OPEN ACCESS

Artificial Intelligence COMMENTARY

The Innovation Informatics

Beyond Euclidean deep learning: Geometric intelligence for Earth observation

Xin Li,[1,2] Xin Lyu,[1,2,*] Zhonghao Chen,[1] Caifeng Wu,[1] Yiwei Fang,[1] Zhennan Xu,[1] Yubao Qiu,[3,4,5,*] and Nan Xu[6,7,*]

[1]College of Computer Science and Software Engineering, Hohai University, Nanjing 211100, China

[2]Key Laboratory of Water Big Data Technology of Ministry of Water Resources, Hohai University, Nanjing 211100, China

[3]International Research Center of Big Data for Sustainable Development Goals, Beijing 100094, China

[4]Key Laboratory of Digital Earth Science, Aerospace Information Research Institute, Chinese Academy of Sciences, Beijing 100094, China

[5]University of Chinese Academy of Sciences, Beijing 100049, China

[6]MNR Key Laboratory for Geo-Environmental Monitoring of Great Bay Area & Guangdong Key Laboratory of Urban Informatics & Shenzhen Key Laboratory of Spatial Smart Sensing and Services, Shenzhen University, Shenzhen 518060, China

[7]School of Architecture and Urban Planning, Shenzhen University, Shenzhen 518060, China

*Correspondence: lvxin@hhu.edu.cn (X.L.); qiuyb@aircas.ac.cn (Y.Q.); xunan2025@szu.edu.cn (N.X.)

Received: May 12, 2026; Accepted: October 30, 2026; Published Online: October 30, 2026; https://doi.org/10.59717/j.xinn-inform.2026.100076

© 2026 The Author(s). This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/).

Citation: Li X., Lyu X., Chen Z., et al. (2026). Beyond Euclidean deep learning: Geometric intelligence for Earth observation. *The Innovation Informatics* **2**:100076.

THE RISE OF AI-DRIVEN EARTH OBSERVATION

Earth observation (EO) is entering an era in which abundant satellite data, artificial intelligence (AI), and growing demand for Earth Intelligence (EI) are converging.[1] Constellations of Earth-orbiting sensors now provide continuous, multi-resolution measurements of climate dynamics, ecosystem change, land use, hazards, and human activities. At the same time, deep learning has become central to EO analysis, from land-cover mapping and object detection to medium-range weather prediction and global environmental monitoring.

Recent models already demonstrate the power of large-scale Earth AI. GraphCast, for example, uses graph-based representations to predict global weather with high skill,[2] while AlphaEarth Foundations produces general geospatial embeddings for mapping and monitoring from sparse labels.[3] These advances show that EO is moving beyond image interpretation toward decision-ready EI. However, multimodal unification does not necessarily imply geometric preservation: a shared semantic space may still lose task-relevant relations such as geodesic distance, graph connectivity, native sampling support, or physically constrained latent structure. We therefore frame Geometric Earth Intelligence (GEI) as a complementary representation principle that preserves and propagates such geometry while learning shared Earth representations. Here, Earth AI denotes AI models for EO and Earth-system data, EI denotes actionable understanding derived from them, and GEI denotes the geometry-aware representation principle connecting the two.

GEI is not intended to compete with multimodal Earth foundation models or to rename existing geometric methods. Its contribution is to ask how these geometric tools can be organized into a shared computational layer that preserves native Earth structures while remaining compatible with modern foundation models and operational EO workflows.

EARTH OBSERVATION IS GEOMETRIC, BUT NOT ALWAYS IN THE SAME WAY

EO data are not simply images of the planet. They are observations embedded in multiple geometric structures. Importantly, geometry is task- and scale-dependent. For many local mapping problems, such as crop classification or building detection within a small region, planar image representations may be sufficient and efficient. Geometry becomes more central when models address regional systems, global processes, long-range interactions, irregular observations, or multimodal fusion.

Four types of structure are especially relevant. First, global atmospheric, oceanic, and climate processes unfold on a spherical surface, where geodesic neighborhoods and spherical meshes can better reflect spatial relations than flat projections. Second, many Earth systems are networked: river basins, ecological corridors, transport systems, urban interactions, and climate teleconnections can be represented as graphs with physical or functional edges. Third, EO observations are often irregular, including satellite swaths, LiDAR point clouds, in situ sensors, and missing data induced by clouds or orbital constraints. Fourth, high-dimensional signals, such as hyperspectral signatures, vegetation states, and environmental variables, often occupy latent manifolds shaped by physical and biochemical processes.

These structures are not abstract mathematical decorations. They determine how information should move through a model, which neighbors should interact, how observations from different sensors should be aligned, and when a learned representation should be constrained by physical or spatial relationships. Treating all EO data as flat images may be computationally convenient, but it can obscure these organizing principles when they are central to the task.

EXISTING EARTH AI AND REMAINING GEOMETRIC GAPS

The success of CNNs, transformers, and Earth foundation models should not be underestimated. Grid-based representations are scalable, easy to pretrain, and effective for many EO benchmarks. The limitation is not Euclidean modeling itself, but the absence of a common representation layer that can preserve multiple geometries when they matter. More specifically, flat projection can distort geodesic neighborhoods on the sphere; generic spatial or raster representations can obscure functional or directional graph topology; resampling can erase the native support and density of irregular observations; and unconstrained embeddings may capture correlations without guaranteeing physically meaningful manifold organization.

This is where current progress also reveals open questions. GraphCast shows that graph-based weather prediction can be powerful, but its geometry is designed for a specific global forecasting setting.[3] AlphaEarth Foundations demonstrates the value of general geospatial embeddings, but a shared embedding space does not by itself guarantee explicit preservation of spherical, graph, irregular-sampling, and physical structures across diverse EI tasks.[4] The broader Earth foundation model literature similarly highlights the need for models that are general, interpretable, multimodal, and physically meaningful.[4] GEI aims to address this complementary gap by treating geometry as a reusable representational layer rather than a collection of isolated methods.

Such a layer would also help clarify evaluation. The question is not whether a geometric model always outperforms a planar one, but where geometric priors improve data efficiency, extrapolation, interpretability, or cross-region transfer. This distinction avoids over-generalization and makes GEI testable rather than purely visionary.

GEI AS A MULTI-GEOMETRY REPRESENTATION FRAMEWORK

Geometric deep learning provides the mathematical basis for this shift by extending deep learning to graphs, manifolds, groups, and other non-Euclidean domains.[5] In the GEI framework, however, the central question is not whether to replace grids with graphs or spheres. The goal is to build a hybrid representation that combines the scalability of grid tokens with explicit geometry-aware modules.

As illustrated in Figure 1, GEI can be instantiated through a five-step implementation scaffold; the sequence itself is not intended as the source of novelty. First, EO inputs are associated with their relevant geometry, such as grids, spheres, graphs, point sets, or latent manifolds. Second, structure-

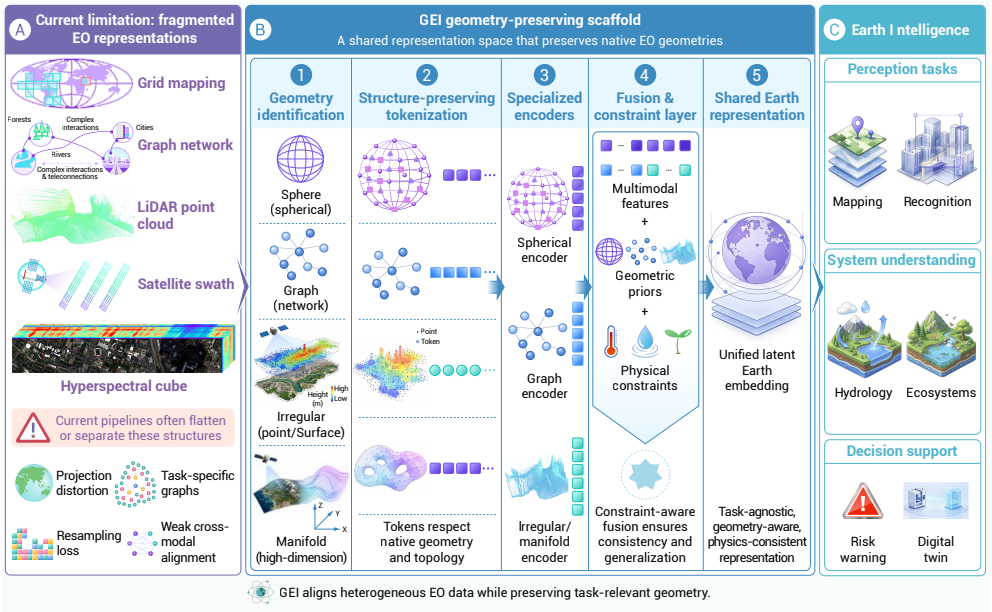


Figure 1. GEI as a geometry-preserving representation scaffold for Earth Intelligence (A) Current EO pipelines often flatten or separate heterogeneous representations, including grid images, river networks, LiDAR point clouds, satellite swaths, and hyperspectral cubes, leading to projection distortion, task-specific graphs, resampling loss, and weak cross-modal alignment. (B) GEI uses a five-step implementation scaffold in which native geometry is identified, preserved during tokenization, used to condition encoding, and retained during multimodal fusion and physical constraint integration before producing a shared Earth representation. (C) The resulting representation supports downstream Earth Intelligence tasks, including perception, system understanding, and decision support.

preserving tokenization converts these inputs into tokens that retain native coordinates, topology, and sampling information. Third, specialized encoders learn representations using spherical, graph, irregular, or manifold-aware operations. Fourth, a fusion and constraint layer aligns multimodal features while retaining task-relevant geometric relations and incorporating geometric priors and physical constraints. Finally, the model produces a shared Earth representation that can be connected to perception tasks, system understanding, and decision support.

The GEI-specific requirement is that task-relevant geometry remains explicit throughout this scaffold rather than being absorbed only implicitly into a shared embedding. Geodesic neighborhoods help preserve global spatial relations; message passing captures system connectivity; coordinate-aware encoders reduce information loss from irregular sampling; and physical priors constrain learned representations toward plausible Earth system behavior. GEI therefore augments, rather than replaces, existing Earth foundation models. It can be implemented through geometry tokens, graph adapters, spherical encoders, point-cloud modules, or manifold regularization inserted into scalable EO pipelines.

Crucially, the proposed framework does not require a single universal graph or segmentation scheme for all tasks. A hydrological application may use river topology, an urban application may use road or administrative relations, and a global climate task may use spherical meshes. GEI instead provides a mechanism for identifying, encoding, and fusing the geometry that is relevant to each task while maintaining a shared representation space for downstream EI.

OUTLOOK: BUILDING GEOMETRY-AWARE EARTH INTELLIGENCE

The practical value of GEI lies in making Earth AI more structure-aware, interpretable, and transferable. Its benefits should be evaluated through concrete benchmarks rather than asserted as universal advantages: global forecasting for spherical effects, watershed prediction for graph structure, LiDAR and swath analysis for irregular sampling, and multimodal mapping for cross-geometry fusion. Such tests should ask not only whether geometry improves accuracy, but under which spatial scales, sampling regimes, and relational dependencies it improves data efficiency, extrapolation, physical consistency, or cross-region transfer.

A realistic path forward is incremental. Near-term systems can add geometry-aware adapters to existing foundation models. Beyond implementation, three questions are especially actionable: Can models select task-relevant geometry rather than impose a universal structure? Can a shared representation preserve multiple geometries without sacrificing multimodal scalability? Which geometric priors deliver measurable gains in cross-region transfer, extrapolation, or robustness to irregular sampling? In this sense, GEI is not a slogan for replacing current EO AI. It is a research agenda for transforming heterogeneous Earth observations into actionable Earth Intelligence through representations that respect the structure of the planet itself.

REFERENCES

1. Reichstein M., Camps-Valls G., Stevens B., et al. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature* **566**:195–204. DOI:10.1038/s41586-019-0912-1
2. Lam R., Sanchez-Gonzalez A., Willson M., et al. (2023). Learning skillful medium-range global weather forecasting. *Science* **382**:1416–1421. DOI:10.1126/science.adi2336
3. Brown C.F., Kazmierski M.R., Pasquarella V.J., et al. (2025). AlphaEarth Foundations: An embedding field model for accurate and efficient global mapping from sparse label data. arXiv preprint. DOI:10.48550/arXiv.2507.22291
4. Zhu X.X., Xiong Z., Wang Y., et al. (2026). On the foundations of Earth foundation models. *Commun. Earth Environ.* **7**:103. DOI:10.1038/s43247-025-03127-x
5. Bronstein M.M., Bruna J., Cohen T., et al. (2021). Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. arXiv preprint. DOI:10.48550/arXiv.2104.13478

FUNDING AND ACKNOWLEDGMENTS

This work was funded in part by the National Key Research and Development Program of China under Grant No. 2024YFC3210800, National Natural Science Foundation of China (Grant No. 62401196), Natural Science Foundation of Jiangsu Province (Grant No. BK20241508), Shenzhen Key Laboratory Program (Grant No. SYSPG20241211173845013), and the Disciplines Breakthrough Project in Aerospace Information and Spatiotemporal Intelligence, MOE, China. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

AUTHOR CONTRIBUTIONS

All authors contributed to the manuscript and approved the final version.

DECLARATION OF INTERESTS

The authors declare no competing interests.