From Topology to LLM Agent Skills

Vector-Quantized Geodesic Trajectories as Skills for Spatial Aware LLM-Driven Agents

Gabriel Turinici

CEREMADE CNRS, Université Paris Dauphine - PSL, Paris, FRANCE

Agentic AI Summit 2026, UC Berkeley, August 1-2, 2026

Overall

- **Goal:** Enhance spatial reasoning like navigation in partially observable and dynamic environments) in LLM-based agents through geometry-grounded skills.
- **Geodesic trajectories:** many geodesics are collected to capture the environment's topology and spatial structure.
- **Vector quantization:** Geodesics are compressed (clustered) into a small set of representative geodesics, each becomes a strategy.
- **LLM interpretation:** The LLM attaches to each strategy a natural-language description to obtain a TOOL.
- **LLM orchestration:** At decision time, the LLM selects the most appropriate tool based on the description, current state and objective.
- **Agentic:** high-level reasoning is handled by the LLM, while low-level execution relies on primitive actions (tools).
- **Key insight:** LLMs are most effective as orchestrators of geometry-derived skills rather than as direct spatial planners.

Rubik's cube analogy

To solve Rubik's cube one uses a combination of several moves fixed in advance called 'algorithms'.
Same set of 'algorithms' works for solving the cube no matter the initial position.

Two known algorithms RUR'URU2R' and FRUR'U'F' used to solve Rubik's cube, cf. `https://alg.cubing.net`

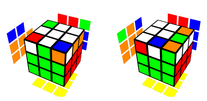

2D grid environment

- Agent: (blue cell).
- Cars: move upward 1 cell / time step.
- Free space: cells can be traversed safely.
- Grey cells: cells are not observed and may contain additional cars.
- Target: goal cell to be reached.
- Yellow: decision is taken at these times.
- Actions: wait, up, down, left, right
- Partial observability: only a local observation window centered on the agent ( ) is visible.

Objective: Reach the goal with minimum expected movements and time.
Multiple advance decisions : environment is fast, evolves between two decisions, at each time we take 5 future actions and act them no matter what.

Figure 1:From left to right: total grid @ $t = 0$; observation @ $t = 0$; observation grid just before reaching target; total evolution including decision times.

Geodesics = shortest-path connections

Geodesics: A geodesic captures the intrinsic geometry of a space by defining minimal paths between locations. Curvature influences the geodesic.

Vector Quantization (VQ)

VQ : Compression of a large set of trajectories into a (small) number of representative prototypes; it summarizes information in a high dimensional dataset.

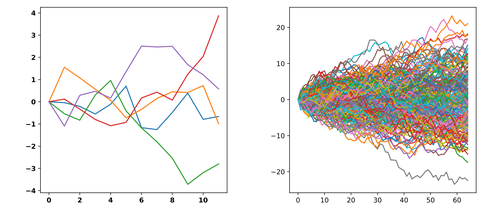

Figure 2:Quantization of the standard Brownian motion, also called cubature. Left: 5 representative trajectories are selected, right: 200 trajectories. [3, 2, 1].

Grid Geodesics

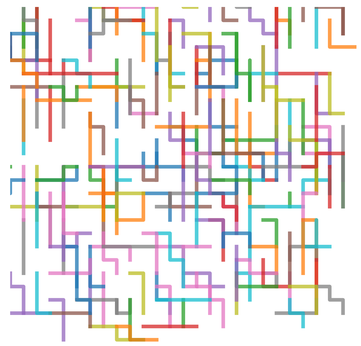

Geodesic Quantization

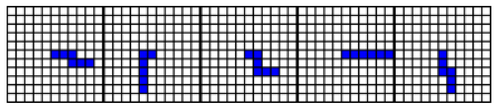

Tool 3 description by LLM

This tool executes a five-step **right–down–down–right–right** maneuver:

1. Move one cell right.
2. Move down one cell.
3. Move down one more cell.
4. Move right.
5. Move right again.

It has no wait actions and produces a net displacement of **three cells right and two cells down**, making strong progress toward a goal in the lower-right direction.

Its main defensive feature is the immediate lane change to the right: this can evade a known or hidden upward-moving car approaching in the starting column. The final two right moves also leave the column used for the downward segment, reducing continued exposure there. Compared with a straight downward tool, it spends less time in our traffic lane and gains information across several columns.

However, the two consecutive downward moves travel against upward-moving traffic in the same column. This tool should therefore be used only when that adjacent column is observed to be clear with sufficient safety margin; a hidden or nearby car below could create a head-on collision. The tool is best suited when the right side is relatively open and the target lies down and to the right.

Results

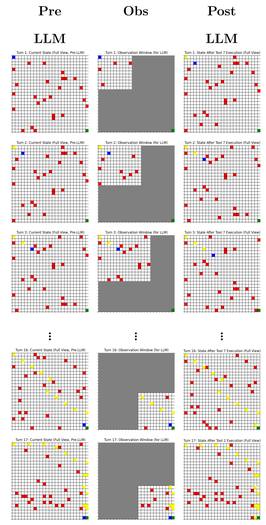

- LLM = OpenAI GPT-5.6 Sol (vision enabled model).
- Result: 3 x more success on target reaching than Sol 5.6 LLM (direct query only with same data).

References (quantization)

[1] Gabriel Turinici. Deep Conditional Measure Quantization. In Ana I. Pereira and al., editors, *Optimization, Learning Algorithms and Applications*, pages 343–354, Cham, 2024. Springer Nature Switzerland. Arxiv:2301.06907, doi : 10.1007/978-3-031-53036-4_24.

[2] Gabriel Turinici. Diversity in Deep Generative Models and Generative AI. In Giuseppe Nicosia and al., editors, *Machine Learning, Optimization, and Data Science*, pages 84–93, Cham, 2024. Springer Nature Switzerland. arXiv:2202.09573.

[3] Gabriel Turinici. Huber-energy measure quantization. *Statistics and Computing*, 35(1):5, November 2024.

Contact Information

- Web: `http://turinici.com`
- Email: `gabriel.turinici@dauphine.fr`