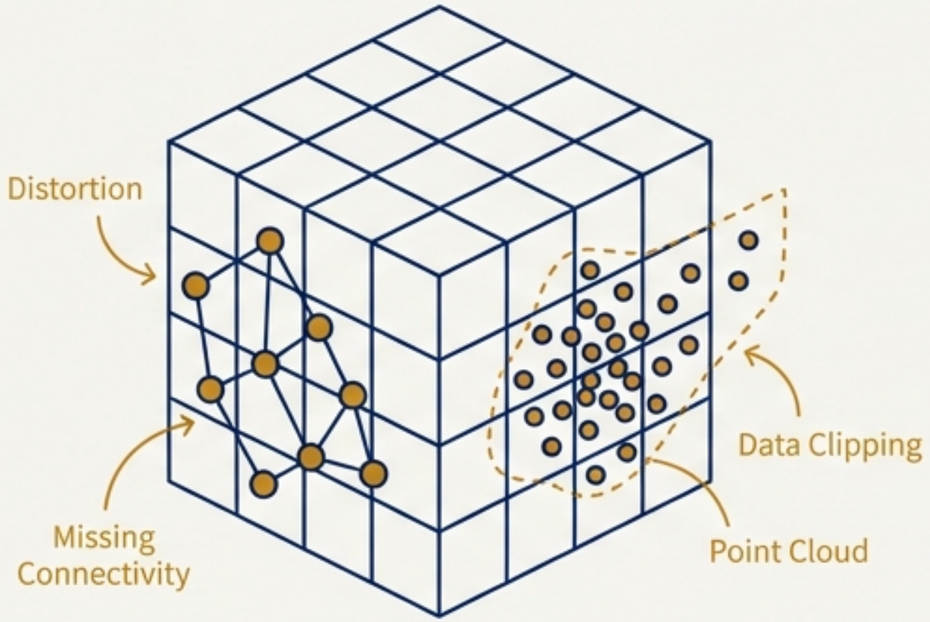


The Grid Assumption: A Fundamental Limitation

Deep learning has mastered grid-like data (images, time series), but most of the world's critical data is not a grid. Molecules, physical systems, social networks, and climate data exist as graphs, point clouds, and manifolds.

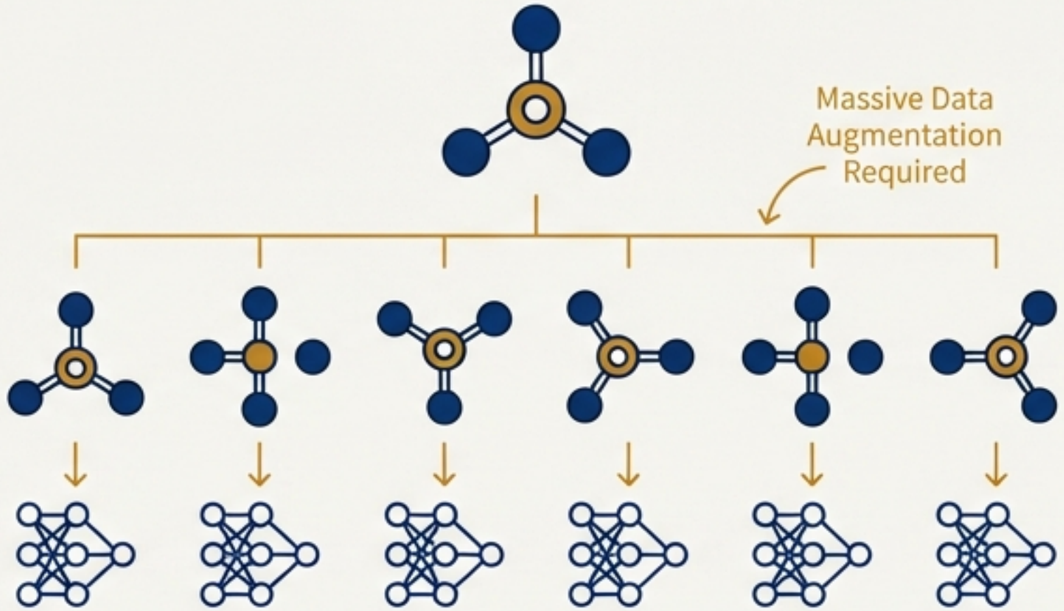
Two Core Weaknesses of Standard Architectures

1. Structural Incompatibility



Methods like CNNs are designed for regular grids. They cannot natively process the irregular connections and variable structure of non-Euclidean data. As noted by Zhao et al., this leads to "excessively smooth output" and a "lack of capacity to model complex spatial relationships."

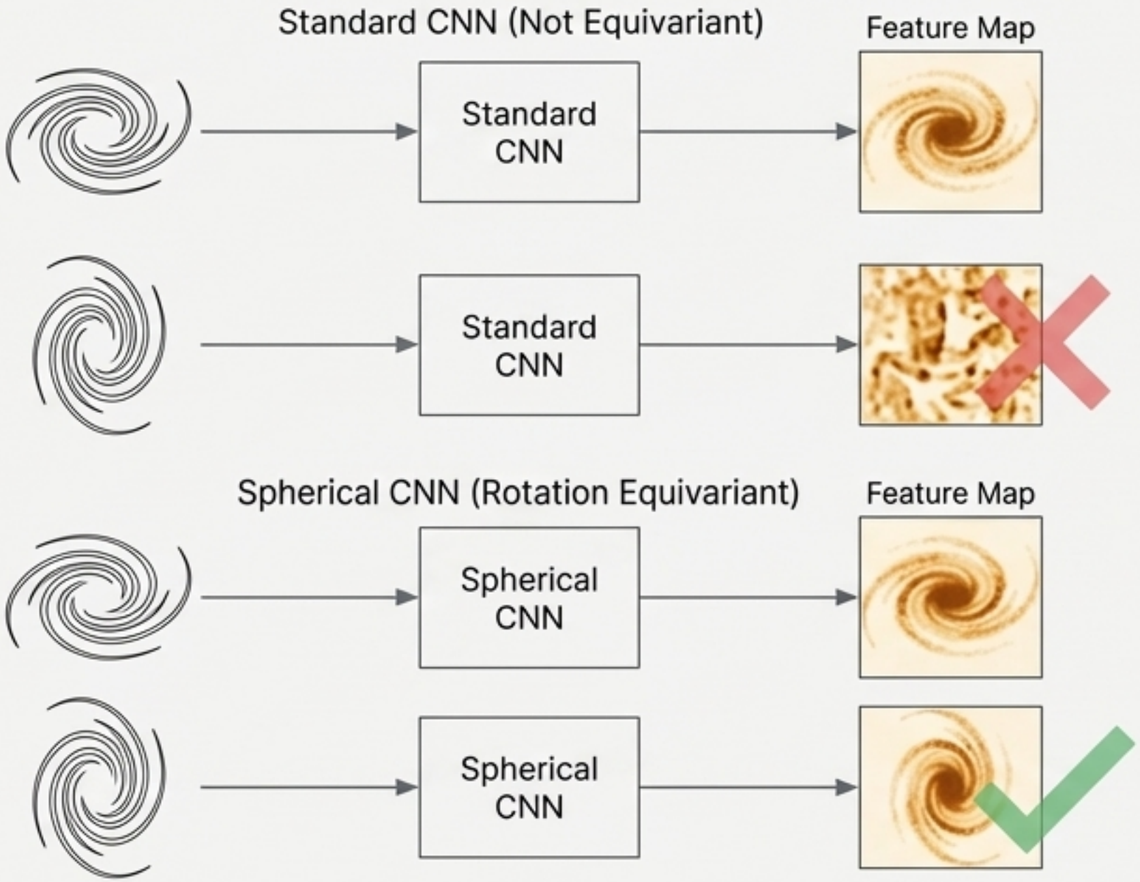
2. Symmetry Blindness



Standard models lack the necessary inductive biases for fundamental physical symmetries. They must learn principles like rotation and translation from scratch, which requires massive data augmentation with no guarantee of generalization. This is highly inefficient and often fails on complex 3D data.

The Guiding Principle: Speak the Language of Symmetry

The solution is rooted in a fundamental principle of physics: **rotation equivariance**. A simple definition: *If the input rotates, the features detected by the network should rotate in exactly the same way*. This preserves the geometric structure of the data, unlike *invariance*, where the output remains static.

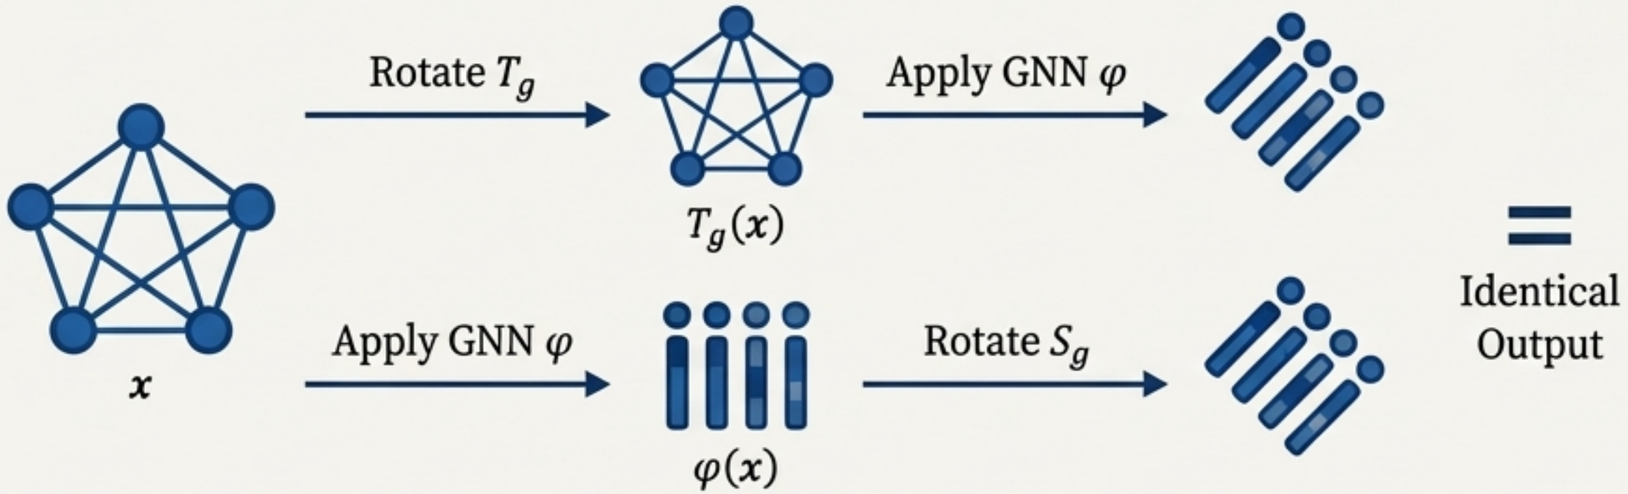


The Solution: Building Symmetries into the Architecture

Key Idea: Instead of forcing a model to learn physics, we build physics into the model. The guiding principle is **Equivariance.**

Definition: A function φ is equivariant if, when the input transforms, the output transforms with it in a predictable way.

Formal Definition: $\varphi(T_g(x)) = S_g(\varphi(x))$
Where T_g is a transformation on the input x, and S_g is the corresponding transformation on the output $\varphi(x)$.





Equivariance

Transforms ***with*** the input.

Preserves richer, directional information (e.g., velocity vectors that rotate with the system).



Invariance

Unaffected by the input's transformation.

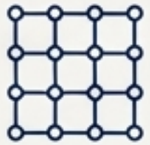
Collapses information into a single, stable output (e.g., the total energy of a system, which is constant regardless of rotation).

A Unifying Blueprint for Deep Learning

This approach provides a common mathematical framework to understand the most successful neural network architectures. It is a “geometric unification” in the spirit of Felix Klein’s Erlangen Program.

> Quote from Bronstein et al., “Geometric Deep Learning” (arXiv:2104.13478): *“Such a “geometric unification” endeavour . . . serves a dual purpose: on one hand, it provides a common mathematical framework to study the most successful neural network architectures, such as CNNs, RNNs, GNNs, and Transformers. On the other hand, it gives a constructive procedure to incorporate prior physical knowledge into neural architectures...”*

Key Takeaway: This isn’t just a new model type; it’s a principled way to build architectures by defining the underlying geometry and symmetry of the data domain.



Grids → Symmetry: Translation → **CNNs**

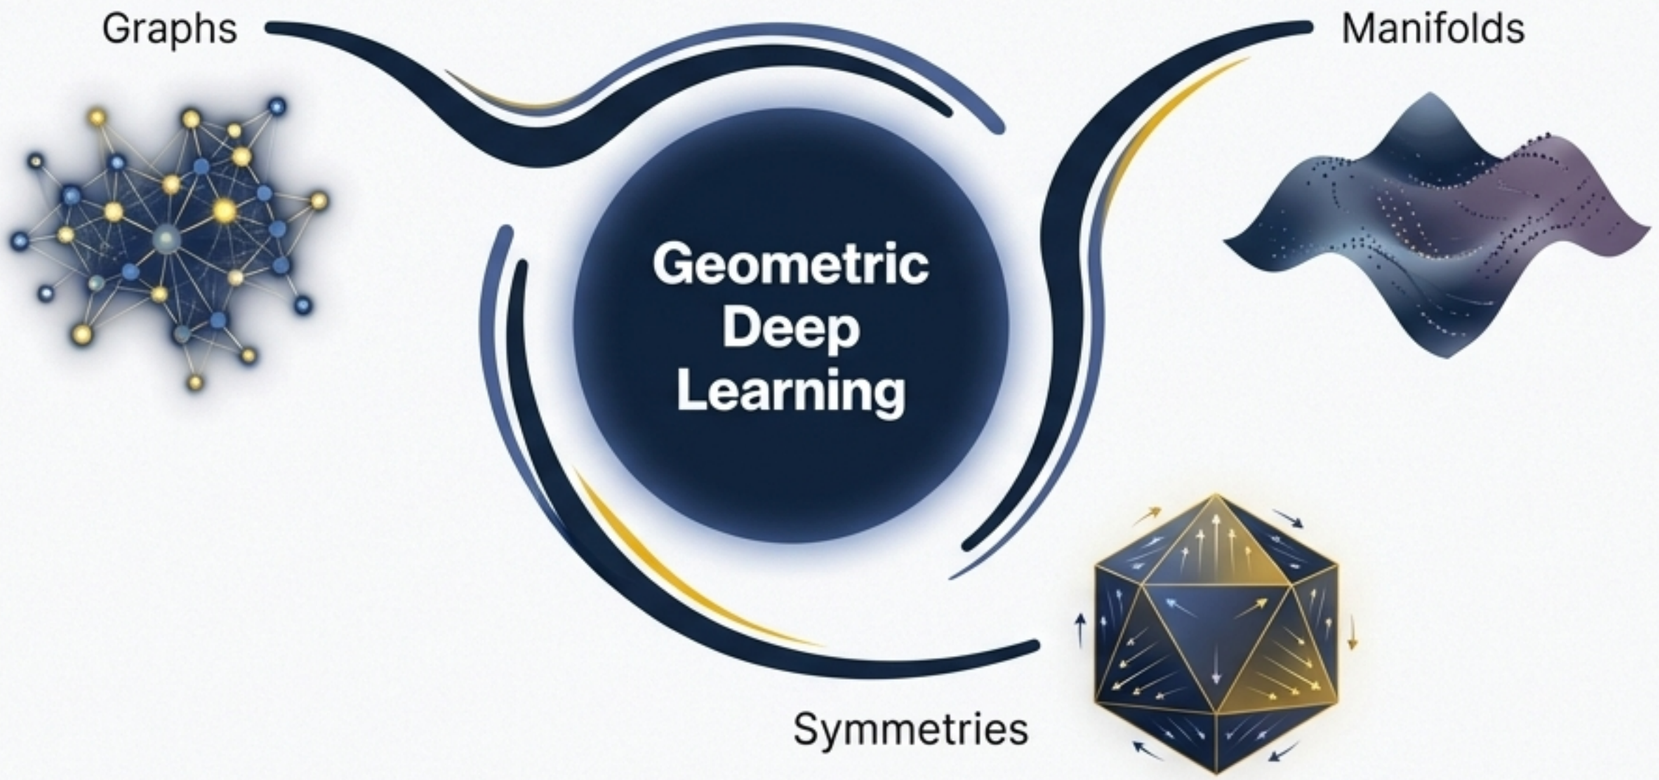


Graphs → Symmetry: Permutation → **GNNs**



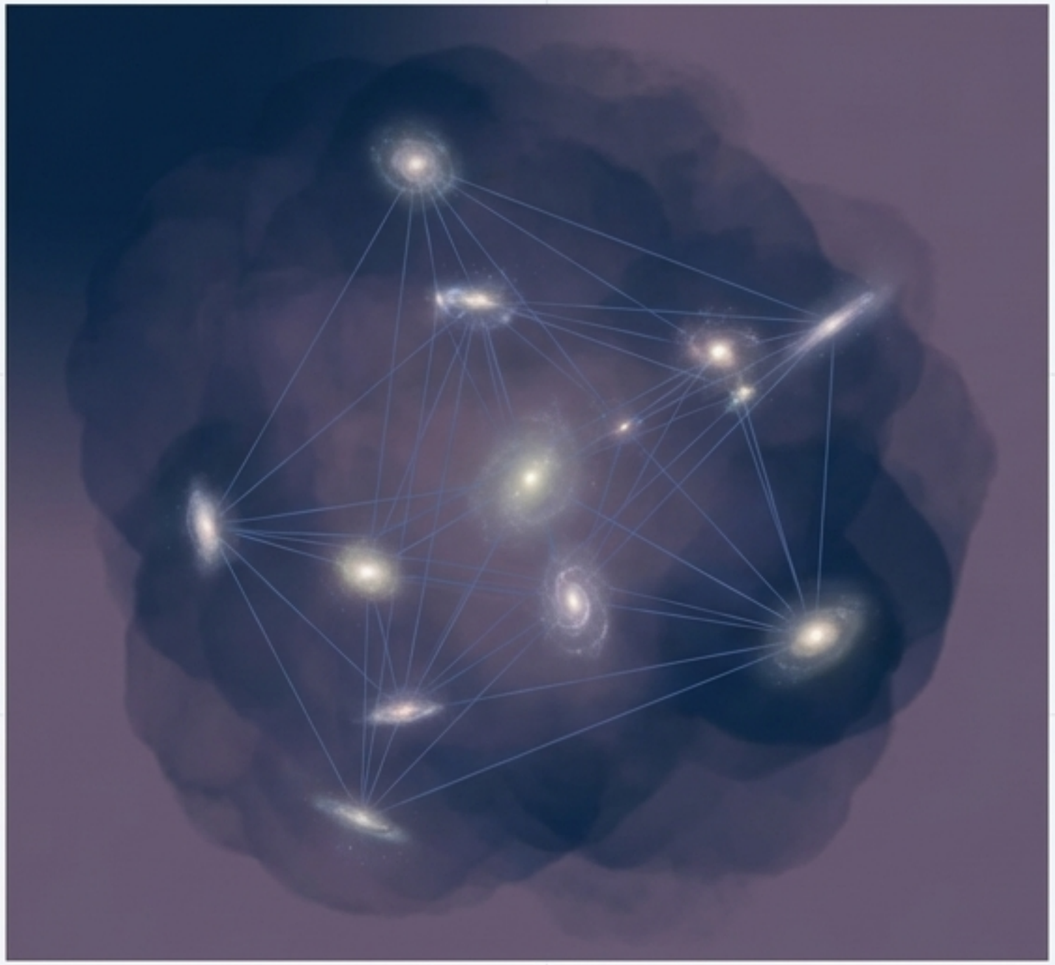
Spheres/3D Space → Symmetry: Rotation, Translation → **Equivariant GNNs**

Geometric Deep Learning: Teaching AI the Language of Nature



Geometric Deep Learning is a class of models designed to operate directly on **non-Euclidean data**. Instead of forcing data onto a grid, GDL leverages the intrinsic geometry and **symmetry** of the problem as a powerful inductive bias. It is an area of research that seeks to exploit inherent data structures using mathematical graphs, group theory, and other geometric concepts.

For Relationships: Graph Neural Networks (GNNs)



Concept:

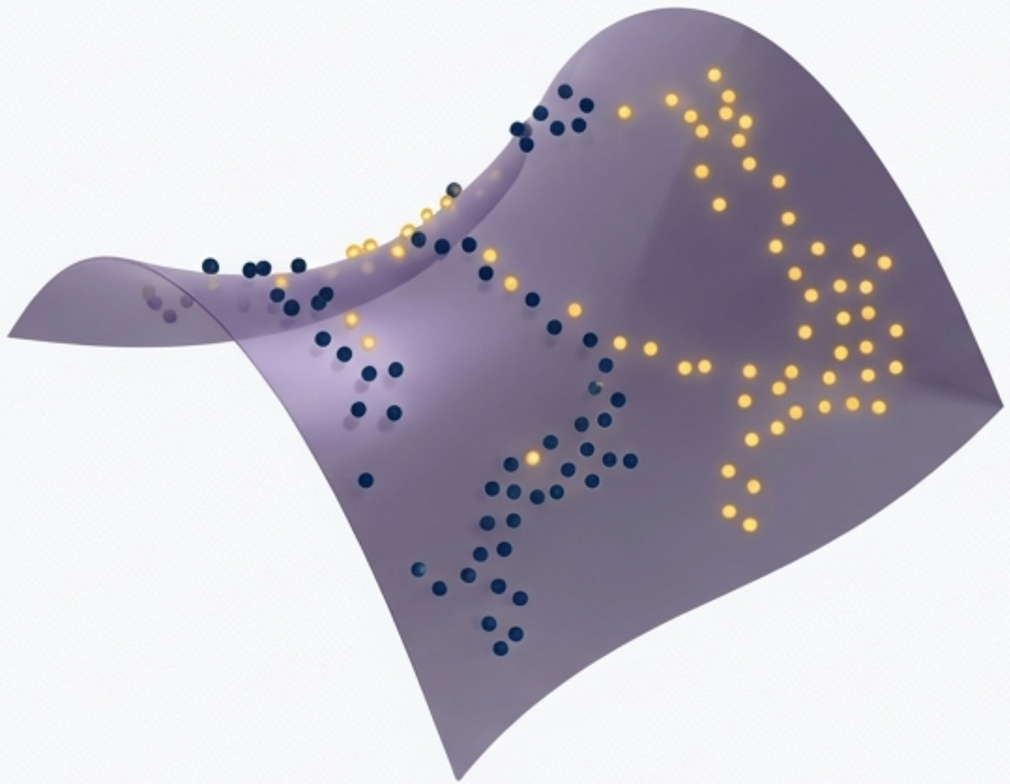
GNNs represent sparsely distributed objects and their complex correlations. By iteratively aggregating information from neighboring nodes through message passing, GNNs capture the spatial correlations and hierarchical relationships inherent in the cosmic web.

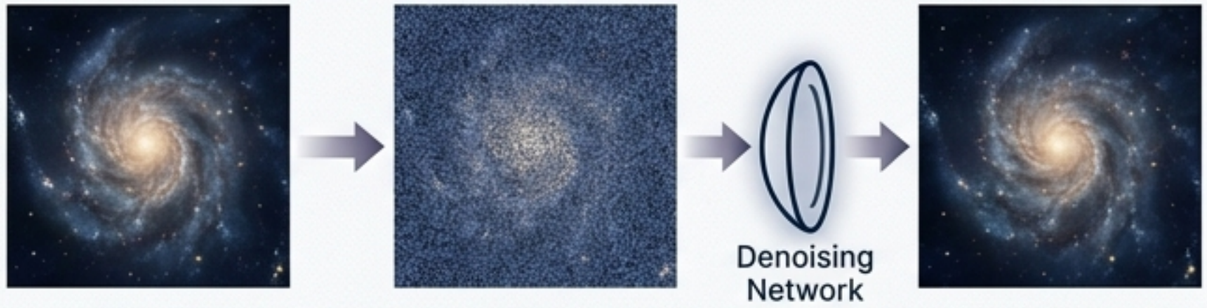
Application:

In cosmology, galaxies are nodes and edges encode their interactions. This allows a model to learn from the entire network structure, not just individual galaxy properties.

> *"GNNs provide a powerful framework by treating the cosmic large-scale structure as a set of graphs."*

For Intrinsic Structure: Manifolds & Diffusion Models





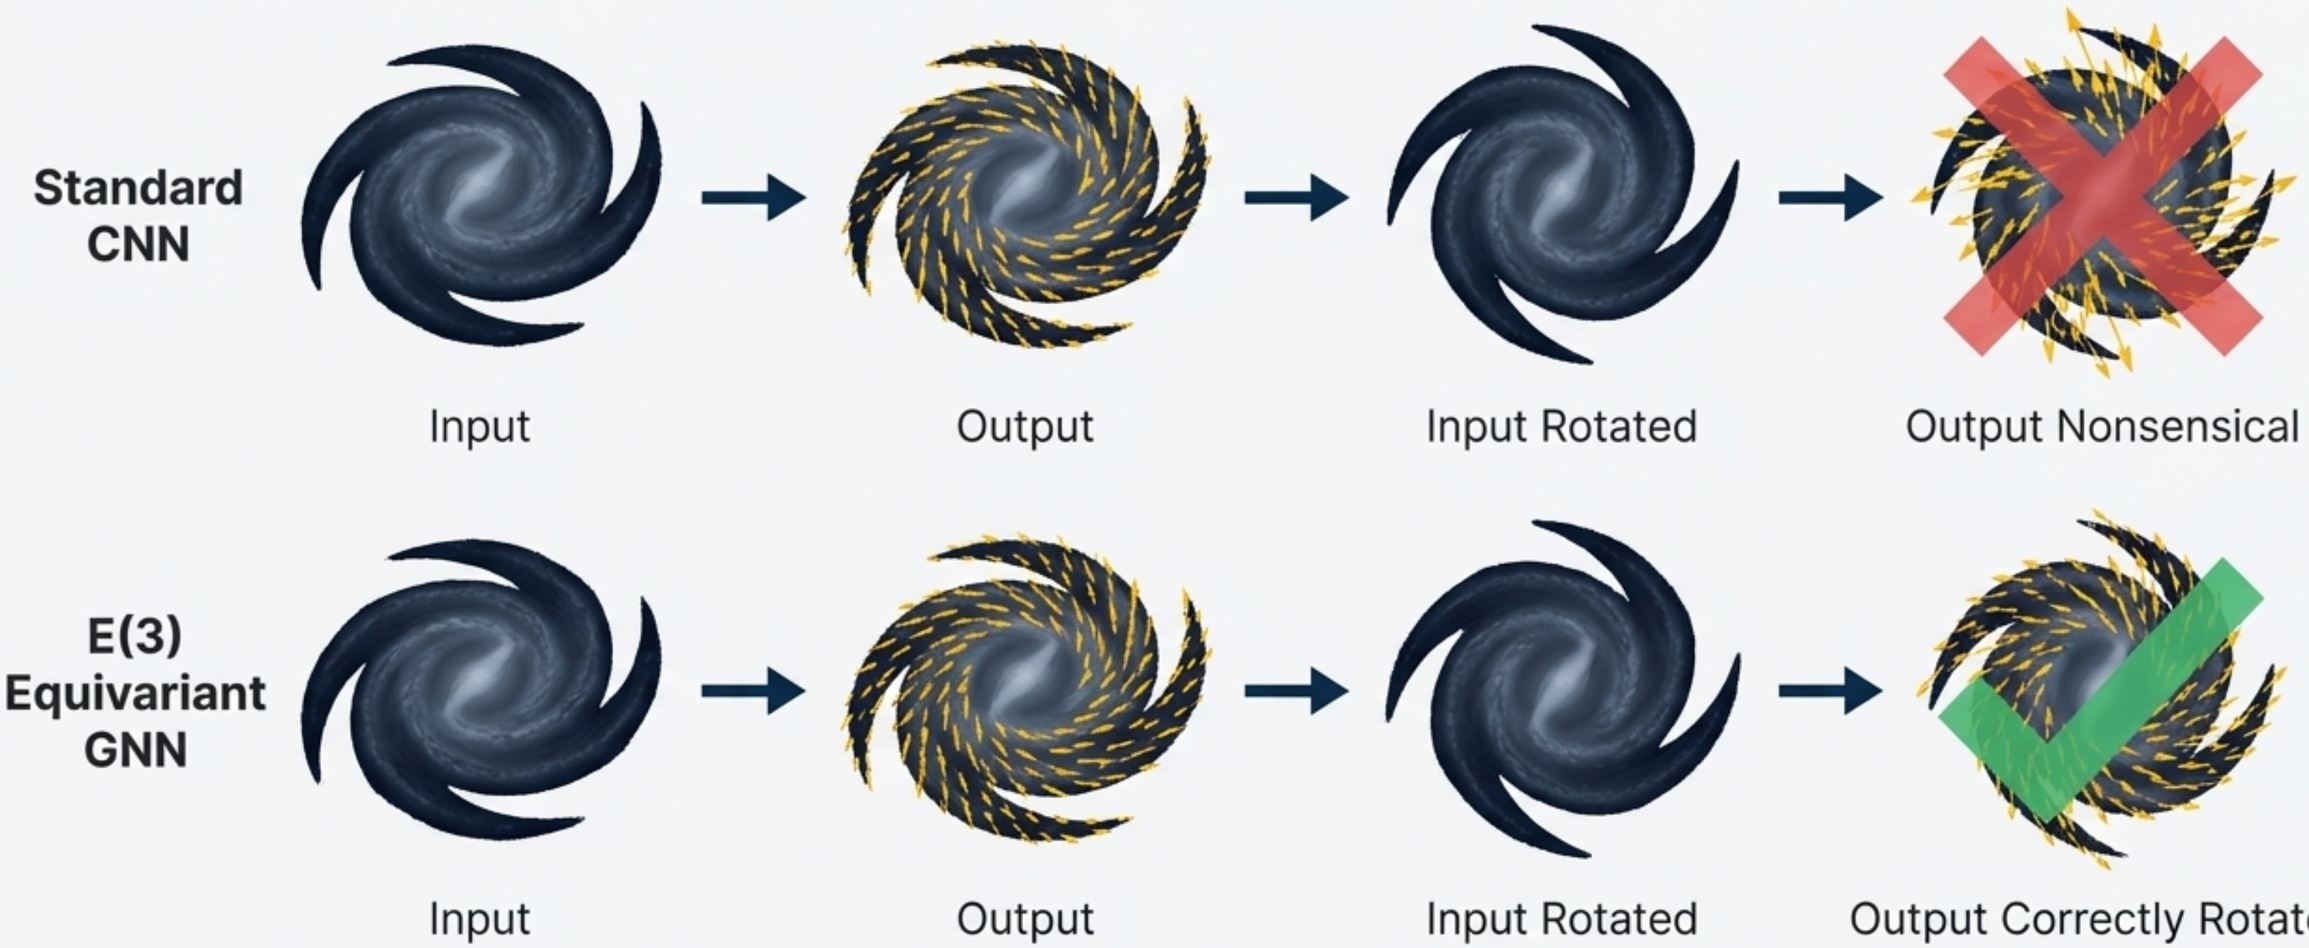
[Original Data] [Progressive Noise Added] [Learned Denoising Process]

Concept: Medium
Many complex systems, like neural populations, evolve on low-dimensional **manifolds**. GDL methods learn the dynamics *on* this **intrinsic surface**.

Generative Power: Medium
Diffusion models learn to reverse a noising process. By training a noising process. By training a network to **denoise** samples, we can generate new, high-fidelity data that respects the **manifold's structure**.

Application (Jagvaral et al.): Medium
A diffusion model on the **SO(3) manifold** (the space of 3D rotations) can generate realistic galaxy orientations, a non-Euclidean property that is otherwise difficult to model.

For Physical Symmetries: E(3) Equivariance



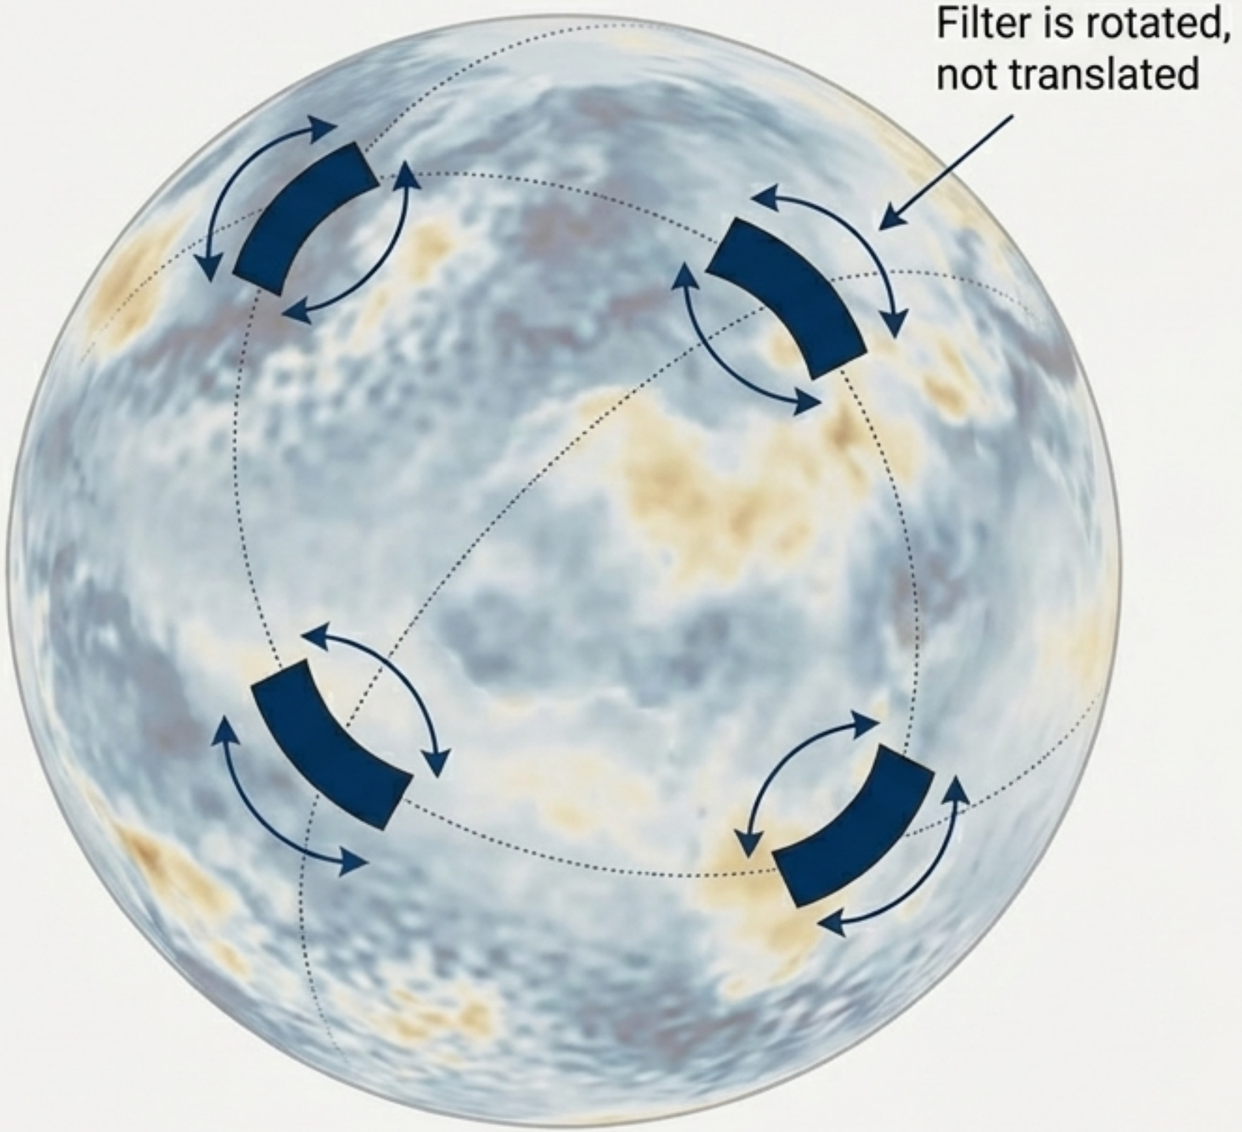
Equivariance ensures that if the input to a model is transformed (e.g., rotated), the output is transformed in the same way. An E(3) equivariant model understands 3D space, respecting the symmetries of rotation, translation, and reflection. This enforces physical consistency by design, rather than hoping the model learns it from data augmentation alone.

Convolution, Reimagined for the Sphere

The Spherical CNN is a new class of network designed specifically for spherical data.

Its core innovation is replacing the 2D sliding filter with a rotation-equivariant operation called **spherical cross-correlation.**

This new operation is designed from the ground up to respect the sphere's geometry, ensuring the principle of equivariance is built directly into the network's architecture.



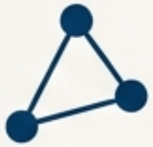
Our Journey of Generalization



1. **The Classic: What Makes CNNs Work?**
   Subtitle: The power and assumptions of the grid.



2. **Beyond the Grid I: Data on a Sphere**
   Subtitle: Introducing Spherical CNNs for rotational symmetry.



3. **Beyond the Grid II: Data as a Network**
   Subtitle: Introducing Graph Neural Networks for arbitrary structures.



4. **The Modern Landscape: GNNs, CNNs, and Transformers**
   Subtitle: A unified view and a principled comparison.



5. **Synthesis & Key Takeaways**
   Subtitle: The principle of matching architecture to data geometry.

Generalization 1: Convolution on the Sphere

Core Idea

A Spherical CNN is an architecture designed for data that naturally lives on a sphere (S^2). It replaces the core operation of CNNs to respect the new geometry.

The New Operation: Rotation

Instead of **translating** a flat filter, we **rotate** a spherical filter. The operation is designed to be equivariant to 3D rotations (SO(3) group).

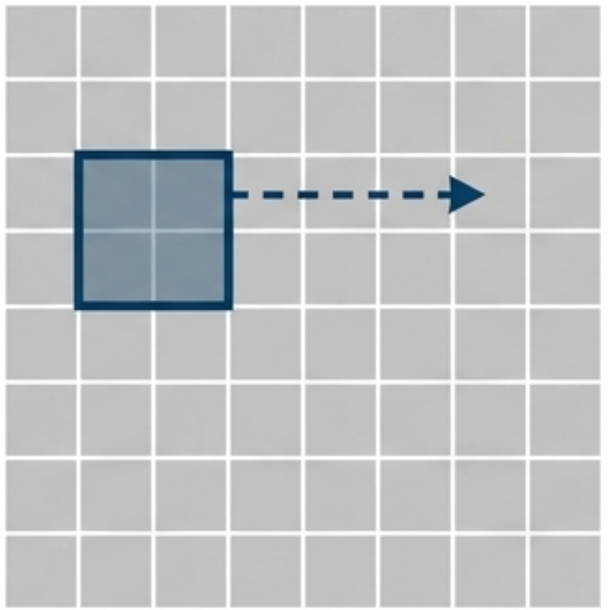
Mechanism

The operation is a spherical cross-correlation. This correlates a spherical patch of the input signal with with a spherical filter at all possible orientations.

Efficiency

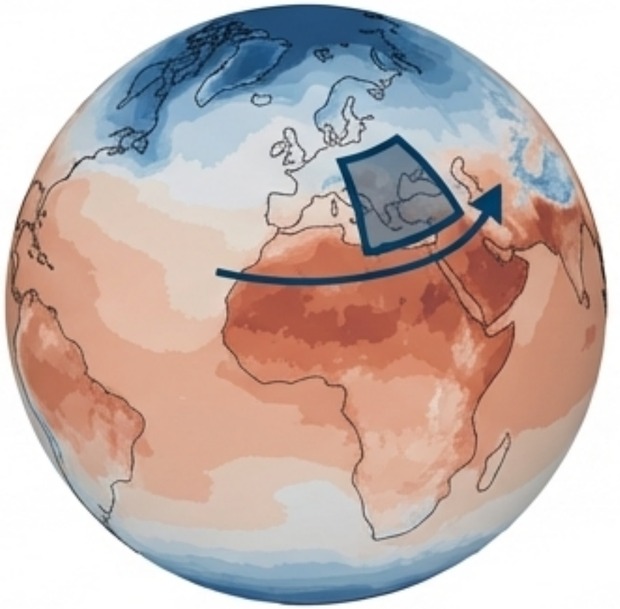
Just as standard CNNs use the Fast Fourier Transform (FFT), Spherical CNNs use a generalized, non-commutative FFT to make this computationally feasible.

Planar Convolution



Translation on $\mathbb{R}^2$

Spherical Convolution

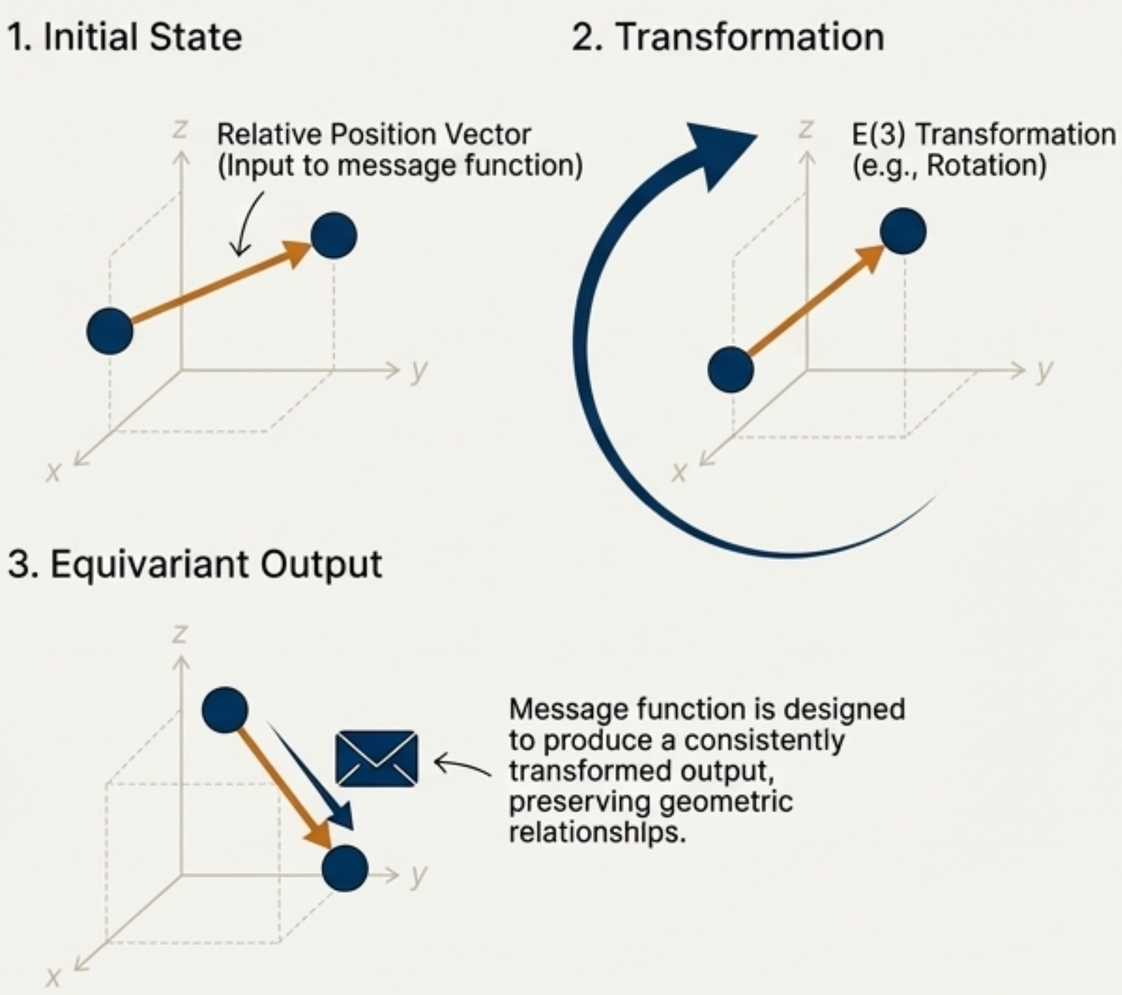


Rotation on $\mathbb{S}^2$

Head-to-Head: Upgrading Symmetry from Grid to Sphere

Feature	Standard CNN	Spherical CNN
Data Domain	Grid ($\mathbb{R}^2$)	Sphere ($\mathbb{S}^2$)
Core Operation	Planar Cross-Correlation	Spherical Cross-Correlation
Key Symmetry	Translation Equivariance	**Rotation Equivariance (SO(3))**
Weight Sharing	Across all translations	Across all rotations & positions
Use Cases	Standard image/video	360° vision, climate science, molecular data

Building Symmetries into the Network Architecture



GNNs can be specifically designed to be **E(3) equivariant**, meaning their operations inherently respect 3D rotations, translations, and reflections—the fundamental symmetries of physical space.

This is not learned through data augmentation; it's an **explicit** architectural constraint. The *message passing* and *node update* functions are mathematically constructed to be equivariant.

These functions use geometric quantities, like the **relative position vectors** between nodes, which transform predictably under E(3) transformations, ensuring the network's outputs transform correctly with its inputs.

Generalization 2: What if there is no regular structure at all?

The Next Challenge

Spheres are regular, but many datasets are simply nodes (entities) and edges (relationships). This is the domain of graphs.

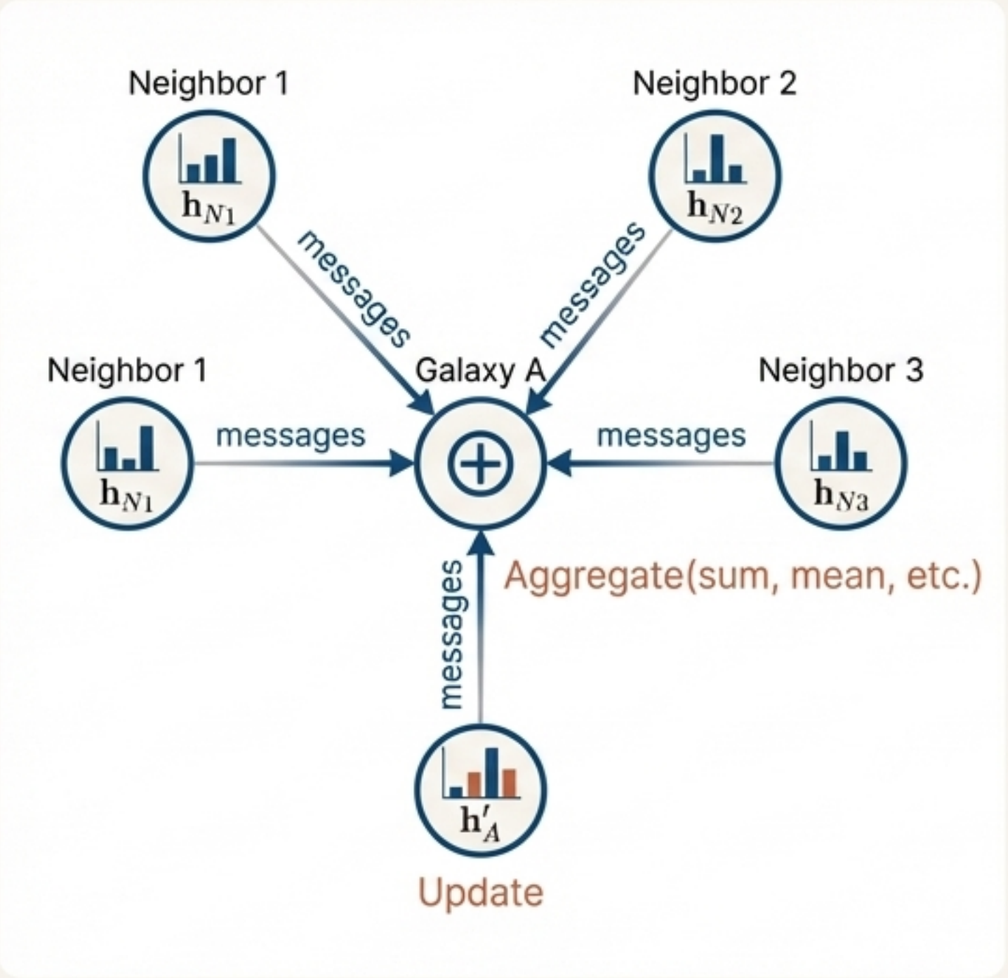
Core Idea: Message Passing

Graph Neural Networks (GNNs) learn features for each node by aggregating information from its neighbors. The “convolution” becomes a process of localized message passing.

From the Source

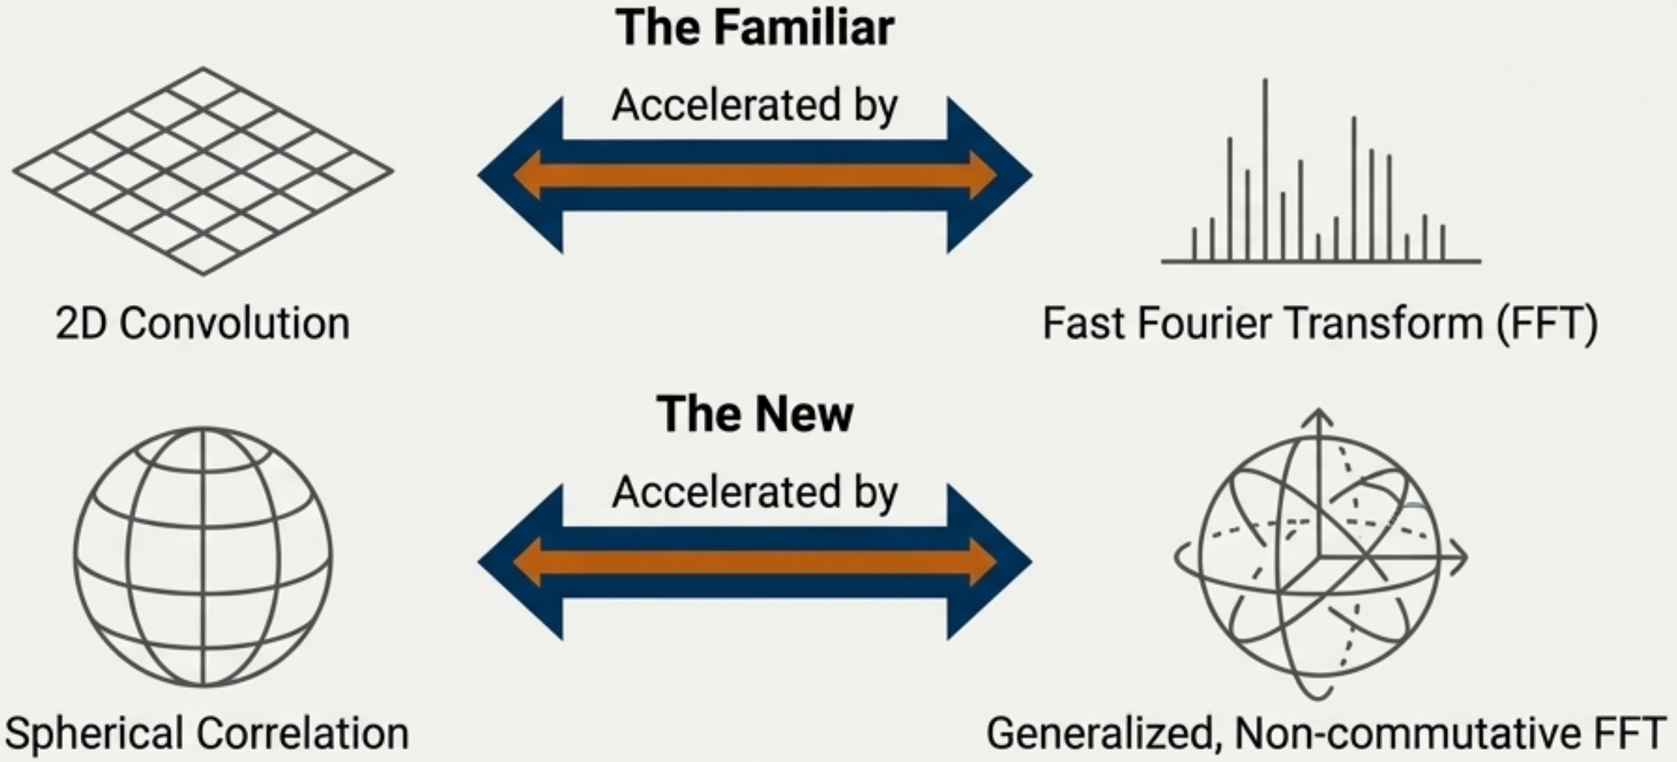
“Graph neural networks (GNNs), based on mathematical graphs, are a natural way to model galaxies, given their ability to model sparsely distributed objects and capture correlations among neighbours.”

- Jagvaral et al.



Overcoming the Complexity Barrier

A key practical challenge is that rotation-equivariant spherical correlations are **computationally intensive** compared to standard 2D convolutions.



The solution is analogous to the standard approach. Just as the Fast Fourier Transform (FFT) dramatically accelerates 2D convolutions, Spherical CNNs can be accelerated by a **generalized FFT**, making them computationally feasible for large-scale practical applications.

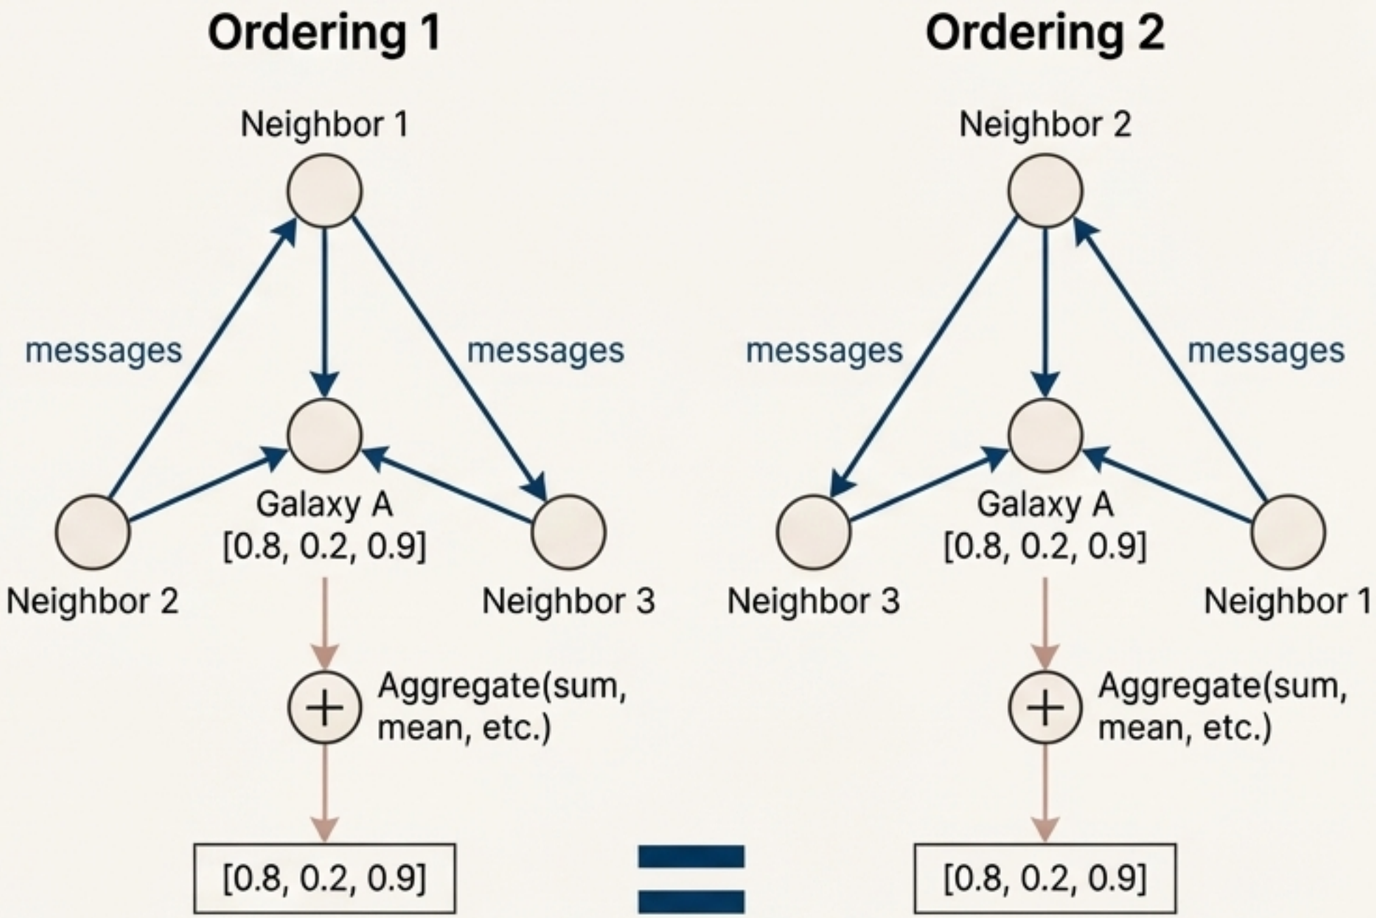
The GNN's Symmetry: Permutation Equivariance

Core Concept

The order in which a node's neighbors are processed does not affect the final output for that node. The aggregation function (sum, mean, max) is inherently invariant to the order of its inputs.

Key Takeaway

This property is what allows a GNN to operate on any arbitrary graph structure, regardless of how we choose to label, order, or draw the nodes and their connections. It learns the structure, not the representation.

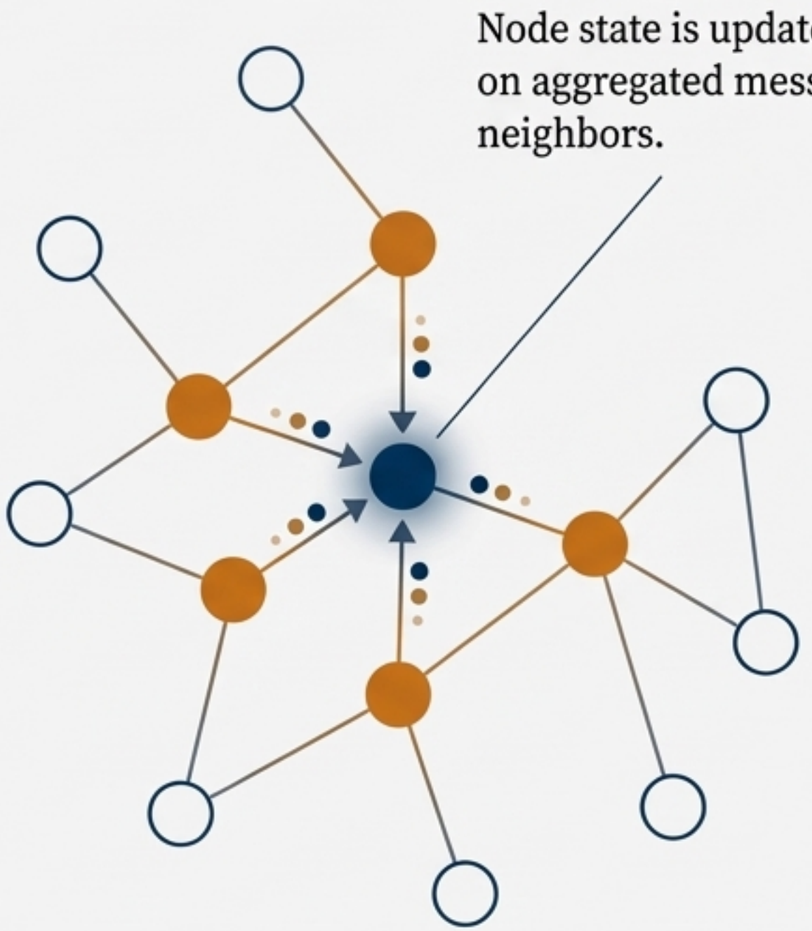


The General Framework: Learning on Any Structure with GNNs

Spherical CNNs are part of a larger field called **Geometric Deep Learning**, which generalizes deep learning to non-Euclidean data like graphs and manifolds.

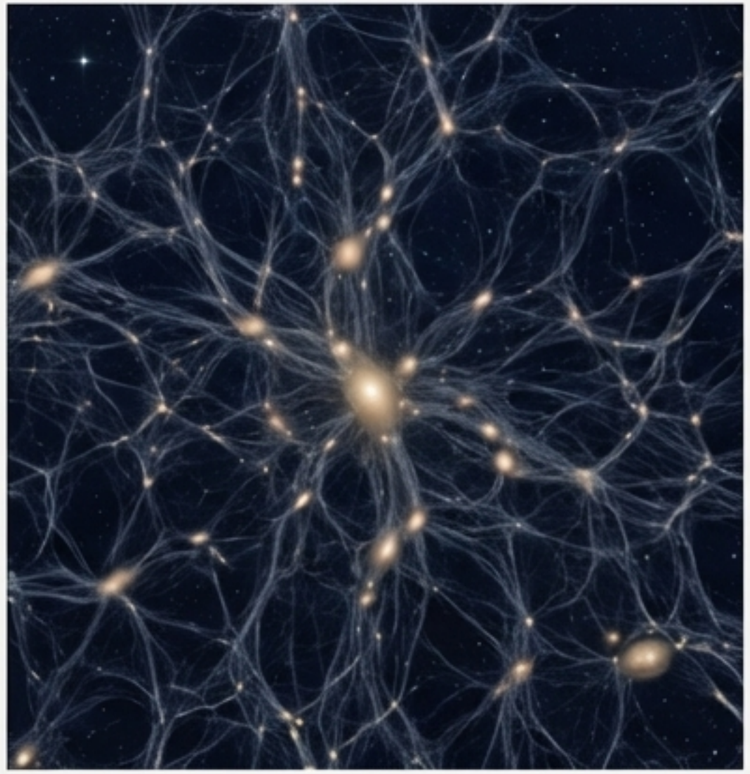
A powerful framework for this is the **Graph Neural Network (GNN).** GNNs operate on a network of *nodes* (objects like galaxies or neurons) connected by *edges* (their relationships). They learn by iteratively passing 'messages' between neighboring nodes.

"GNNs…are a natural way to model galaxies, given their ability to model sparsely distributed objects and capture correlations among neighbours."



From the Cosmos to Cognition: A New Lens on Discovery

By designing models that respect the underlying geometry of data, Spherical CNNs and equivariant GNNs are unlocking new insights in some of humanity's most complex scientific challenges.



ASTROPHYSICS



NEUROSCIENCE



EARTH SCIENCE

Case Study: Respecting the Physics of the Cosmos with E(3) GNNs

The Problem

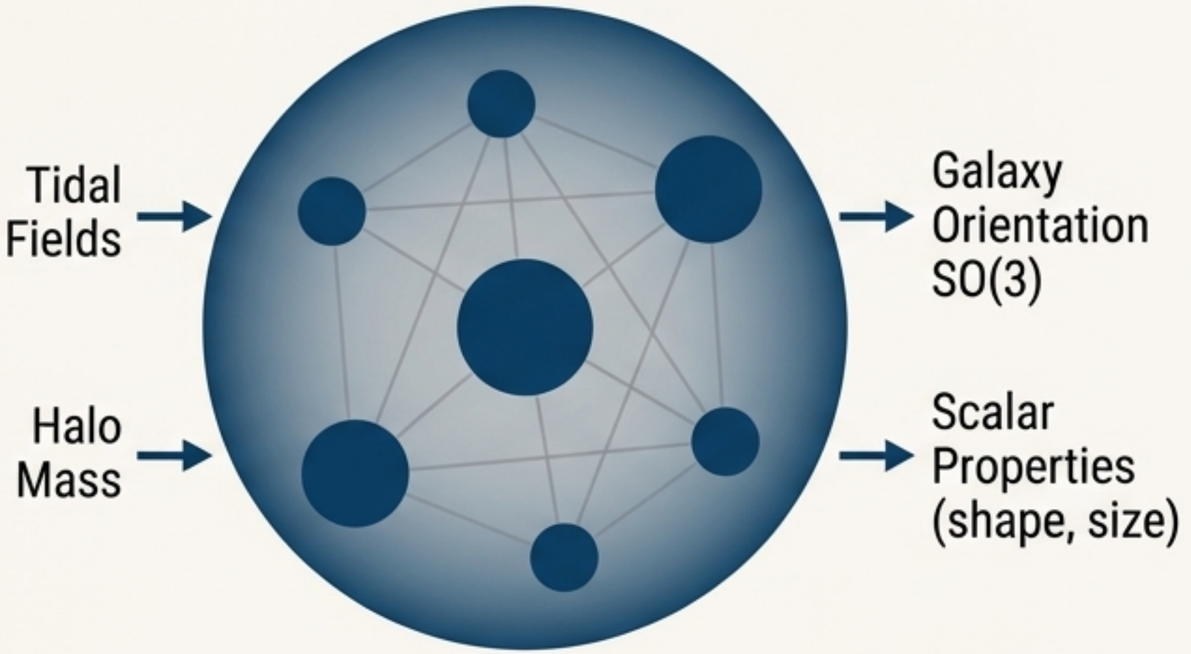
Model the “galaxy-halo connection” to predict galaxy orientation and shape based on its local cosmic environment. This is critical for correcting biases in weak gravitational lensing studies.

Why GNNs?

The connections between galaxies and their surrounding dark matter subhaloes form a graph.

The Critical Insight: E(3) Equivariance

The laws of physics are the same regardless of how you rotate, translate, or reflect the entire system. An E(3) Equivariant GNN builds this fundamental **physical symmetry** directly into its architecture. This is an explicit, analytically imposed symmetry, not one learned through data augmentation.

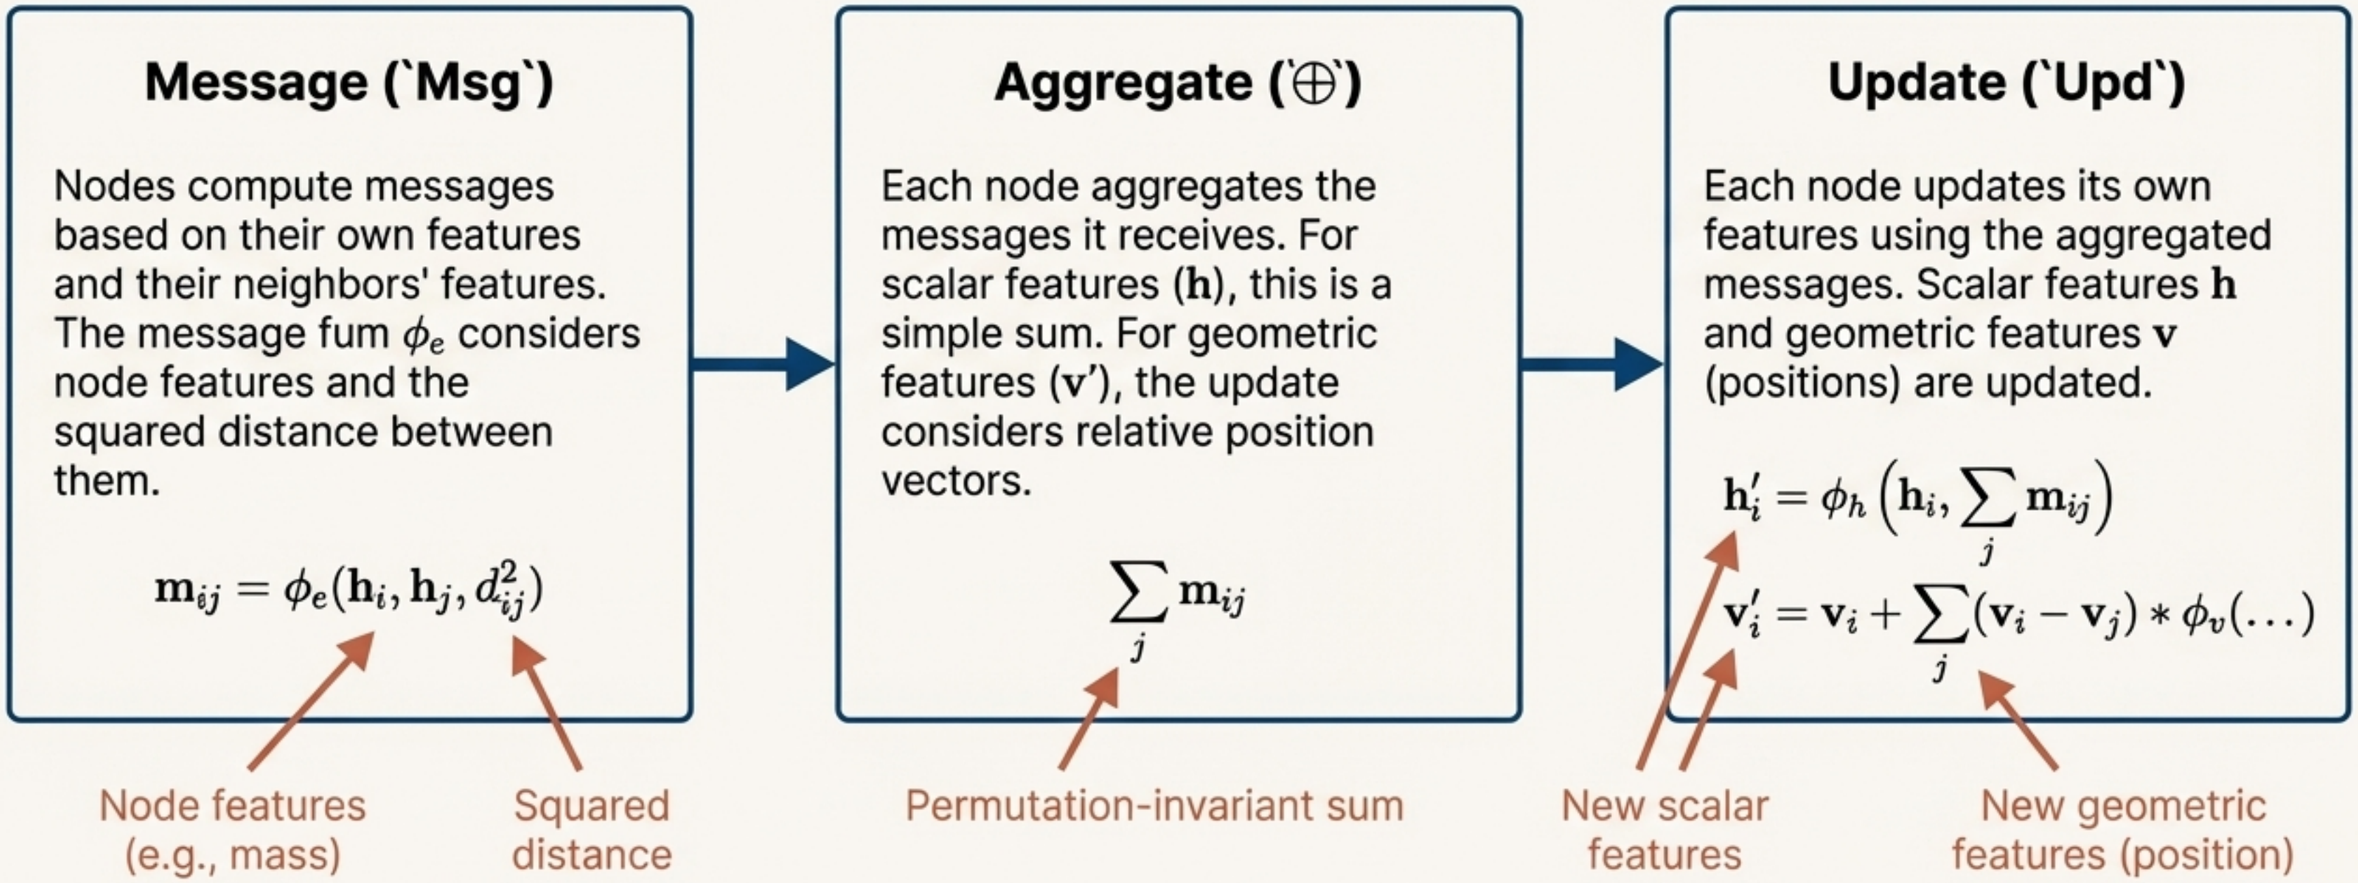




Physics-informed architecture ensures predictions remain consistent under rotation and translation.

How an E(3) GNN ‘Sees’ the Universe

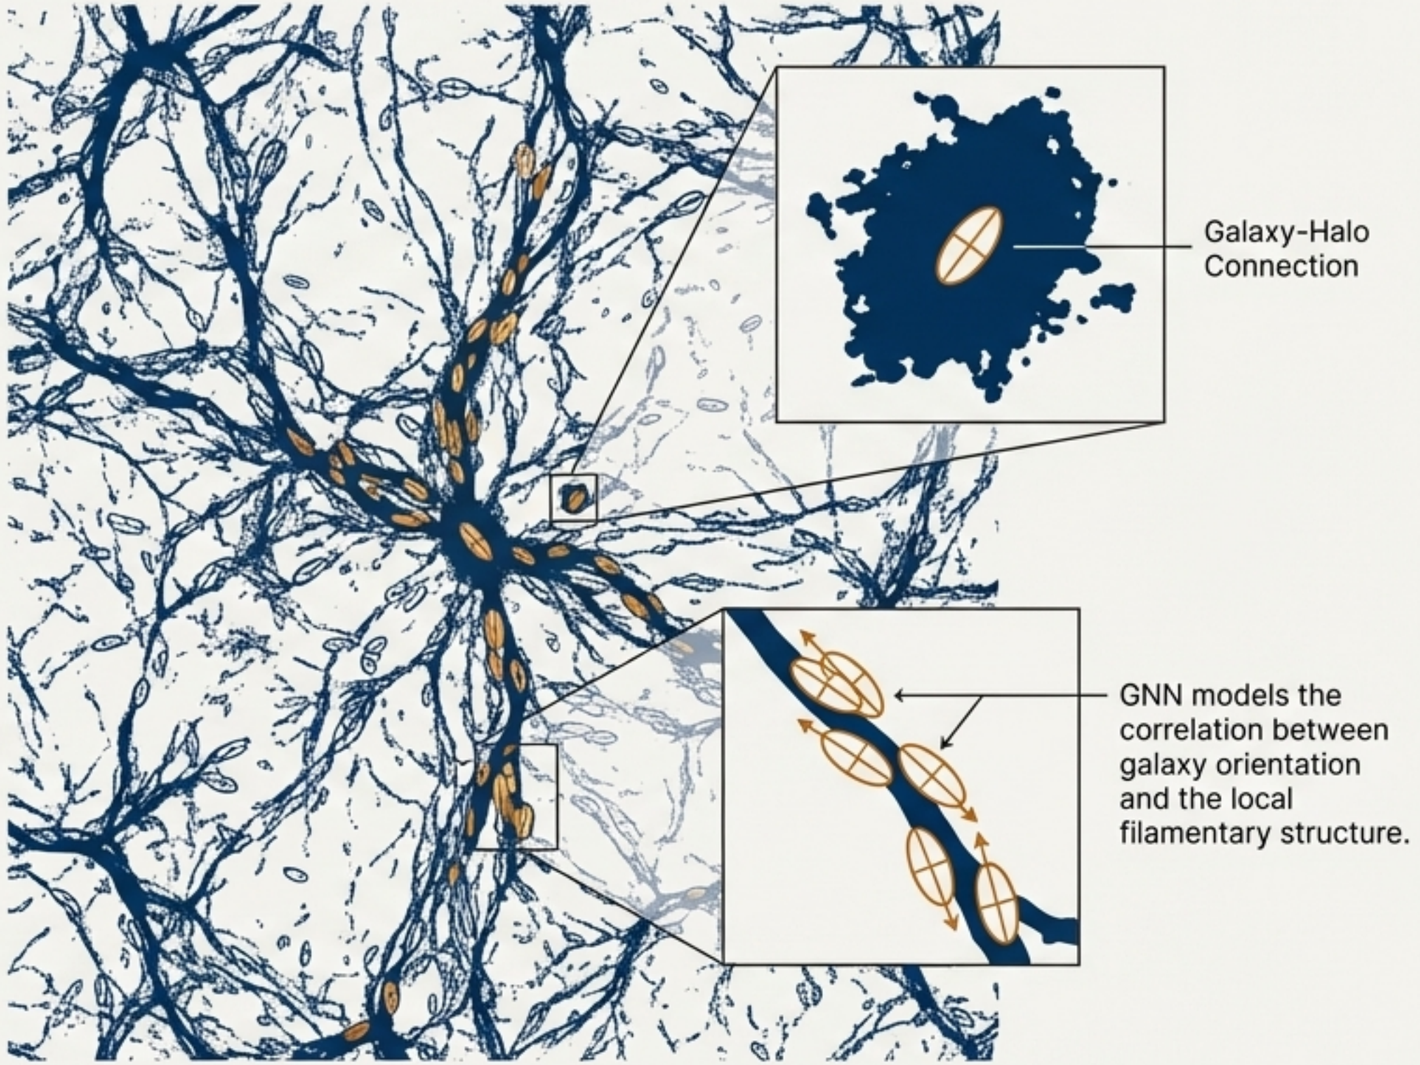
A GNN layer in three steps. Stacking these layers allows information to propagate further across the graph, increasing the model's 'receptive field'.



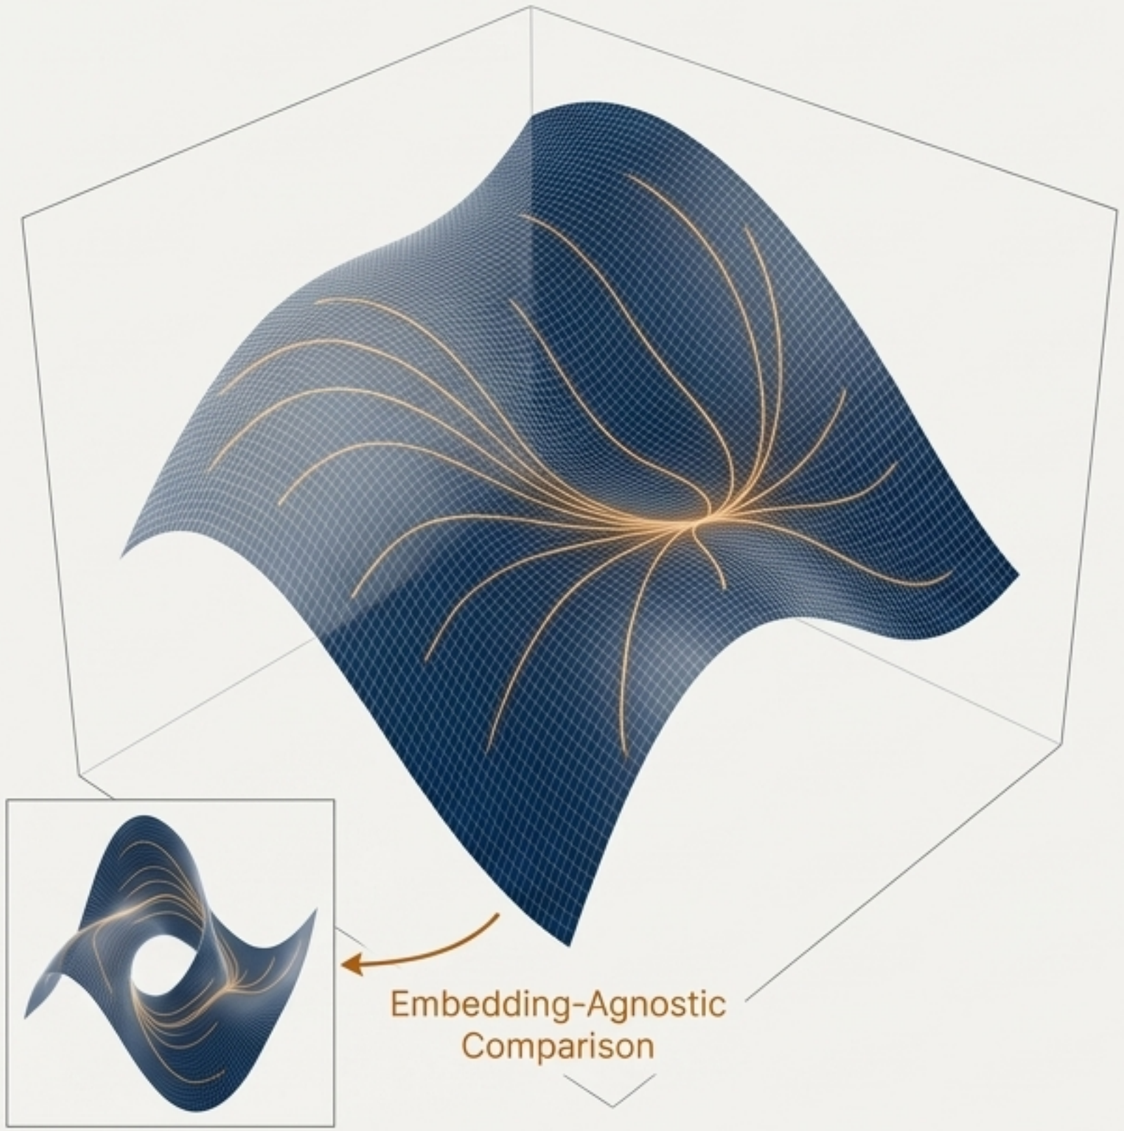
Use Case 1: Mapping the Universe's Hidden Symmetries

The Problem:
The orientation of galaxies is not random. Their **intrinsic alignments (IA)** with the large-scale cosmic structure are a crucial cosmological signal. However, simulating the complex galaxy formation physics responsible for IA across vast cosmic volumes is computationally intractable.

The Solution:
E(3) equivariant GNNs can directly model the 3D positions and orientations of galaxies within their dark matter halos. By respecting the Euclidean symmetries of the universe, these models learn the complex 'galaxy-halo connection' from high-resolution simulations (like IllustrisTNG) and can generate realistic galaxy catalogs for large-volume, dark-matter-only simulations.



Use Case 2: Decoding the Geometry of the Brain



The Problem:
Neural activity during cognitive tasks evolves on complex, low-dimensional "neural manifolds." Different animals might use different neural implementations for the same task, resulting in differently shaped manifolds that are difficult to compare directly.

The Solution:
Geometric deep learning methods (like MARBLE) learn **“embedding-agnostic” representations** of the dynamics on these manifolds. By focusing on the intrinsic local structure, the model finds the common computational process.
This enables robust comparison of cognitive functions across subjects and sessions without supervision from behavioral data.

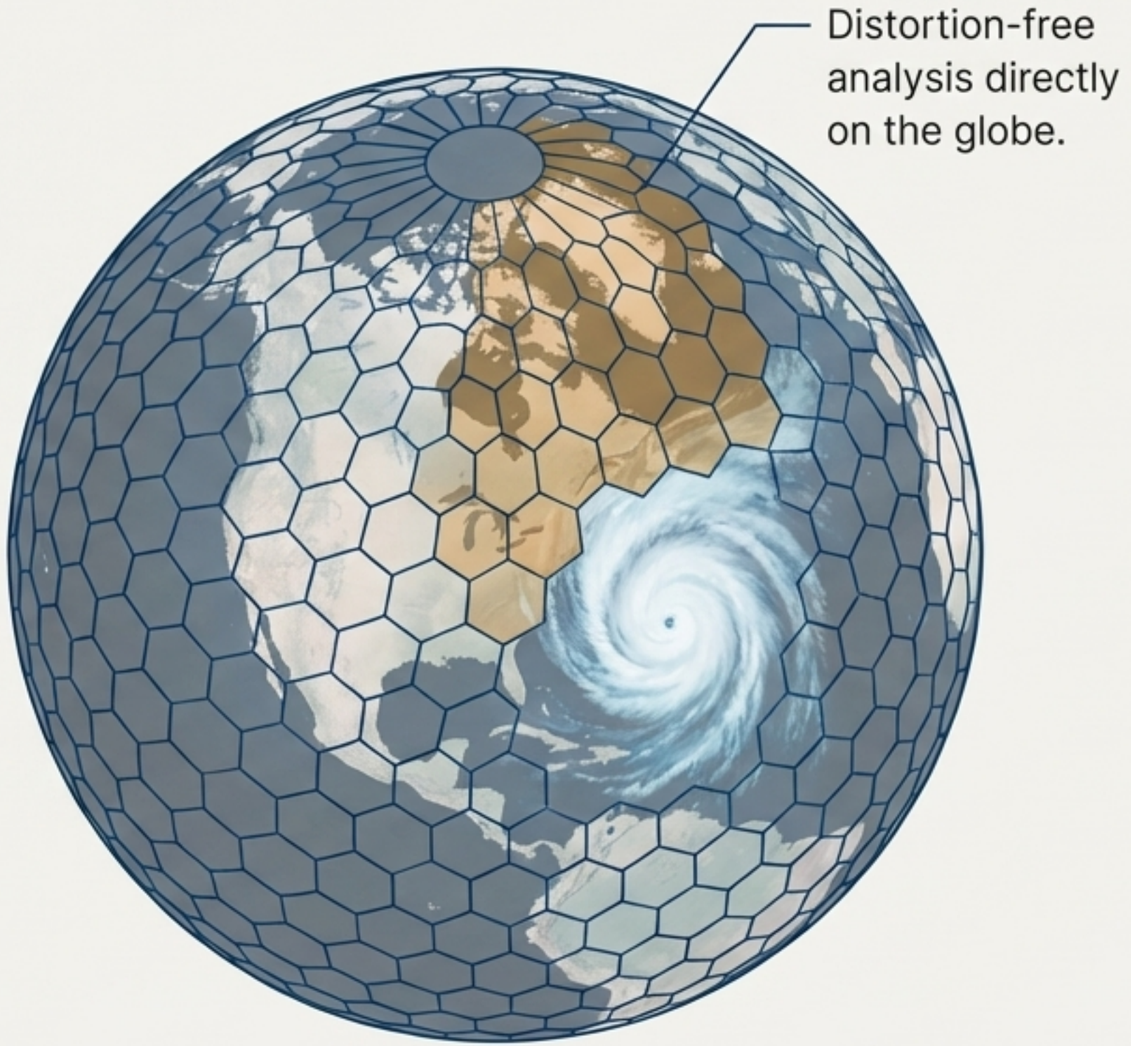
Use Case 3: A Clearer Forecast for a Spherical Earth

The Problem:

Global weather and climate patterns evolve on a sphere. Projecting this data onto flat maps for standard CNNs introduces severe distortions, especially at the poles. This corrupts the representation of weather systems and leads to less accurate forecasts.

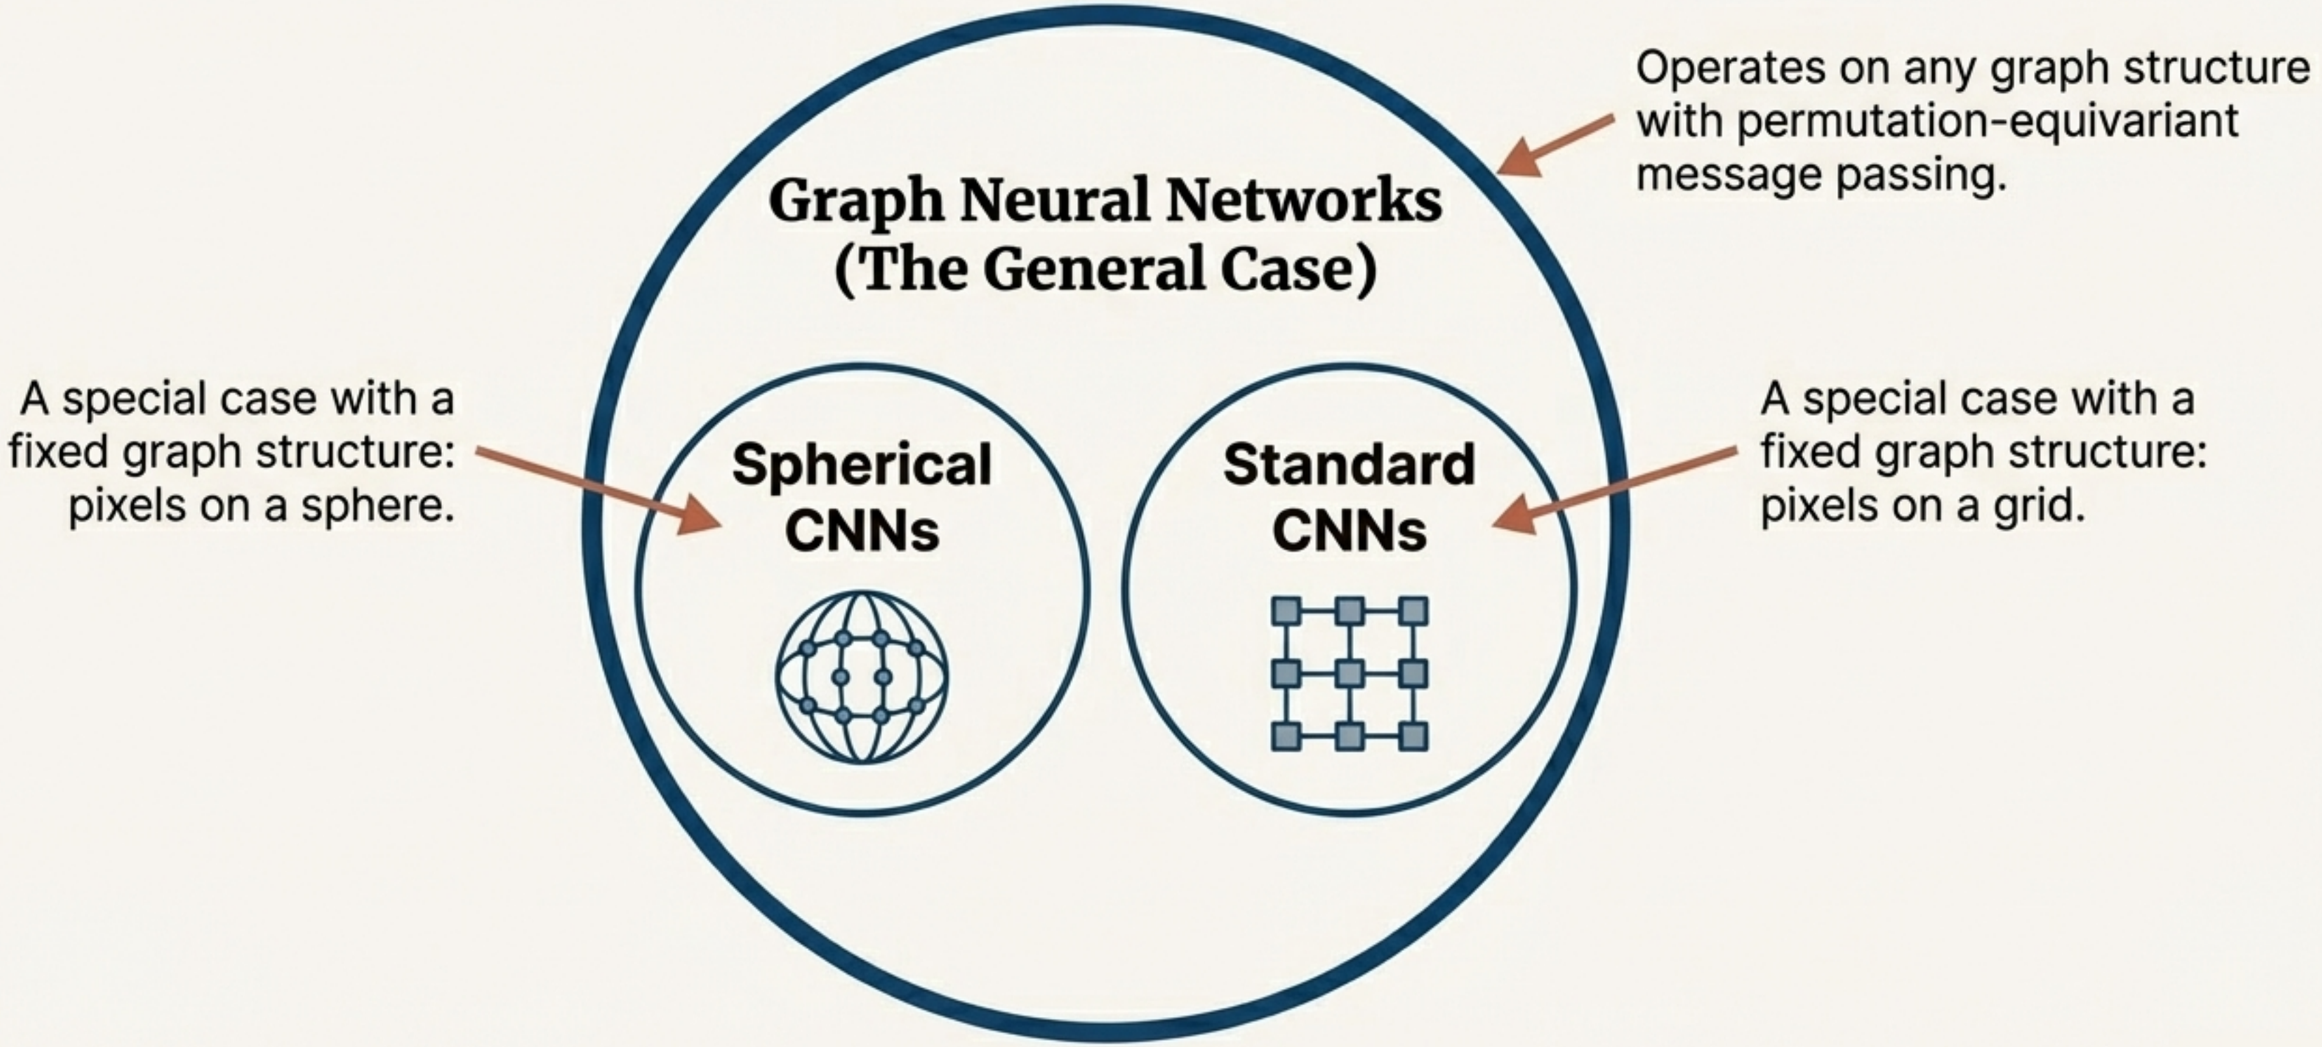
The Solution:

Spherical CNNs process meteorological data directly on the sphere. This preserves the true geometry and spatial relationships of weather systems like hurricanes and jet streams, leading to improved forecasting accuracy and longer lead times for severe weather events.



A Unified View: The Family of Convolutions

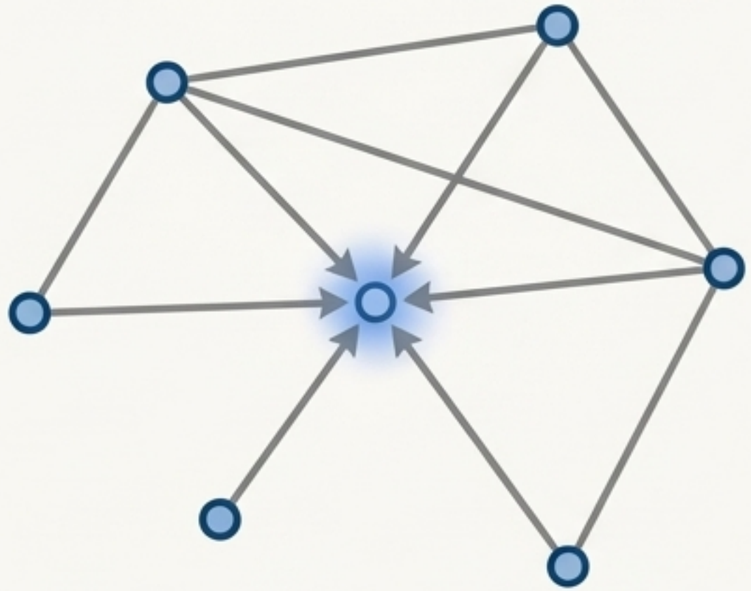
Core Insight: GNNs are the most general form of convolution. Other architectures are special cases optimized for regular data structures.



The Other Path: What About Transformers?

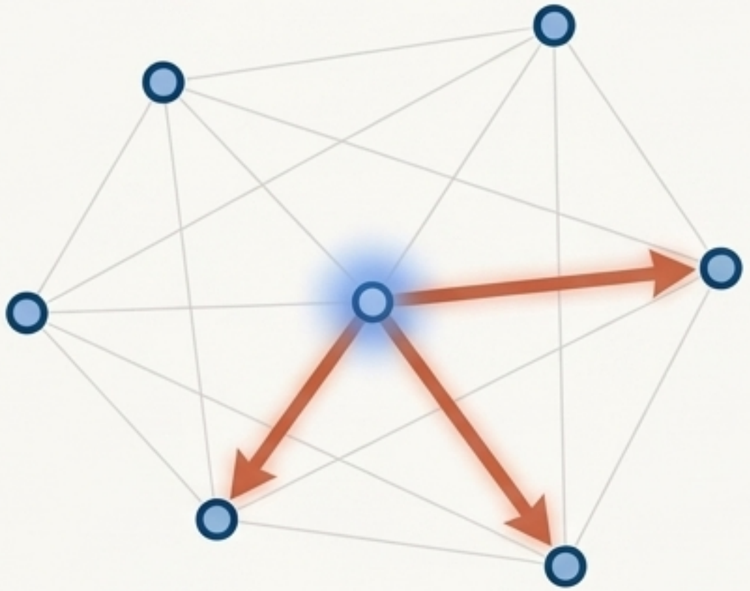
Core argument: A Transformer's self-attention mechanism can be viewed as operating on a dense, implicit graph where every node is connected to every other. The attention mechanism dynamically learns the edge weights—it learns the graph structure from the data itself.

Graph Neural Network



Explicit & Sparse Structure.
Relies on a predefined graph.

Transformer



Implicit & Dense Structure.
Attention learns the important connections.

Head-to-Head: Structure vs. Attention

Feature	Graph Neural Network (GNN)	Transformer
Graph Structure	**Explicit & Sparse.** A crucial inductive bias provided to the model.	**Implicit & Dense.** Assumes a fully-connected graph; attention learns the edge weights.
Computational Cost	Scales with number of edges **O(E)**. Efficient for sparse natural graphs.	Scales quadratically with number of nodes **$O(N^2)$**. Can be prohibitive for large N.
Message Passing	Explicit aggregation from **predefined**, **local** neighbors.	A weighted sum of values from **all other nodes**, with weights determined by attention.
When to Use	When a meaningful, **known sparse structure** exists (molecules, physical systems, social networks, cosmic web).	When relationships are **not known a priori** and long-range dependencies are key (language, vision).

The Modern Architecture Spectrum

	CNN	GNN	Transformer
Data Structure	Regular Grid	Arbitrary Graph	Sequence / Set (Complete Graph)
Neighborhood	Local, Fixed, & Shared	Explicit, Variable, & Sparse	Global, Dynamic, & Learned
Key Symmetry	Translation	Permutation (or other groups like E(3))	Permutation
Inductive Bias	**Strong:** Locality & spatial hierarchy built-in.	**Strong:** Relational structure is built-in.	**Weak:** Attention can learn any structure.
Ideal For...	Images, Video	Molecules, Networks, Physical Systems	Text, Vision (ViT), General Sequences

From Grids to Geometry: Final Takeaways

1. Architecture Must Follow Data Geometry.

The most powerful and efficient models are those whose architectures explicitly respect the intrinsic **symmetries** of the data—be it **translation** for grids, **rotation** for spheres, or the **relational structure** of graphs.

2. GNNs Represent a Fundamental Generalization.

Standard and Spherical CNNs are not distinct technologies but can be understood as **highly specialized GNNs** designed for the fixed, regular graphs of grids and spheres.

3. The Core Divide is Structure vs. Attention.

The choice between modern **GNNs** and **Transformers** is a trade-off. **GNNs** leverage known, **explicit structure** for efficiency and **strong inductive bias**. **Transformers** forgo this bias to **dynamically learn implicit relationships** from data, at a higher computational cost.

The Ecosystem: Getting Started with PyTorch Geometric

A rich ecosystem of open-source tools makes Geometric Deep Learning accessible for research and application.

Featured Library: PyTorch Geometric (PyG)

A library built on PyTorch to easily write and train Graph Neural Networks.

- **Comprehensive**: Includes a vast collection of state-of-the-art GNN models, layers, and operators.
- **Scalable**: Mini-batch loaders and multi-GPU support for operating on giant graphs with millions of nodes.
- **Flexible**: An easy-to-use message passing API makes it simple to design and experiment with new architectures.

PyG



Example: A 2-Layer GCN in PyG

```
import torch
from torch_geometric.nn import GCNConv
from torch_geometric.datasets import Planetoid

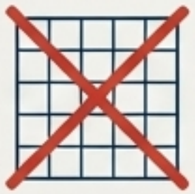
dataset = Planetoid(root='.', name='Cora')

class GCN(torch.nn.Module):
    def __init__(self):
        super().__init__()
        self.conv1 = GCNConv(dataset.num_features, 16)
        self.conv2 = GCNConv(16, dataset.num_classes)

    def forward(self, data):
        x, edge_index = data.x, data.edge_index
```

A New Foundation for AI Perception

Key Takeaways



THE PROBLEM

Flat projections distort data, breaking standard CNNs.



THE INSIGHT

Equivariance provides the guiding principle: respect the data's native geometry.



THE SOLUTION

Spherical CNNs & Equivariant GNNs build symmetry into the architecture.



THE IMPACT

Unlocking transformative applications across fundamental science.

The Road Ahead

Geometric Deep Learning is a rapidly expanding field, poised to become a foundational tool for AI in science, robotics, and beyond.

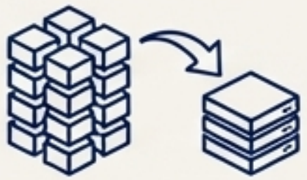
The Road Ahead

Geometric Deep Learning is a rapidly expanding field, poised to become a foundational tool for AI in science, robotics, and beyond.

The Power of Priors: Smarter Architectures, Not Just Bigger Data

The success of Geometric Deep Learning demonstrates a fundamental principle for the future of AI.

By building physical and geometric symmetries (inductive biases) directly into our network architectures, we create models that are:



More Data-Efficient: They don't need to learn fundamental laws of the universe from scratch.

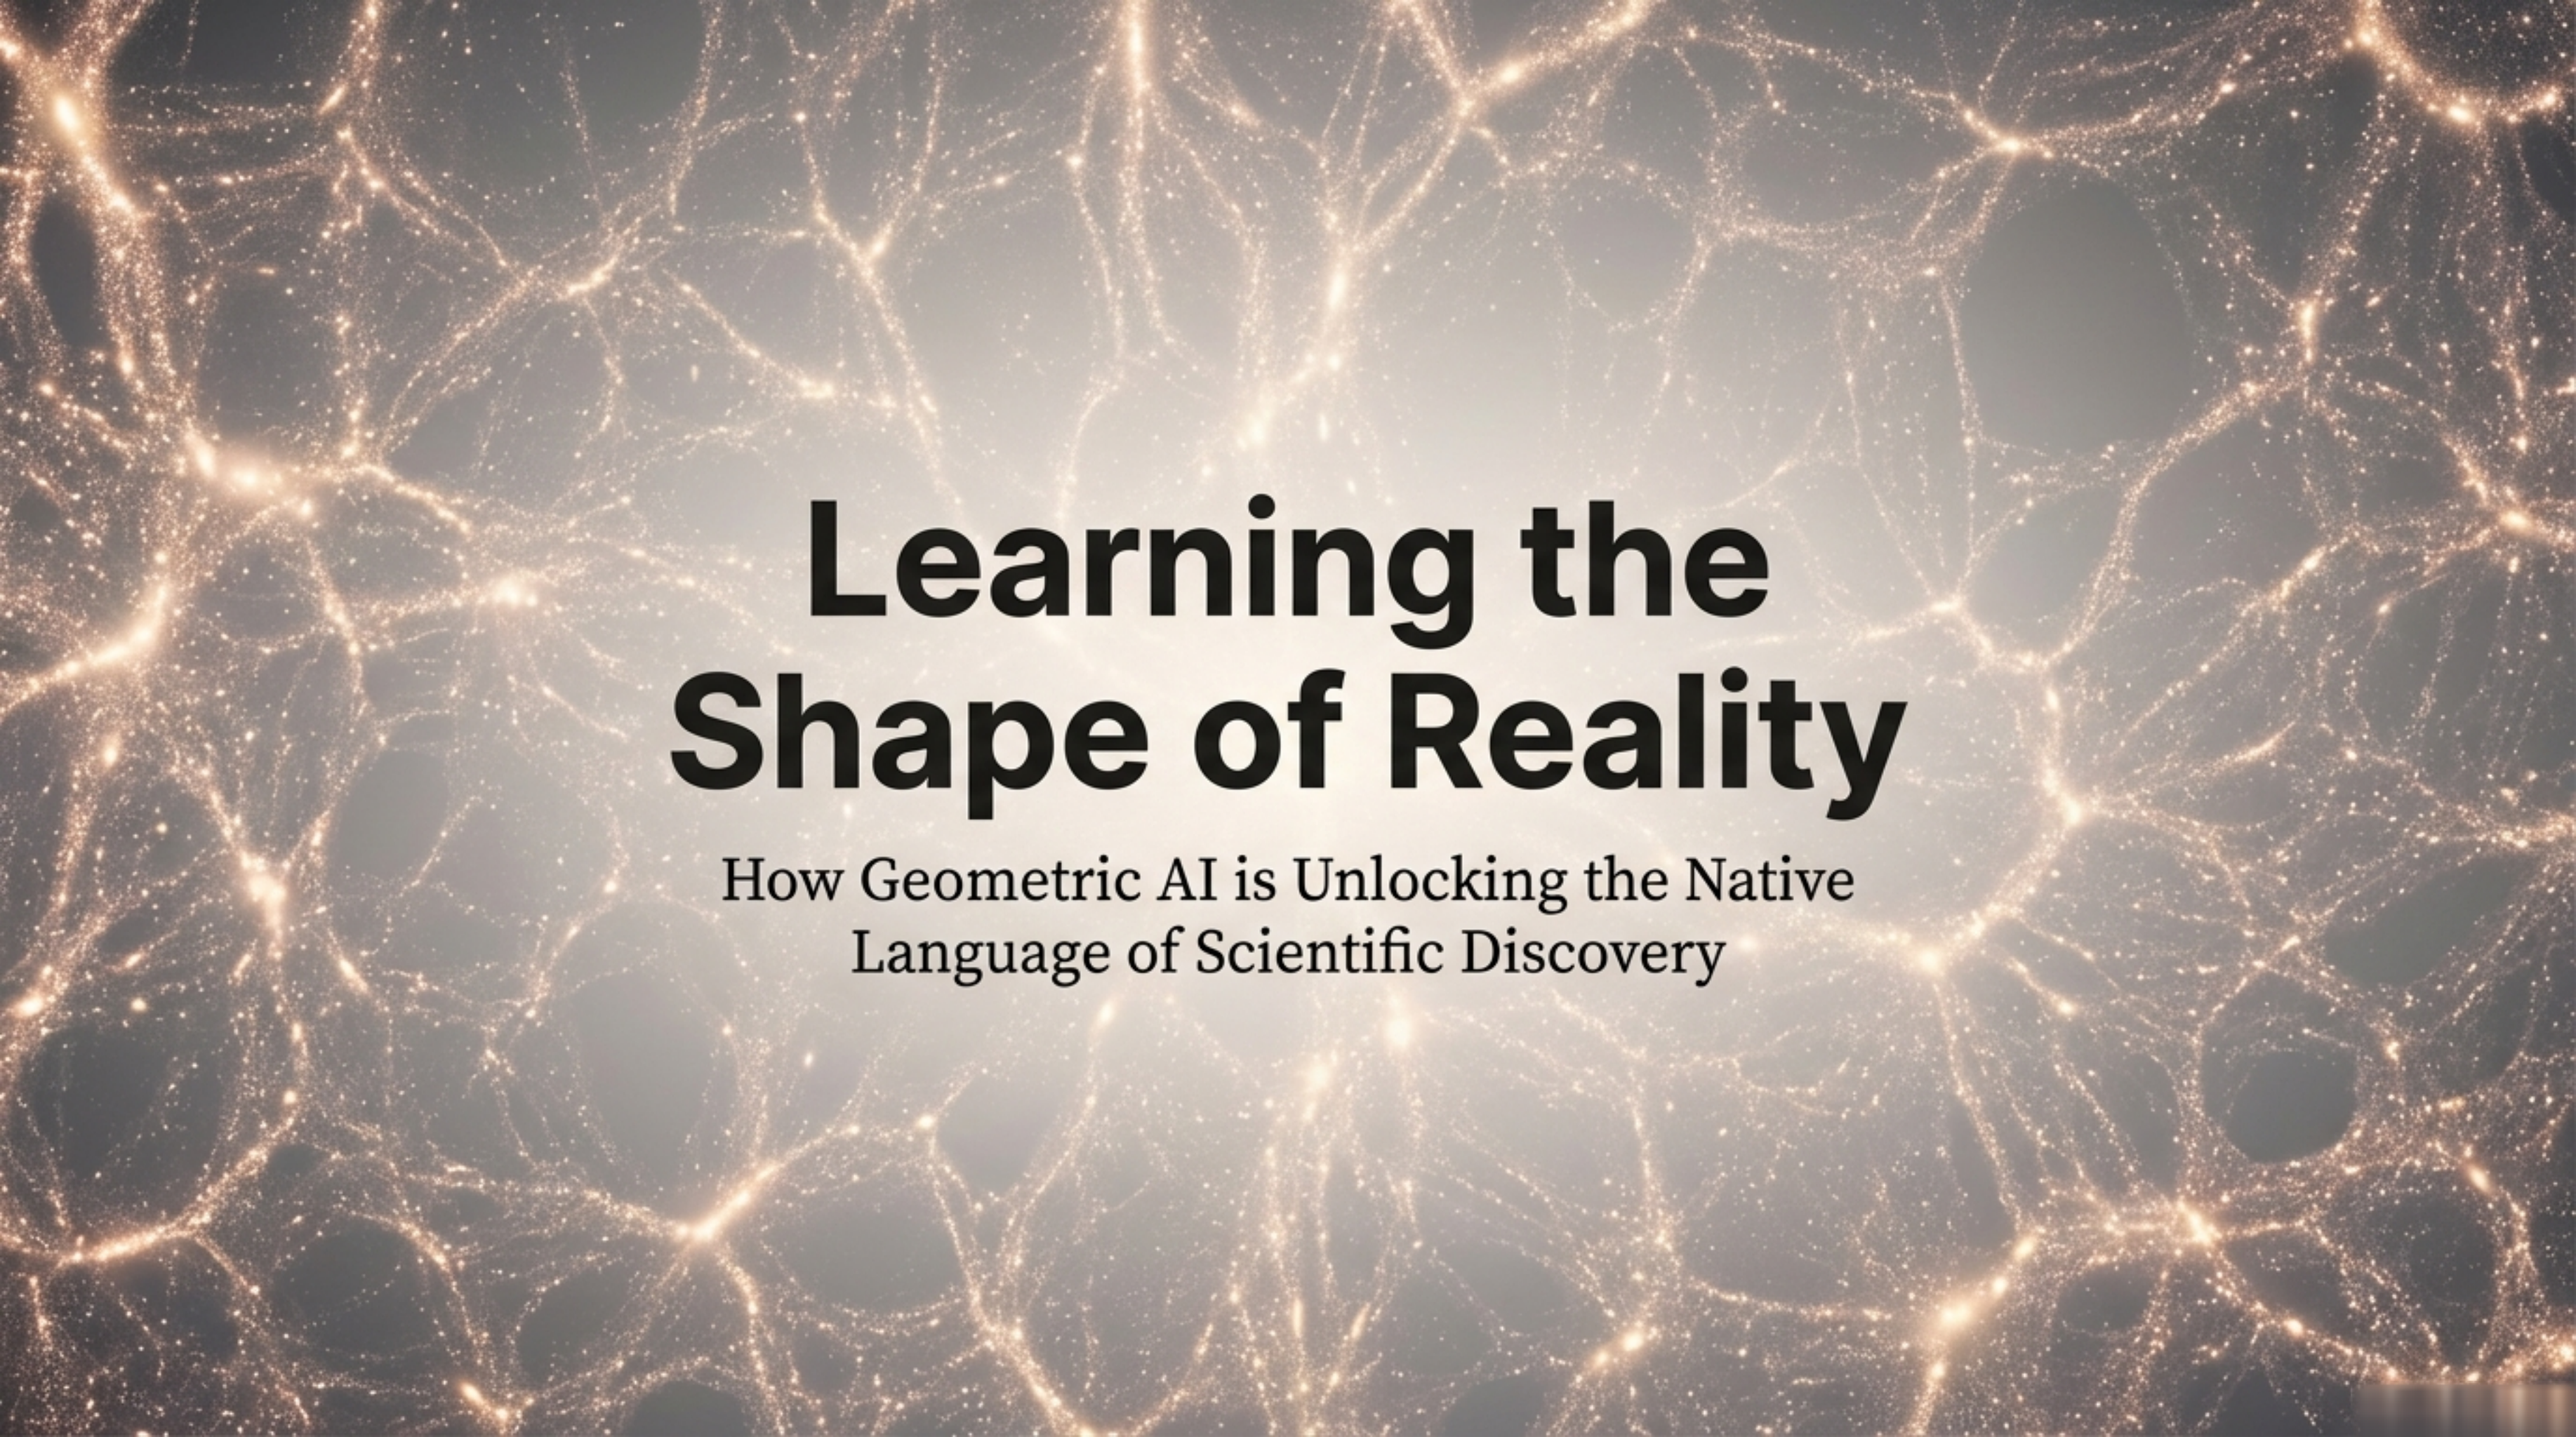


More Robust & Generalizable: They are guaranteed to respect the symmetries of the data domain.



Higher Performing: They achieve state-of-the-art results on complex data where traditional methods fail.

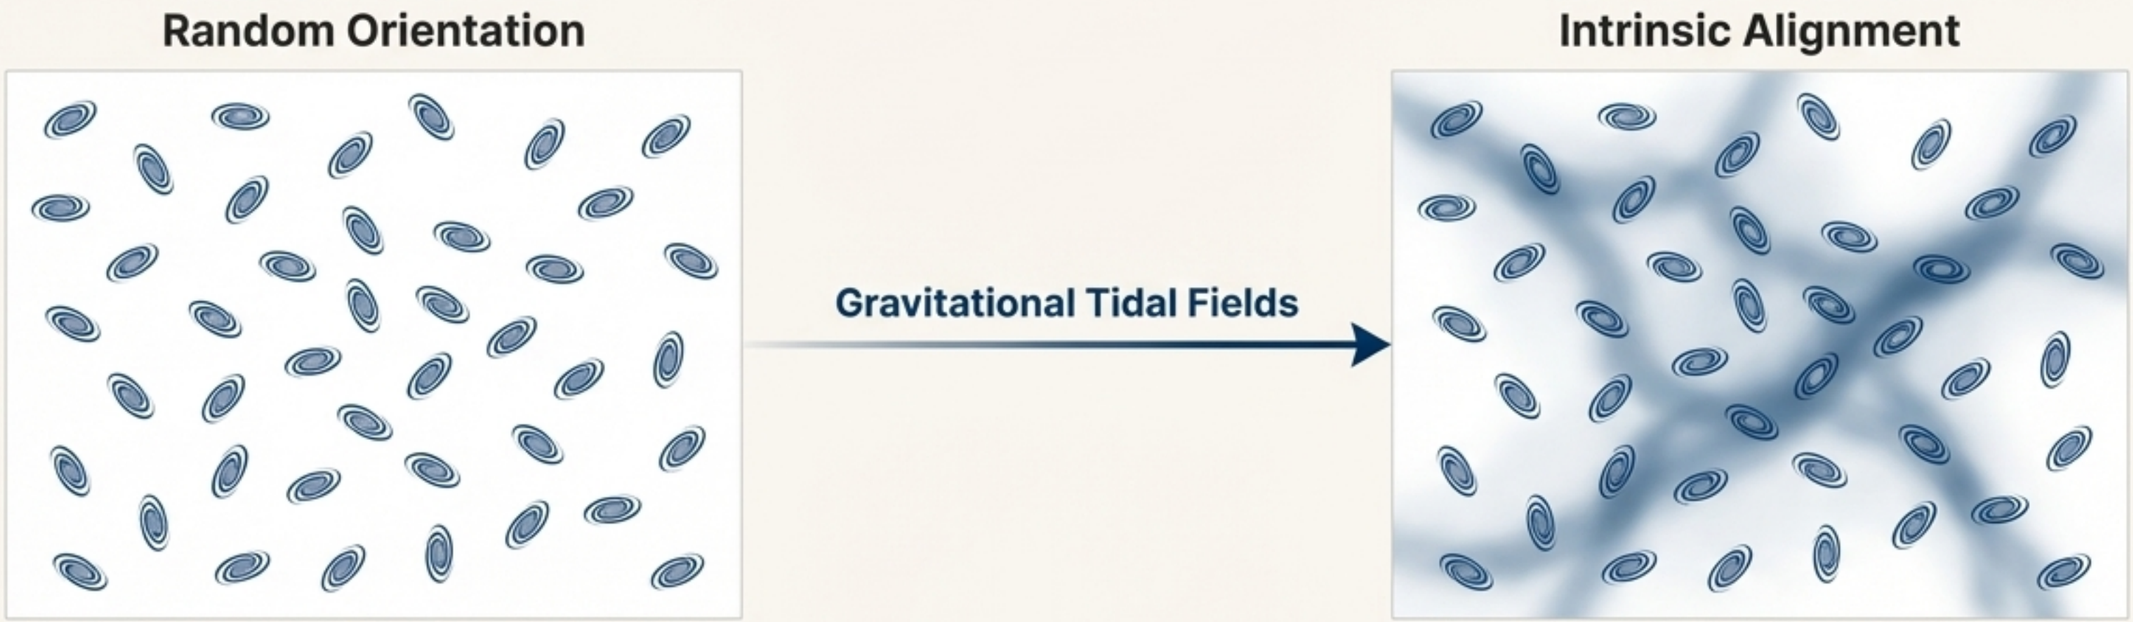
Concluding Thought: This represents a shift from a purely data-driven paradigm to a knowledge-driven one, where principled design guided by the structure of the physical world leads to the next generation of intelligent systems.


Learning the
Shape of Reality
How Geometric AI is Unlocking the Native
Language of Scientific Discovery

The Next Generation of Cosmic Surveys Demands a New Approach

- Surveys like LSST and Euclid will generate unprecedented volumes of data to probe the nature of dark matter and dark energy.
- Our ability to extract cosmological insights is limited by our ability to model subtle, systematic effects.
- One of the most critical systematics is a phenomenon called Intrinsic Alignments.





The Cosmic Contaminant: Intrinsic Alignments

The Problem:

Galaxies don't orient randomly. They physically align with the large-scale structure of the universe.

The Consequence:

This physical alignment mimics the signal of gravitational lensing, the very tool we use to map dark matter.

The Bias:

As stated by Jagvaral et al., this effect "can falsely imitate weak gravitational lensing signals and bias cosmological analyses."

Impact in Astrophysics: Modeling the Galaxy-Halo Connection

Source: Jagvaral et al., *Monthly Notices of the Royal Astronomical Society*, 2025

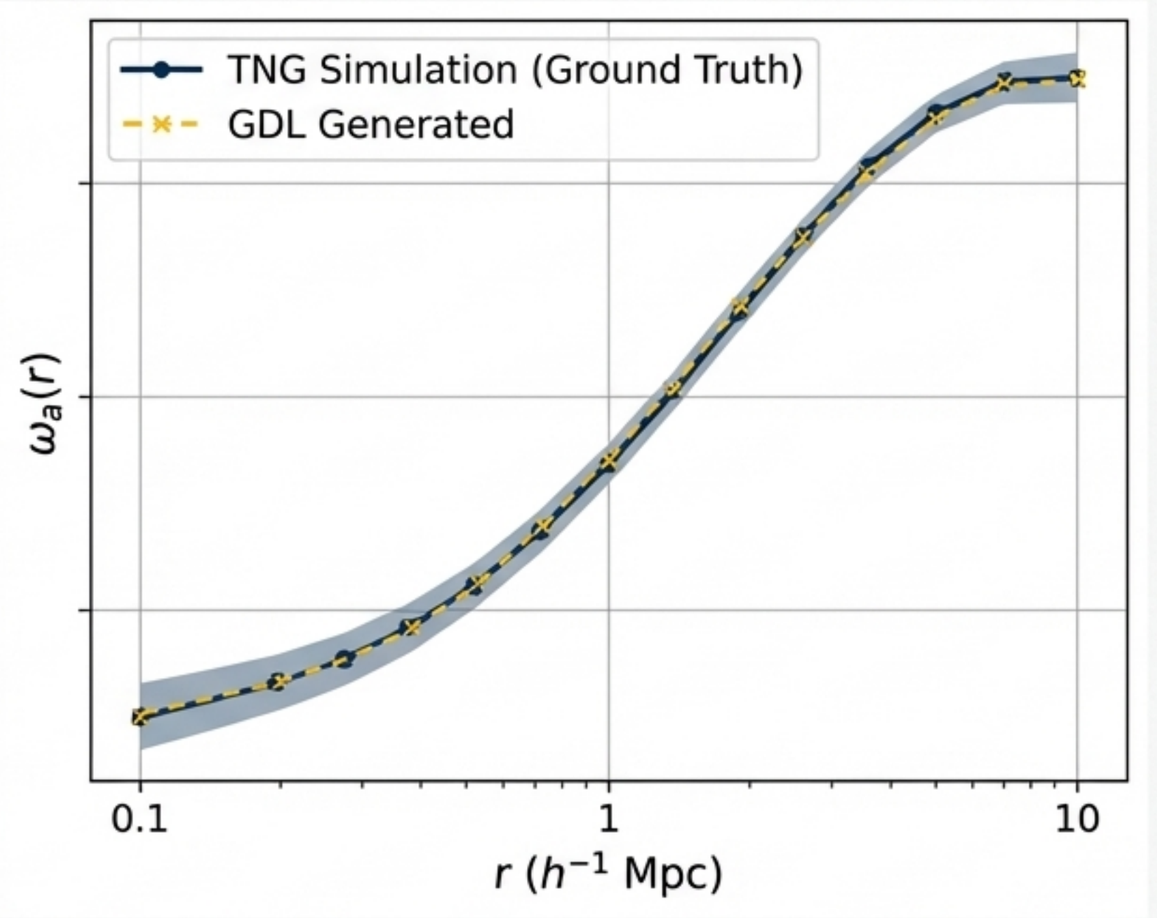
Challenge

Galaxy orientations, or "intrinsic alignments" (IA), are a major source of systematic bias in weak gravitational lensing surveys like the Rubin Observatory LSST. Simulating the complex baryonic physics that drives these alignments across vast cosmological volumes is computationally intractable with traditional methods.

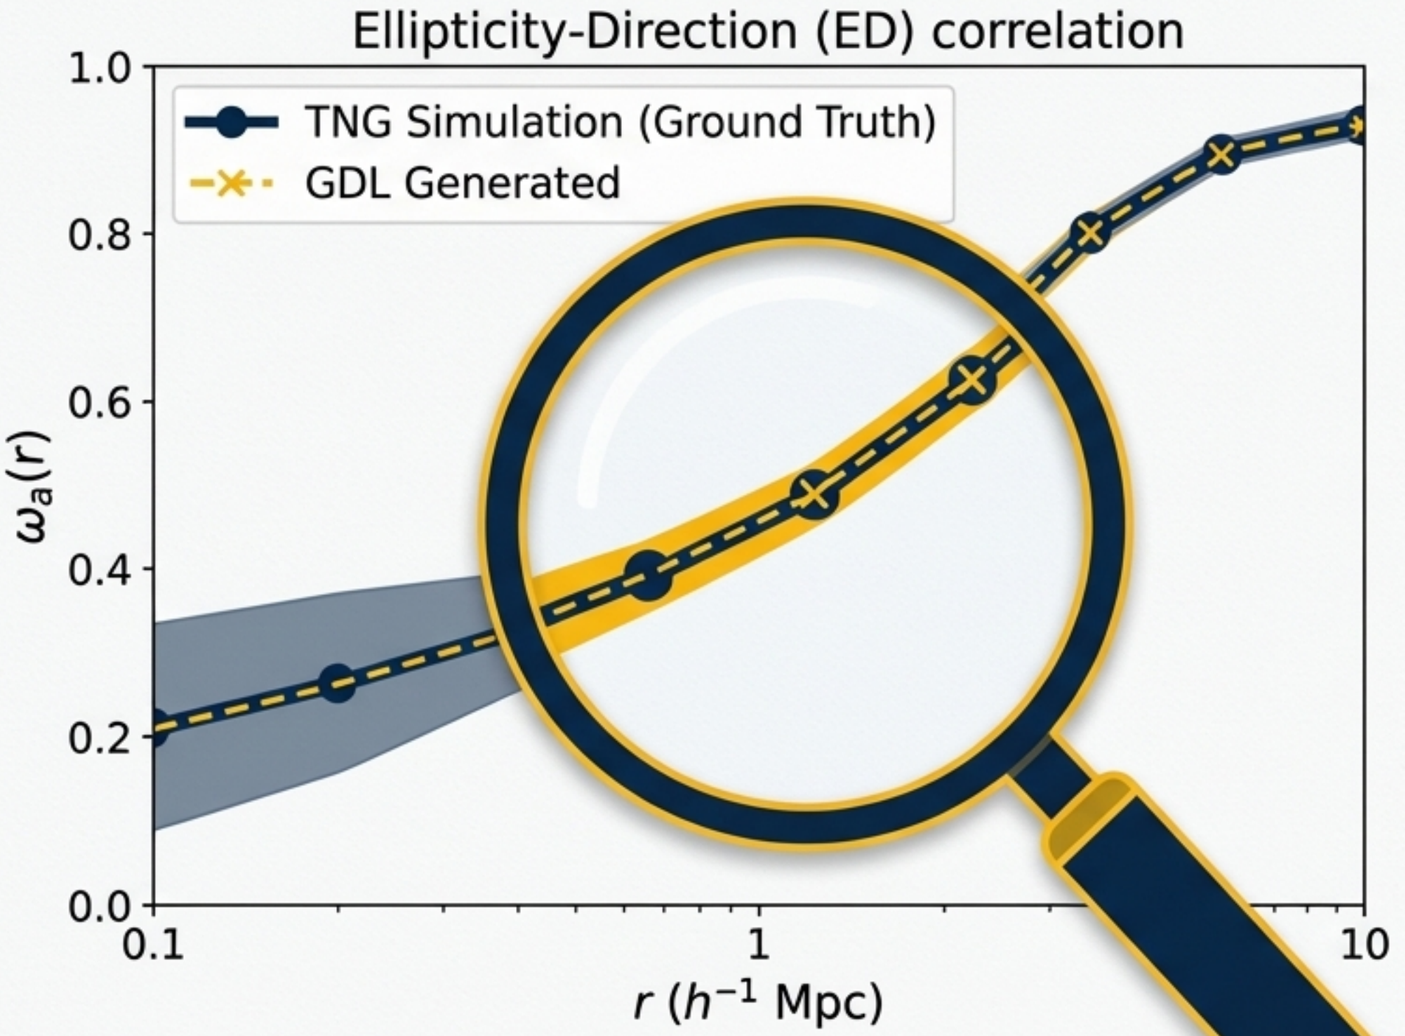
GDL Method

An E(3) equivariant GNN is combined with a diffusion generative model on the SO(3) manifold. The model learns to "paint" realistic 3D galaxy shapes and orientations onto dark matter halos from the IllustrisTNG-100 simulation, conditioned on the local tidal field.

Ellipticity-Direction (ED) Correlation



Result: A Near-Perfect Match at a Fraction of the Cost



The generative model reproduces the target joint distributions with high statistical agreement. The generated and true values for geometric quantities are statistically consistent across all scales.

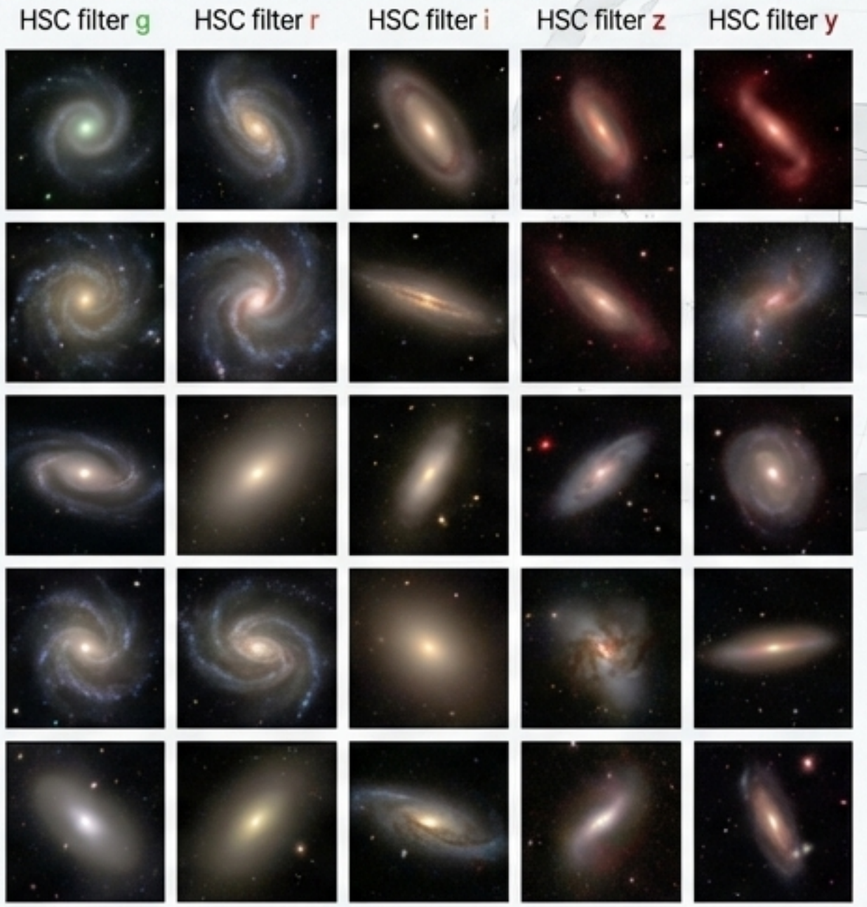
Quantitative Proof (from Jagvaral et al., Table 2)

- Wasserstein-1 distances show agreement of the probability distributions to better than **1 percent**.
- The ratio of means and variances between generated and true scalar properties are all **~1.00**.

Impact

This provides a robust and computationally efficient emulator for producing synthetic galaxy catalogues, essential for validating analysis pipelines for next-generation surveys like LSST and Euclid.

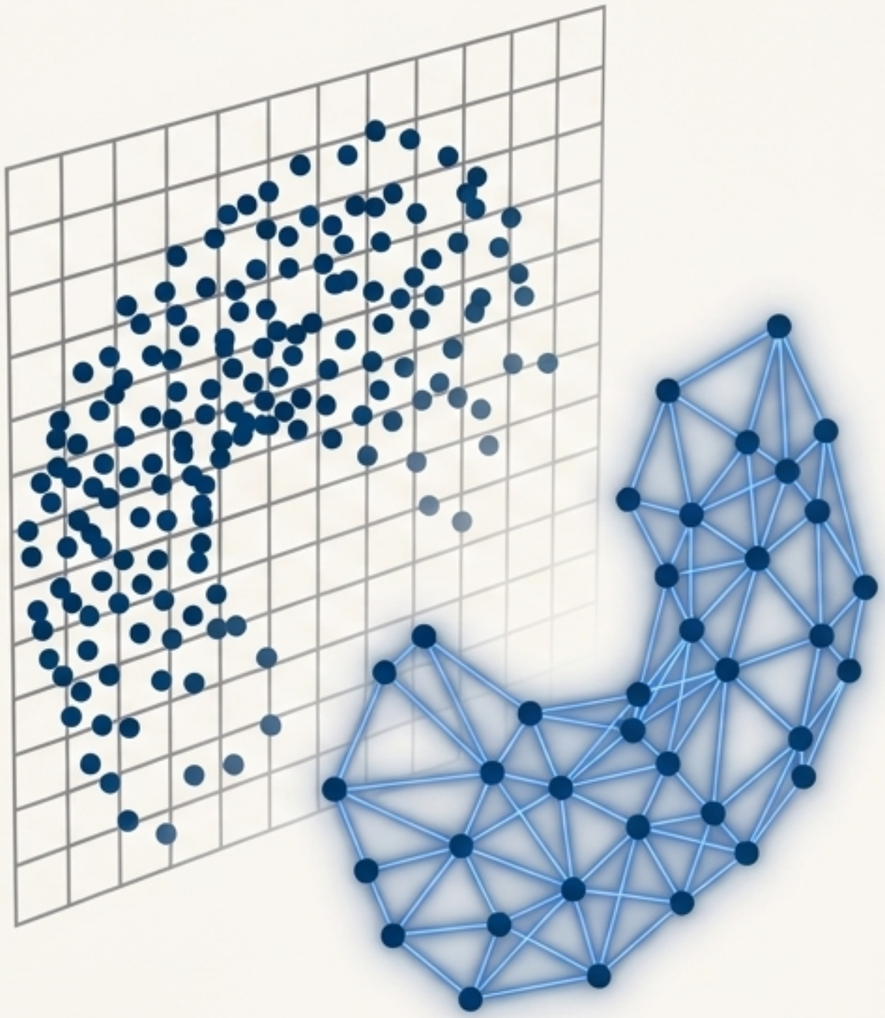
The Enabling Ecosystem: Fueling Discovery with ML-Ready Data



These advanced models are only possible with large-scale, meticulously curated datasets. The GalaxiesML dataset is a prime example of this critical infrastructure.

- **286,401** galaxy images from the Hyper-Suprime-Cam (HSC) Survey.
- Spectroscopically confirmed **redshifts** (z=0.01 to 4) as ground truth.
- Rich morphological parameters: **Sérsic index**, **ellipticity**, **half-light radius**.
- Crucially supports upcoming surveys like **LSST and Euclid**.

This resource empowers astrophysicists and data scientists to train and validate the next generation of cosmological models.

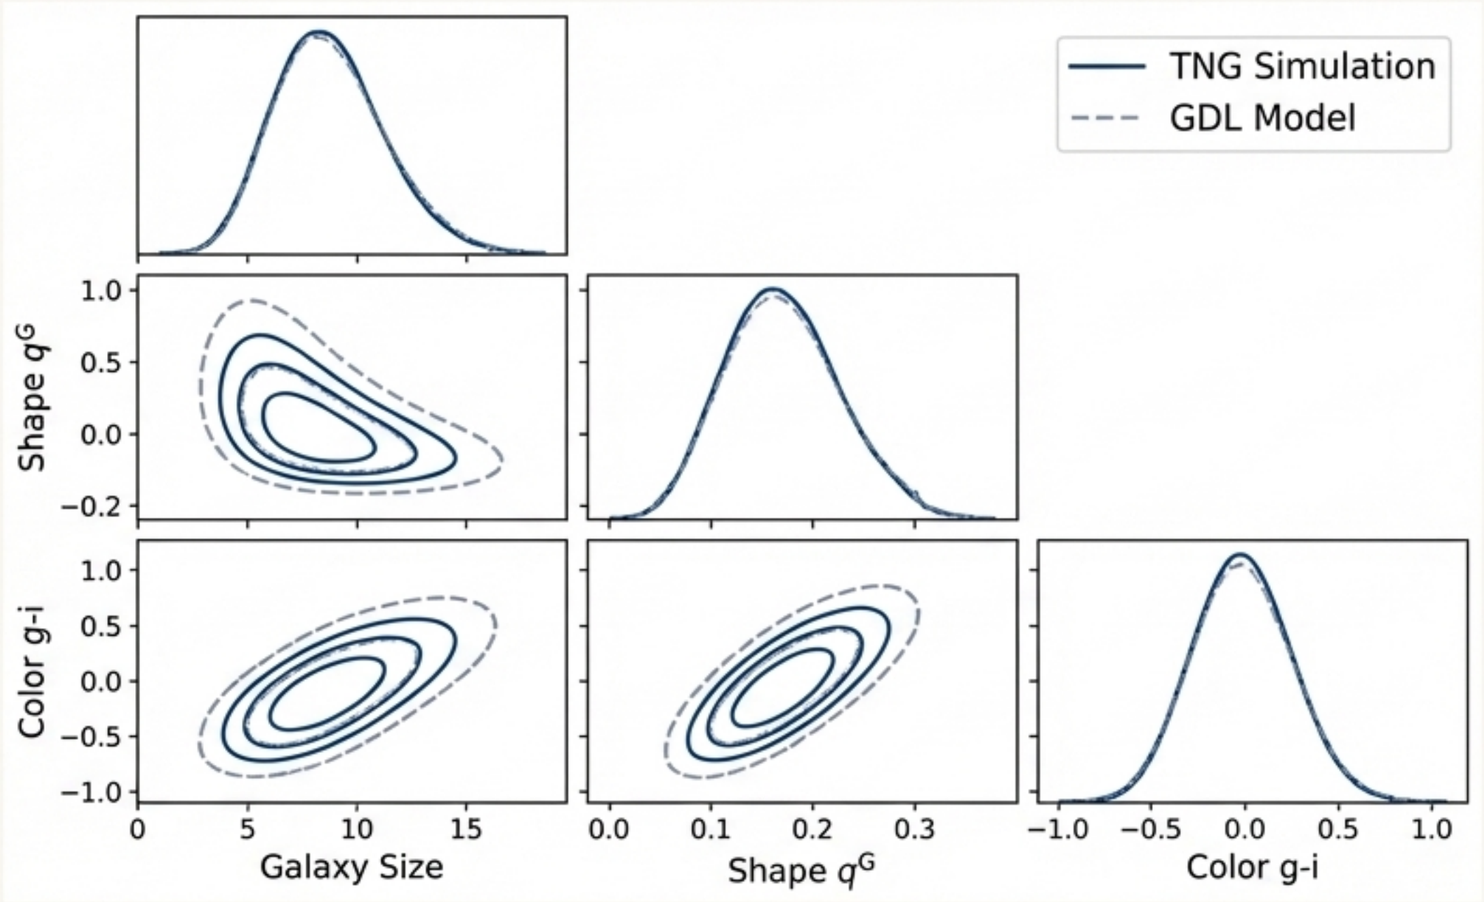


A New Paradigm: Geometric Deep Learning

Instead of forcing data onto a grid, GDL models learn on the data's native structure—graphs, manifolds, and other geometric objects.

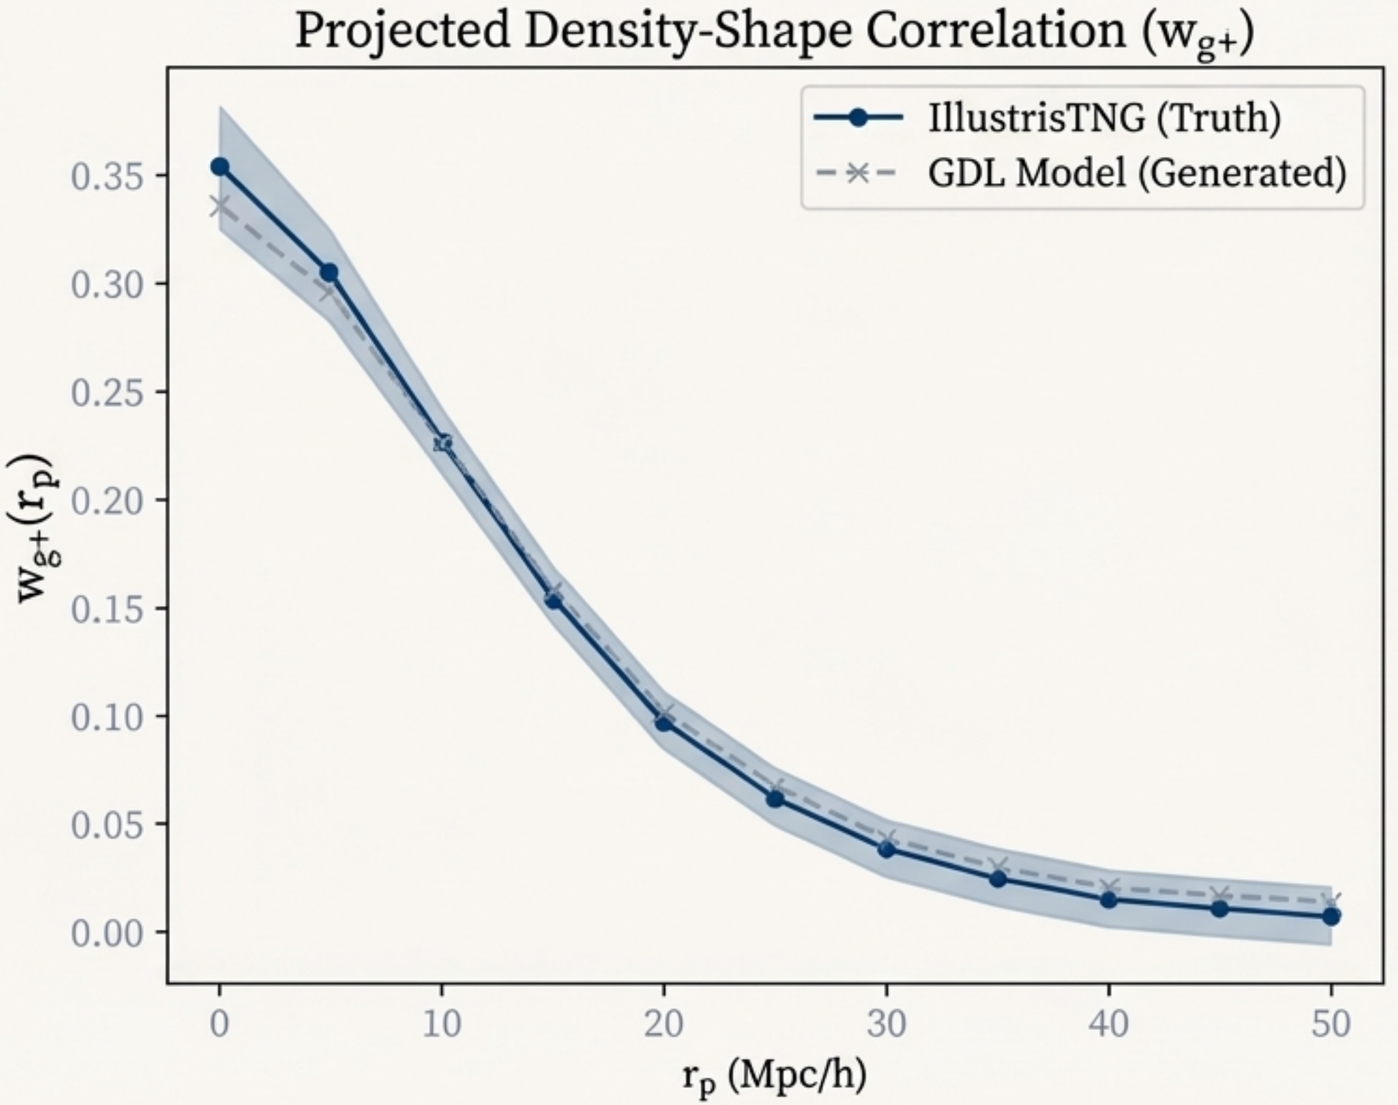
Graphs: Model relationships and interactions in the cosmic web.
Symmetry: Explicitly respect the physical laws of the universe, like rotation and translation (E(3) equivariance).
Manifolds: Learn distributions on curved spaces, like the 3D orientations of galaxies [SO(3)].

The Result: The Model Learns the Joint Distribution of Galaxy Properties



Key Takeaway: Our model generates synthetic galaxies with sizes, shapes, and colors that are statistically consistent with the complex hydrodynamical simulation (IllustrisTNG).
Quote: “Wasserstein-1 distances indicate per cent-level (or better) agreement in the 1D distributions.”

Proof: The Model Captures the Physics of Cosmic Alignment



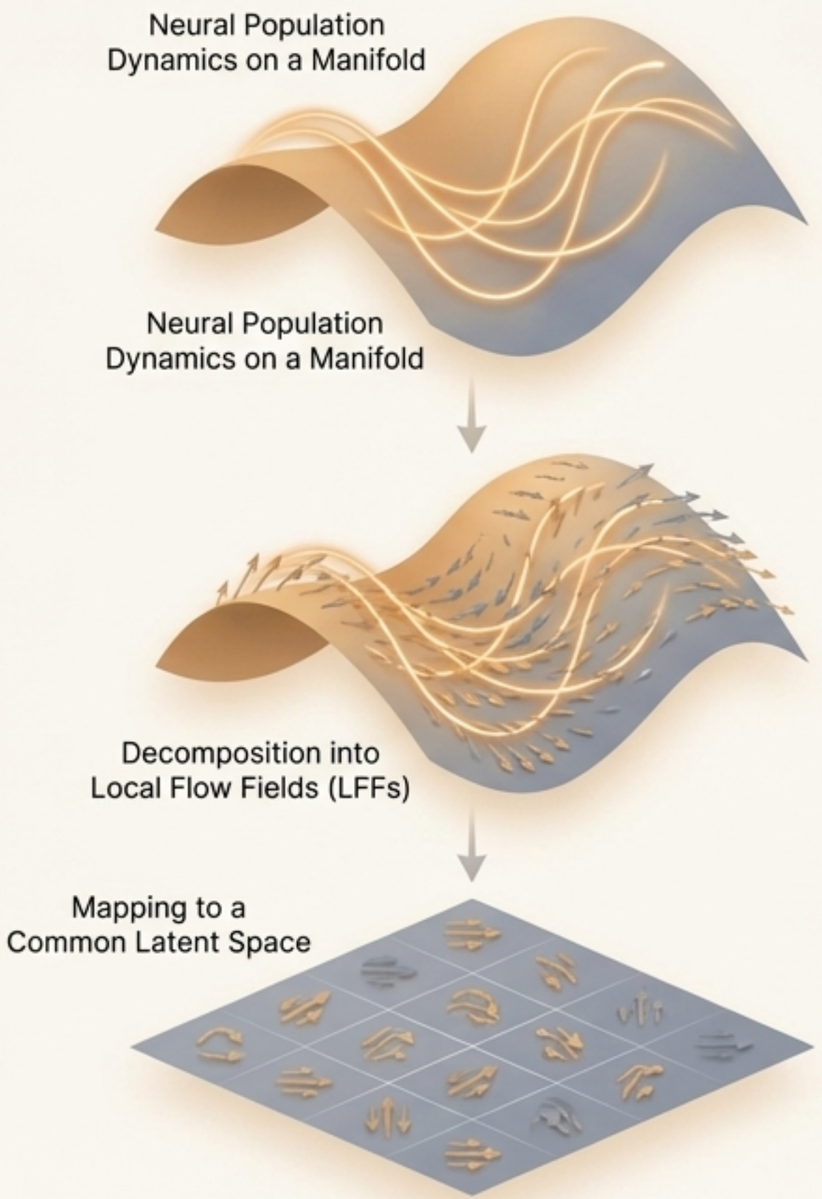
The model correctly reproduces the correlation between galaxy positions and their intrinsic shapes (`$w_{g+}$`). This is a highly non-trivial statistic that depends on complex baryonic and gravitational physics at non-linear scales.

Conclusion: The GDL model successfully learned the galaxy-halo connection for intrinsic alignments.





These Geometric Principles Are Not Limited to the Cosmos



The Same Approach is Decoding the Brain

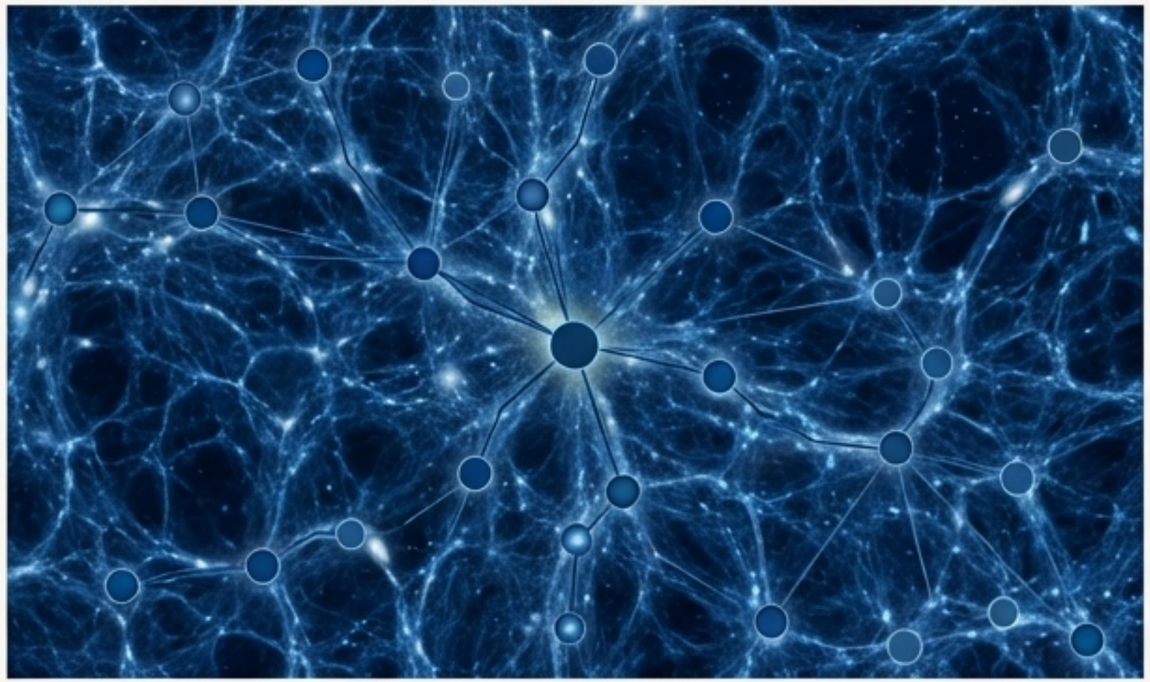
Introducing MARBLE: A method that uses geometric deep learning to find interpretable representations of neural population dynamics.

The Core Idea: Like our cosmology model, MARBLE assumes neural dynamics evolve on low-dimensional manifolds. It learns the "flow fields" on these manifolds.

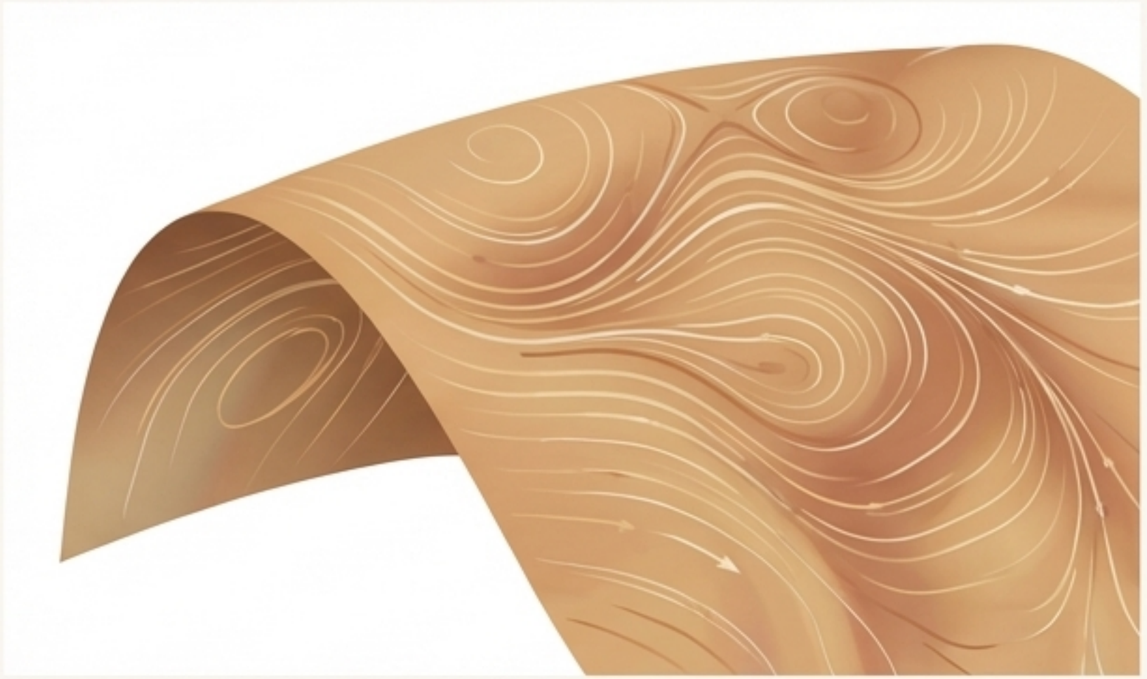
MARBLE "decomposes on-manifold dynamics into local flow fields and maps them into a common latent space using unsupervised geometric deep learning."

Two Domains, One Unifying Principle

Cosmology: Spatial Geometry



Neuroscience: Dynamical Geometry



- **Cosmology:** We modeled the static structure of the cosmic web and the orientation of galaxies on the SO(3) manifold. The geometry was spatial.
- **Neuroscience:** We modeled the temporal dynamics of neural populations as flows on a latent manifold. The geometry was dynamical.
- **The Unifying Idea:** In both cases, success came from building models that explicitly understood the system's inherent geometry and symmetries.

A Unified Framework for Geometric Science

Domain	Fundamental Geometry	GDL Application
 **Astrophysics**	3D Euclidean Space (E(3)) & Rotations (SO(3))	Simulating Galaxy Intrinsic Alignments for weak lensing surveys.
 **Neuroscience**	Low-Dimensional Latent Manifolds	Discovering consistent, cross-animal neural representations for behavior.
 **Climate Science**	The Sphere (S^2)	Overcoming projection distortions for more accurate global weather forecasting.

The same core principles of Geometric Deep Learning—respecting relationships, intrinsic structure, and symmetries—provide a common language to tackle fundamental challenges across disparate scientific fields.

Further Reading & Key Sources

Spherical CNNs
Cohen, T. S., Geiger, M., Köhler, J., & Welling, M. (2018). *Spherical CNNs*. ICLR.

Equivariant GNNs in Cosmology
Jagvaral, Y., Lanusse, F., & Mandelbaum, R. (2025). *Geometric deep learning for* galaxy-halo connection: a case study for galaxy intrinsic alignments. Monthly Notices of the Royal Astronomical Society.

Geometric DL in Neuroscience
Gosztolai, A., Peach, R. L., et al. (2025). *MARBLE:* interpretable representations of neural population dynamics using geometric deep learning. Nature Methods.

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