GridMind: A DGGS-Grounded Multi-Agent System for Spatial Reasoning

Mingke Erin Li
University of Calgary
Calgary, Canada
mingke.li@ucalgary.ca

Jinya Wang
University of Calgary
Calgary, Canada
jinya.wang@ucalgary.ca

Steve H.L. Liang
University of Calgary
Calgary, Canada
liangs@ucalgary.ca

ABSTRACT

Spatial reasoning depends on relations such as adjacency, containment, and hierarchy, while Large Language Models (LLMs) infer these relations unreliably from coordinates, maps, or text. Existing agentic GeoAI systems often address this limitation by reconstructing spatial relations for each query, without a persistent spatial substrate for reasoning. We present GridMind, a Discrete Global Grid Systems (DGGS) grounded multi-agent system that makes spatial relations a stable part of the architecture. To our knowledge, GridMind is the first system to use DGGS topology as a standing multi-agent substrate for spatial reasoning. It indexes space with DGGS and models both grid cells and spatial features as agents. Adjacency, hierarchy, and containment are derived from cell identifiers rather than inferred by the LLM. In the demonstration, users interact with an oil-and-gas asset-management sandbox to place sensors, replay leak events, observe selective sensor activation, and ask agents why and how they acted. GridMind shows that grid-grounded spatial relations can make multi-agent spatial reasoning explainable, scalable, and transferable to other DGGS-indexed domains.

CCS CONCEPTS

• Information systems → Geographic information systems; • Computing methodologies → Multi-agent systems

KEYWORDS

Discrete Global Grid Systems, multi-agent systems, spatial reasoning, large language models, asset monitoring

ACM Reference Format:

Mingke Erin Li, Jinya Wang, and Steve H.L. Liang. 2026. GridMind: A DGGS-Grounded Multi-Agent System for Spatial Reasoning. In *The 34th ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '26), November 3–6, 2026, Riverside, CA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3841645.3843111

SIGSPATIAL '26 November 03--06, 2026, Riverside, CA, USA

ACM ISBN 979-8-4007-2950-8/2026/11
https://doi.org/10.1145/3841645.3843111

1 Introduction

Spatial reasoning derives new spatial facts or decisions from known relationships such as adjacency, containment, proximity, and hierarchy. In geographic contexts, such reasoning should respect Tobler's First Law of Geography, which states that nearby things tend to be more related than distant things [1]. However, Large Language Models (LLMs) remain unreliable when spatial relationships must be inferred from raw coordinates, maps, or unstructured text. Recent benchmarks show that strong models answer fewer than one quarter of basic spatial cognition questions correctly [2], and LLMs generated coordinates fail to preserve the distance-decay structure implied by Tobler's law [3]. These limitations suggest that spatial relations should be explicitly grounded rather than implicitly inferred by language models.

We present GridMind, a multi-agent system that grounds spatial reasoning in Discrete Global Grid Systems (DGGS). A DGGS hierarchically partitions the Earth's surface into uniform cells whose topology encodes containment, hierarchy, and adjacency [4]. In GridMind, cells across resolutions act as agents, while spatial features are represented by agents bound to the cells that contain them. Spatial reasoning is driven by grid topology: 1) adjacency supports propagation between neighboring cells, 2) hierarchy connects coarse summaries with fine-scale detail, and 3) containment binds each feature to its corresponding cell. DGGS operators resolve these spatial relations explicitly, while the LLM invokes the operators, reasons over their outputs, and provides the natural-language interface by interpreting questions, routing them to relevant agents, and rendering grounded memory as explanations. We demonstrate GridMind in an oil-and-gas asset-management scenario involving production facilities, methane emissions, and sensors, where grid-grounded relationships support selective sensor activation, resource conservation, and sensor-placement decisions.

This paper makes two contributions. First, it grounds spatial relations in DGGS topology, so that adjacency and hierarchy are obtained from grid structure rather than inferred by a language model. Second, it models grid cells and spatial features uniformly as agents that coordinate through decentralized, event-driven communication. We realize these ideas in an interactive sandbox where users can run spatial scenarios, observe agent behavior, and question cell, facility, and sensor agents in natural language.

2 Related Work

2.1 Spatial Reasoning Limitations of LLMs

Recent work consistently shows that large language models remain weak at spatial reasoning. Yang et al. [2] evaluated spatial cognition across landmark, route, and survey knowledge and found that GPT-4-turbo answered fewer than one quarter of questions correctly. On the MapEval benchmark, no foundation model exceeded 67% accuracy, and all models lagged human performance by more than 20% [5]. GeoAnalystBench further shows that models are better at planning spatial-analysis workflows than at executing them correctly [6]. Mai et al. [3] trace part of this limitation to spatial representation, showing that LLM-generated coordinates do not preserve the distance-decay structure described by Tobler's law. These findings expose a gap between language-model reasoning and the structured spatial relations required for geographic analysis.

2.2 Grounding Spatial Reasoning

One response is to augment language models with external spatial structure or computation. Yang et al. [2] coupled an LLM with deterministic GIS algorithms through a mental-map builder, improving spatial-cognition accuracy from below 25% to roughly 70%. Retrieval and workflow-based systems follow a similar direction. GeoAgentic-RAG retrieves spatial data and decomposes queries across specialized agents [7], while Spatial-Agent parses questions into workflow graphs grounded in spatial core concepts such as neighborhood, hierarchy, and containment [8]. These systems improve reliability by adding spatial tools, retrieval, or workflow structure, but the relevant spatial relations are still assembled for each query rather than provided as a persistent substrate.

3 The GridMind System Design

3.1 DGGS Topology as Spatial Grounding

GridMind represents space using the rearranged Hierarchical Equal Area isoLatitude Pixelization (rHEALPix) DGGS (Figure 1), which partitions the Earth's surface into a hierarchy of cells through aperture-9 refinement [4, 9]. Cells at the same resolution have equal area, so emission totals are directly comparable across cells and a hotspot flag carries the same meaning anywhere on the grid. Each cell has a string identifier consisting of a face symbol followed by a sequence of digits. Each appended digit selects one of nine child cells, so refinement to the next resolution appends one digit to the identifier (Figure 1).

Containment maps a feature at a longitude λ_f and latitude φ_f to the unique cell that contains it at resolution r. Hierarchy follows from the identifier prefix, where the parent cell is obtained by removing the final digit, and its child cells are obtained by appending digits $0, \dots, 8$. Adjacency is defined by the ring of cells surrounding a cell at the same resolution. These grounding operators are summarized in Equations 1-4:

$$c_r(f) = I_r(\varphi_f,\ \lambda_f) \tag{1}$$
$$parent(s) = s_{1:|s|-1} \tag{2}$$
$$child(s) = \{s \cdot k \mid k = 0, \dots, 8\} \tag{3}$$
$$N(c) = ring(c, 1) \tag{4}$$

where f is a spatial feature, $c_r(f)$ is the cell that contains feature f at resolution r, I_r maps a location to its containing rHEALPix cell at resolution r, s is a cell identifier, $parent(s)$ and $child(s)$ denote the parent cell and the set of child cells of s, $s \cdot k$ denotes the identifier extended by a digit k, and $N(c)$ is the set of cells adjacent to c. Because these operators are computed from DGGS topology, the spatial relations that drive agent interaction are exact and stable, independent of language-model inference.

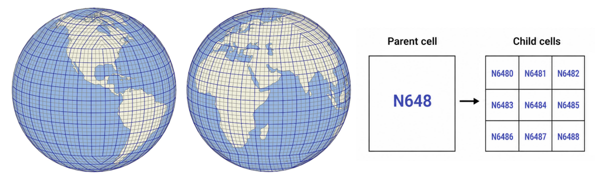

Figure 1: The rHEALPix DGGS tessellating the globe into multi-resolution cells with hierarchical indexing

3.2 Agents and Communication Channels

GridMind models both grid cells and spatial features as agents. Cell agents exist at each DGGS resolution, so the same region is represented by coarse cells that summarize broad conditions and finer cells that capture local detail. Feature agents represent domain entities such as facilities or sensors and are bound to their containing cells. Each agent maintains local state and responds only to events that reach it, such as a leak onset, sensor reading, neighbor message, or user query. Computation is decentralized and event-driven, where agents remain idle until triggered, and inactive regions consume no processing resources.

Figure 2 shows the system architecture design. Agents communicate through three topology-defined channels, including inward, lateral, and vertical. The inward channel connects a feature agent with its containing cell, allowing facilities and sensors to report their status. The lateral channel connects adjacent cells at the same resolution, supporting communication among neighboring regions. The vertical channel connects parent and child cells, carrying information across resolutions.

These channels carry the system's spatial reasoning, where agents use topology-defined relations to decide how to act. Along the lateral channel, a cell with a leaking facility can notify adjacent cells that emissions may propagate, allowing nearby sensors to activate when local conditions require attention. Along the vertical channel, child cells aggregate emission totals to parent cells, giving coarse cells regional summaries. The same hierarchy can be followed downward to trace a high parent total to the contributing child cells. In the methane emission scenario, this reasoning enables selective sensor activation, conserving power while supporting asset-management decisions.

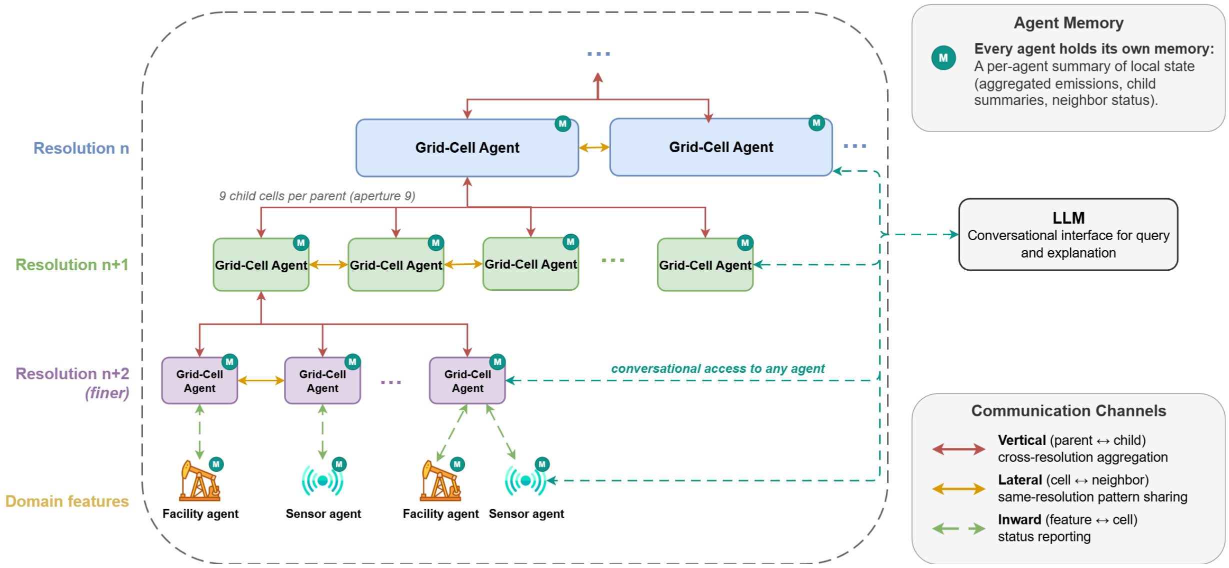

Figure 2: GridMind architecture with cell and feature agents communicating over the inward, lateral, and vertical channels

3.3 Agent Memory Management

Each agent reasons over a memory that summarizes the scenario from its own perspective. Rather than storing a running message history, the memory is computed as a fixed summary of the compiled scenario, including its features, events, and cell states. The summary function selects scenario records within the agent's scope and applies fixed aggregations to them. A DGGS operator defines this scope, which for a cell agent corresponds to the prefix subtree in Equation 3. We define agent memory in Equation 5:

$$M(a) = \mu(S, a) \tag{5}$$

where $M(a)$ is the memory of the agent a, S is the compiled scenario, and μ is the fixed summary function. Each agent's view is reproducible and grounded in the same underlying records, while the language model operates only on the resulting summaries and does not alter the spatial relations they encode.

4 Implementation

GridMind is implemented as a web application with a Python backend and a browser-based front end. The spatial substrate is built with the DGGAL library [10], which assigns each feature to its containing rHEALPix cell and provides the parent-cell, child-cell, and neighbor lookup functions that define the communication channels. The demonstration uses the 2023 British Columbia Energy Regulator (BCER) Leak Detection and Repair (LDAR) dataset (https://www.bc-er.ca/files/documents/LDAR-data/2023-BC-LDAR-Data.xlsx). Facilities, leak events, and repair records are indexed into the grid and compiled into the scenario from which agent memories are computed.

The backend uses FastAPI to run the event-driven simulation and stream reasoning updates to the client through server-sent events. The front end is built with vanilla JavaScript and Mapbox GL JS, rendering the grid, spatial features, and live reasoning trace while supporting natural-language interaction with individual agents. The language model is accessed through a single interface, with GPT-4o-mini as the default.

5 Demonstration

The demonstration places the audience in the role of an operator monitoring an oil-and-gas region drawn from the BCER LDAR dataset, where production facilities are distributed across British Columbia and indexed into rHEALPix cells at multiple resolutions. The interface contains scenario controls on the left, an interactive map at the center, a scenario timeline beneath the map, a live reasoning panel along the bottom, and an agent chat panel on the right (Figure 3). To begin, the user places sensors on the map, each of which binds to its containing cell, sets a time window, and presses Play. Leak onsets and repairs within the window are then replayed as hourly discrete events, allowing the scenario to unfold step by step. The scenario timeline places a mark at each hour when an event occurs, and the user can move backward and forward along it. Selecting a mark replays the map and the reasoning panel at that point in the run.

As the simulation runs, each event propagates through the three topology-defined channels. A leaking facility reports inward to its containing cell, the cell updates its state and flags a hotspot, and the flagged cell notifies adjacent cells laterally that emissions may spread. A neighboring cell activates a nearby sensor only when the notification or the sensor's own readings warrant it, which appears on the map as a change in sensor state. In parallel, cell states roll up vertically, so that coarse cells summarize the conditions of their child cells. The reasoning panel records each conclusion as it is drawn, identifying the spatial relation used and the reason for the resulting action.

The map supports spatial progressive disclosure across resolutions. At coarse zoom levels, users see regional cells that

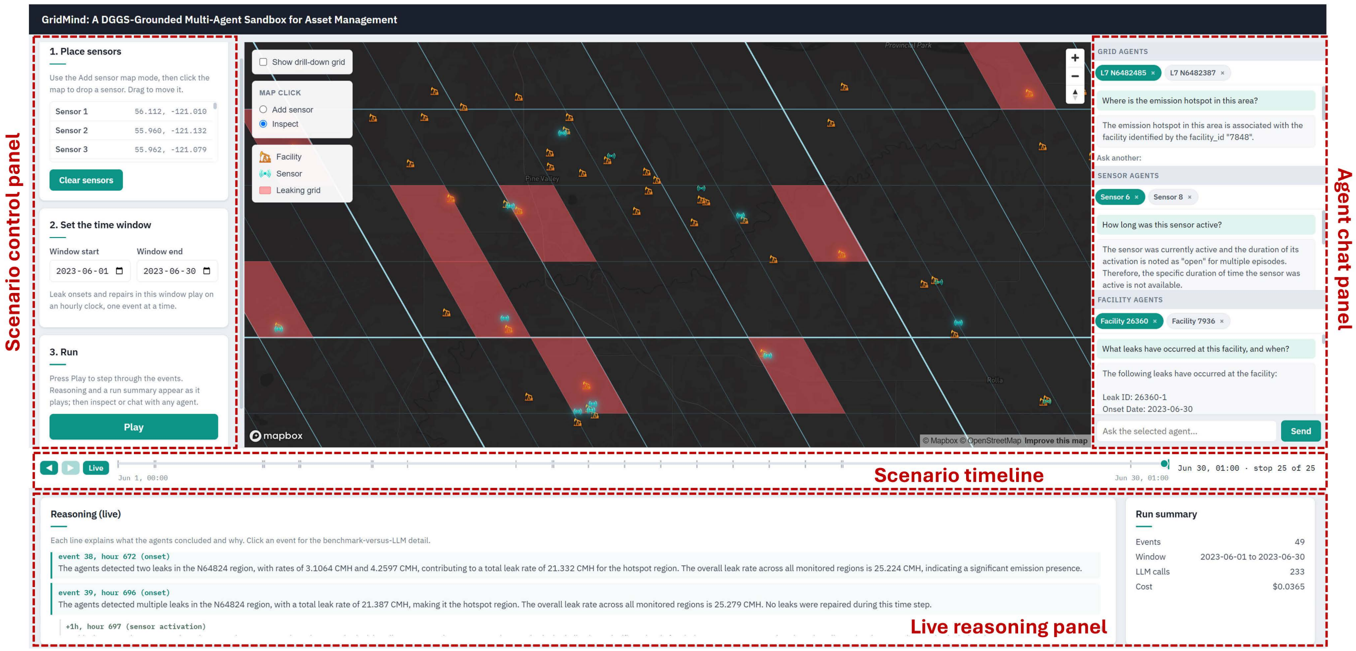

Figure 3: User interface of the GridMind

summarize broad conditions. Drilling into a flagged cell reveals finer child cells and the facilities or sensors responsible for the alert. Because computation is event-driven and activation is selective, only a subset of cells and sensors is active at a given time. The display, therefore, shows monitoring effort concentrating where leaks occur rather than being distributed uniformly across all assets.

After a scenario has played, users can question individual agents through the chat panel. A cell agent can explain why it raised an alert, what it aggregated from its child cells, and at which resolution the hotspot emerged. A sensor agent can explain why it was activated, including whether it was triggered by a neighboring cell. A facility agent can report its leak onset and repair status. Each answer is generated from the agent's grounded memory. The demonstration shows GridMind's two main ideas in operation, including DGGS-grounded spatial relations and a uniform population of cell and feature agents coordinating through decentralized, event-driven communication.

6 Conclusion and Future Work

We presented GridMind, a multi-agent system that grounds spatial reasoning in DGGS topology. GridMind derives adjacency, hierarchy, and containment from rHEALPix cell identifiers and represents grid cells and spatial features as agents that coordinate through decentralized, event-driven communication. In an oil-and-gas methane scenario, agents aggregate conditions across resolutions, selectively activate sensors, and answer questions from grid-grounded memory. Agents are addressable views over the compiled scenario, so resident state scales with active leaks and sensors rather than the total number of grid cells. Prefix indexing, cached ancestor aggregates, and batched narrative calls can further reduce cost. Future work will increase agent autonomy and compare LLM reasoning over raw coordinates with DGGS-grounded reasoning to quantify the value of structural spatial grounding.

REFERENCES

[1] Waldo R. Tobler. 1970. A computer movie simulating urban growth in the Detroit region. Economic geography 46, sup1, 234–240.
[2] Anran Yang, Cheng Fu, Qingren Jia, Weihua Dong, Mengyu Ma, Hao Chen, Fei Yang, and Hui Wu. 2025. Evaluating and enhancing spatial cognition abilities of large language models. International Journal of Geographical Information Science 39, 9, 2009–2044. https://doi.org/10.1080/13658816.2025.2490701
[3] Gengchen Mai, Weiming Huang, Jin Sun, Suhang Song, Deepak Mishra, Ninghao Liu, Song Gao, Tianming Liu, Gao Cong, Yingjie Hu, Chris Cundy, Ziyuan Li, Rui Zhu, and Ni Lao. 2024. On the Opportunities and Challenges of Foundation Models for GeoAI (Vision Paper). ACM Transactions on Spatial Algorithms and Systems 10, 2, 1–46. https://doi.org/10.1145/3653070
[4] Kevin Sahr, Denis White, and A. Jon Kimerling. 2003. Geodesic discrete global grid systems. Cartography and Geographic Information Science 30, 2, 121–134. https://doi.org/10.1559/152304003100011090
[5] M.L. Dihan, M.T. Hassan, M.T. Parvez, M.H. Hasan, M.A. Alam, M.A. Cheema, M.E. Ali, and M.R. Parvez. 2024. Mapeval: A map-based evaluation of geo-spatial reasoning in foundation models. arXiv preprint arXiv, 2501.00316.
[6] Qianheng Zhang, Song Gao, Chen Wei, Yibo Zhao, Ying Nie, Ziru Chen, Shijie Chen, Yu Su, and Huan Sun. 2025. GeoAnalystBench: A GeoAI Benchmark for Assessing Large Language Models for Spatial Analysis Workflow and Code Generation. Transactions in GIS 29, 7. https://doi.org/10.1111/tgis.70135
[7] Chao Liang, Yuanzheng Cui, Run Shi, Guixiang Zha, Xin Yin, Mingzhong Xiao, Dong Xu, Xuejun Duan, and Bo Huang. 2026. GeoAgentic-RAG: A Multi-Agent framework for autonomous geospatial reasoning and visual insight generation with LLM. International Journal of Applied Earth Observation and Geoinformation 147. https://doi.org/10.1016/j.jag.2026.105195
[8] R. Bao, C. Yang, D. Yu, Z. Tang, G. Mai, and L. Zhao. 2026. Spatial-Agent: Agentic Geo-spatial Reasoning with Scientific Core Concepts. arXiv preprint arXiv, 2601.16965.
[9] Robert Gibb. 2016. The rHEALPix discrete global grid system. In *Proceedings of the IOP Conference Series: Earth and Environmental Science, 9th Symposium of the International Society for Digital Earth (ISDE)* October 5–9, 2016, Halifax, Canada. IOP Publishing Ltd, 012012. https://doi.org/10.1088/1755-1315/34/1/012012
[10] Jerome St-Louis. 2025. DGGAL: the Discrete Global Grid Abstraction Library. Retrieved September 30, 2025, from https://github.com/ecere/dggal