D-3577

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

Mingke Erin Li · Jinya Wang · Steve H. L. Liang

Department of Geomatics Engineering | University of Calgary | mingke.li@ucalgary.ca

01 Background and motivation

THE PROBLEM

Spatial reasoning derives new spatial facts and decisions from geographic relationships such as adjacency, containment, proximity, and hierarchy. These relationships should respect **Tobler's First Law of Geography**, which states that nearby things tend to be more related than distant ones. Although large language models (LLMs) are increasingly used in GeoAI applications, they remain unreliable when inferring spatial relationships directly from coordinates, maps, or text. This limitation highlights the need to explicitly ground spatial relationships instead of relying on LLM inference alone. Existing approaches address parts of this limitation:

- **Autonomous GeoAI systems** integrate LLMs with deterministic GIS tools to compute spatial relationships. However, many assemble these relationships for individual queries and operate primarily in response to user requests.
- **Grid-based agent models**, including cellular automata, typically apply predefined local rules at a single spatial resolution rather than enabling LLM-driven reasoning across scales.

Existing approaches have not provided a persistent spatial structure that grounds LLM agents to reason, communicate across resolutions, and respond autonomously to events as they occur.

OUR APPROACH

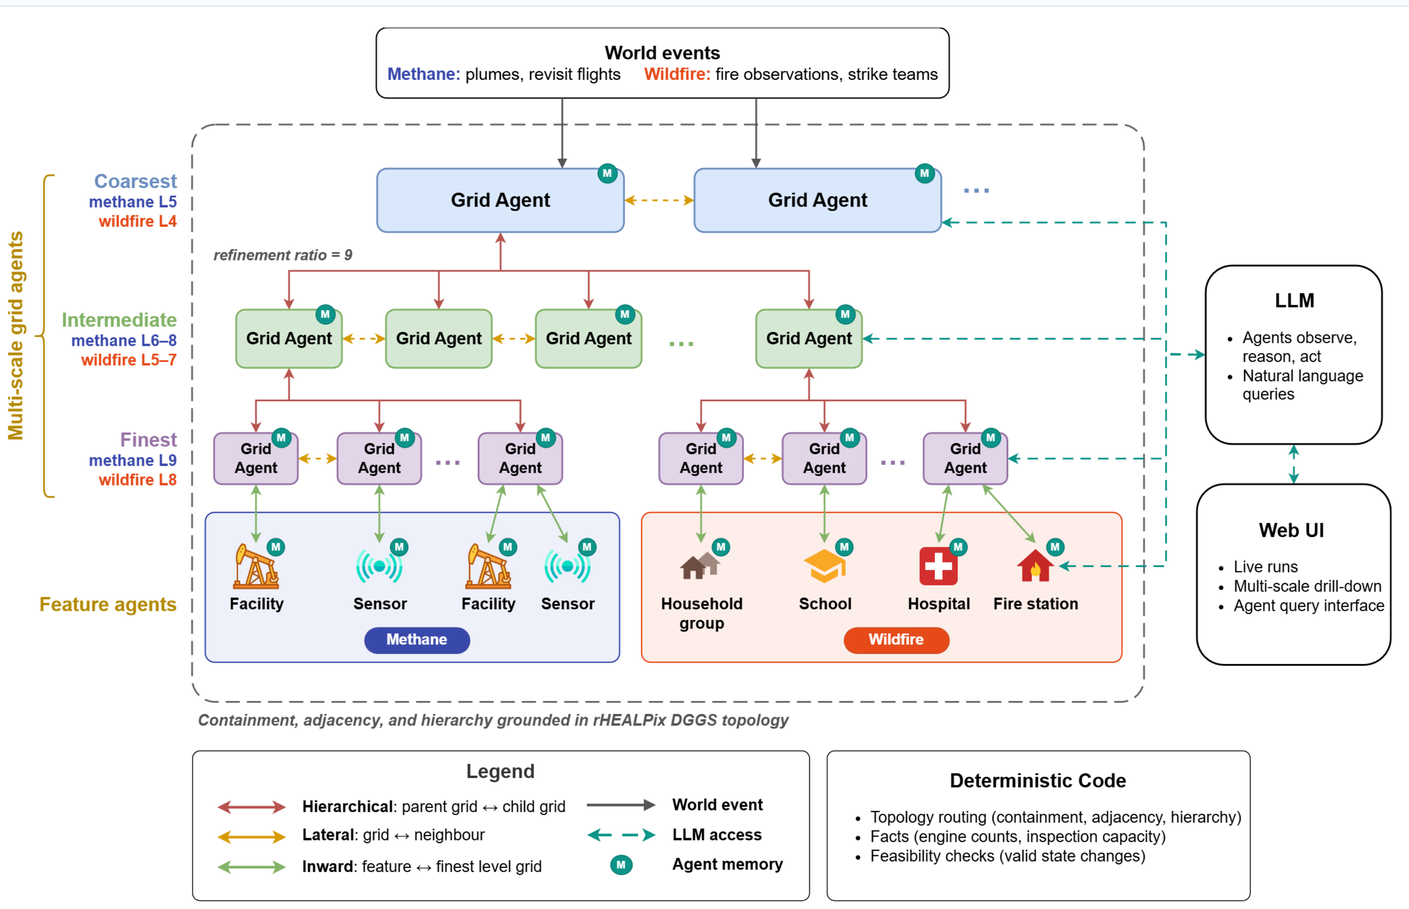
GridMind is a multi-agent system that uses a Discrete Global Grid System (DGGS) as a persistent substrate for spatial reasoning, communication, and memory. A DGGS partitions the Earth's surface into hierarchical grids whose identifiers support deterministic spatial operations. GridMind adopts **the rearranged Hierarchical Equal Area isoLatitude Pixelization (rHEALPix) DGGS**, which provides equal-area grids and a congruent hierarchy in which each child grid's identifier contains its parent's identifier as a prefix.

GridMind represents space and spatial entities through two types of agents. **Grid agents** operate at multiple DGGS resolutions, while **feature agents** represent geographic entities such as oil and gas facilities, sensors, buildings, and fire stations, each bound to its containing grid. Agents communicate through three channels defined by the DGGS topology.

- **Inward**, between feature agents and their containing grid agents
- **Lateral**, between neighbouring grid agents
- **Hierarchical**, between parent and child grid agents

Deterministic code computes spatial relationships from DGGS grid identifiers, while the LLM interprets incoming messages, reasons using the agent's persistent memory, conducts actions, and explains its decisions in natural language. Each agent maintains its state across events, enabling autonomous, event-driven responses. Users can also question any agent at any point during a run.

02 System architecture



Messages by communication channel

CHANNEL	METHANE	WILDFIRE
World	Plume observations Revisit results	Fire observations Strike-team arrivals
Hierarchical	Plume routing, case decisions District referrals Inspection assignments	Regional warnings, engine requests Request escalation Dispatch decisions
Lateral	Neighbour candidate requests Neighbour alerts, closure notices	Neighbour fire alerts
Inward	Candidate and monitoring requests Investigation requests Operator responses	Fire and watch notices Protective actions Dispatch responses

Persistent agent memory

Agents	METHANE	WILDFIRE
Coarsest grid agents	Referred cases Daily inspection capacity	Available engines Strike-team pool
Intermediate grid agents	Child-grid summaries Referrals in transit	Child-grid summaries Pending engine requests
Finest grid agents	Plume observations Open cases and revisit results	Fire state and open requests Engines on site
Feature agents	Facility: records and case history Sensor: monitoring assignment	Household, school, hospital: status Fire station: engine availability

03 Application scenarios

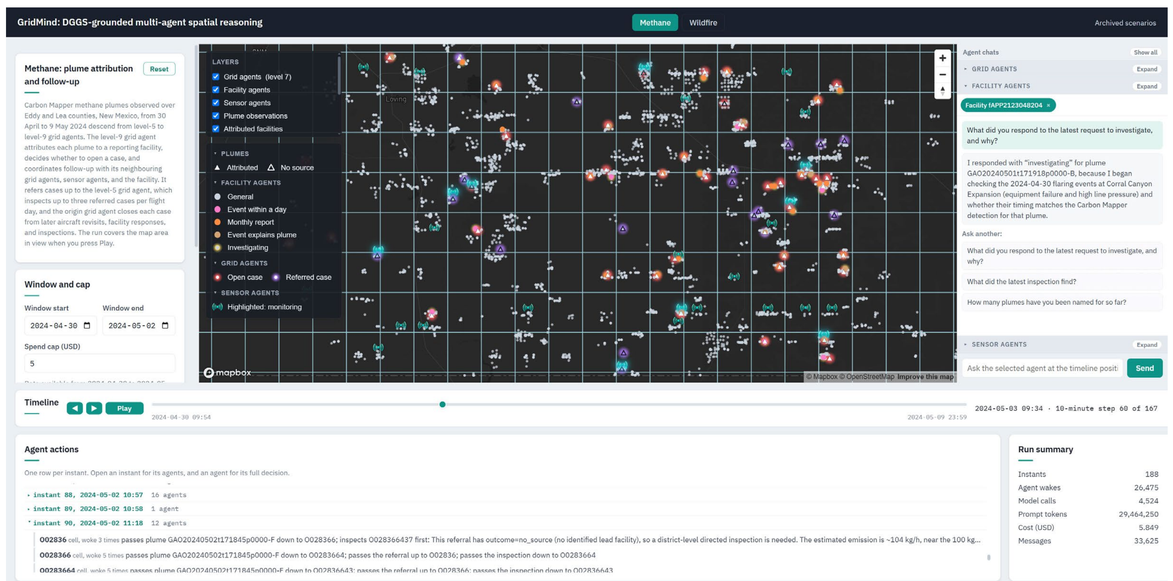
A | Methane plume attribution and follow-up

DATA & SETTING

This scenario replays Carbon Mapper airborne surveys over Eddy and Lea counties in the Permian Basin, New Mexico, covering 579 methane plume observations across eight flight days from 30 April to 9 May 2024. Each observation includes its time, emission rate, origin, and spatial extent, while subsequent surveys provide revisit observations. Public records from the New Mexico Oil Conservation Division (OCD) provide well and facility locations, reported venting and flaring events, and monthly vented and flared volumes. We place 30 sensor agents in level 9 grids with the most plume observations, and grid agents operate across rHEALPix resolutions 5 to 9.

AGENT WORKFLOW

Plume routing → Source attribution → Case opening and sensor tasking → Investigation and referral → Daily revisit and closure



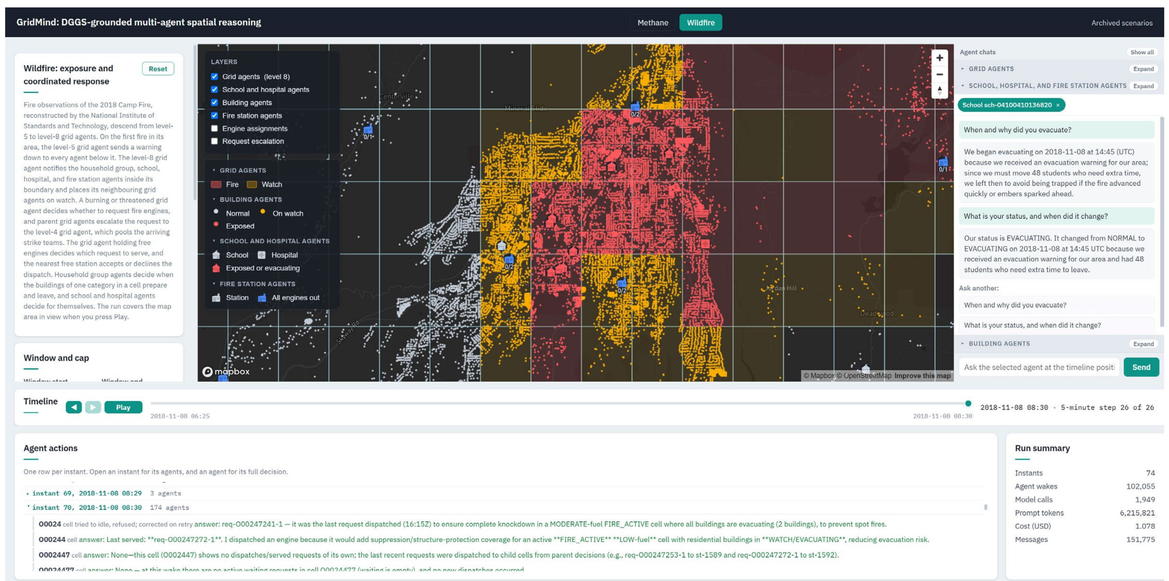
B | Wildfire exposure and coordinated response

DATA & SETTING

This scenario replays the early progression of the 2018 Camp Fire in Butte County, California, using 537 fire observations from the National Institute of Standards and Technology (NIST), recorded between 06:25 and 10:00 on 8 November 2018. The CAL FIRE damage survey provides 18,453 buildings, while facility records identify 16 fire stations with 22 engines. Additional agents represent schools and hospitals. Environmental conditions are provided by LANDFIRE fuel classes and gridMET weather data. Grid agents operate across rHEALPix resolutions 4 to 8, with additional strike teams available at level 4.

AGENT WORKFLOW

Regional warning and fire alerts → Protective action → Engine requests → Hierarchical dispatch → 30-minute review and release



CONCLUSION

The grid grounds spatial relations. The LLM agents decide how to act.

GridMind uses rHEALPix topology to ground spatial relationships and organize agent communication through inward, lateral, and hierarchical channels. Each agent maintains persistent memory across events, supporting LLM-driven decisions and natural language explanations. The same core architecture supports methane plume attribution and coordinated wildfire response across different spatial and temporal scales, with potential applications in other domains where spatial features interact through their containing grids. Future work will incorporate line features such as roads and pipelines, environmental context maintained at coarser grid resolutions to inform agent reasoning, and lateral alerts weighted by neighbourhood size.