

AI for Simulation-Based Training

Lessons from applied collaboration in defense simulation*

Luís Alvaro de Lima Silva | Federal University of Santa Maria (Brazil)

The ELLIIT Focus Period Symposium on Robust AI for Science and Industry, Halmstad University, Sweden | September 2026

*The ideas and opinions expressed in this presentation are solely my own and do not necessarily reflect the views or official positions of other members of my research group or our project partners in the Brazilian Army.



AI

SIMULATION

TRAINING

DEFENSE

Who I am

**Applied AI researcher
working across simulation,
health, cybersecurity,
computer games, and
environmental applications.**

- ▶ Professor at the Federal University of Santa Maria – UFSM.
- ▶ Department of Applied Computing, Technology Center.



Federal University of Santa Maria – UFSM



A large public university in southern Brazil.

- ▶ Around 26,000 students.
- ▶ More than 300 courses.
- ▶ More than 2,000 faculty members.

Part I

The SIS-ASTROS project

UFSM and Brazilian Army: a long-term partnership

A simulation ecosystem for training with the Brazilian ASTROS Artillery System.

- ▶ Computer-Based Training software (CBT).
- ▶ Cabin/operational simulators for ASTROS artillery batteries.
- ▶ Virtual Tactical Simulator.
- ▶ Interoperability with other Brazilian Army simulators.



Prof. Raul
Distributed
systems



Profa. Usandra
Software
engineering



Prof. Pozzer
Computer
graphics



Prof. Edison
Artificial
intelligence



Prof. Mateus
Computer
engineering



Prof. João Baptista
engineering



Prof. Luis Alvaro
Artificial
intelligence


The value of simulation-based training

- ▶ Practice **tactical decision-making** before real-world operations.
- ▶ Improve **situational awareness** and **terrain understanding**.
- ▶ Train **coordinated deployment** and **command-and-control** procedures.
- ▶ Evaluate **routes, positions, and environmental constraints**.
- ▶ Support **instructor-driven challenges** and **After Action Review**.
- ▶ Reduce **cost, risk, and logistics** while enabling repeatable training.

Design challenges

- ▶ **Different users** (e.g., military users at different echelons) need **different levels of abstraction**.
- ▶ **Realism must serve learning**, not simply visual fidelity.


Ten-year evolution



The application matured - and the research and development questions matured with it.

From Live to Virtual and Constructive Simulation

- ▶ **Live:** real people operating real systems.
 - ▶ soldiers conducting a field exercise with **real vehicles**, radios, and equipment.
- ▶ **Virtual:** real people interacting with simulated systems and environments.
 - ▶ a crew operating a vehicle simulator or **officers planning and executing missions in a 3D tactical environment**.
- ▶ **Constructive:** simulated entities and higher levels of abstraction.
 - ▶ **commanders controlling battalions, brigades, or larger forces represented as simulated units** on a digital battlefield.




The Virtual Tactical Simulator bridges lower-level virtual simulation and higher-level constructive simulation.

SIS-ASTROS: more than a single simulator

An integrated simulation ecosystem

- ▶ Supports training at different levels: **technical**, **operational**, and **tactical**.
- ▶ **Connects multiple simulators** instead of treating each training tool as an isolated system.
- ▶ A central contribution is the **Virtual Tactical Simulator**, designed for Position Recognition and Occupation activities.




Virtual Tactical Simulation: the bridge between map and terrain

Three synchronized interfaces

- ▶ **Tactical Digital Table**
80-inch multi-touch 2D interaction.
- ▶ **Wall-TV**
146-inch 3D visualization of the terrain.
- ▶ **Instructor Terminal**
scenario configuration, monitoring, and intervention.

The trainee works primarily with the map, while the 3D environment exposes terrain details that may influence tactical decisions.




Training is more than executing a simulation

Instruction mode


- ▶ **Instructor** defines explicit **learning goals**.
- ▶ **Scenarios** teach procedures and tactical reasoning.
- ▶ The instructor may inject **Simulated Military Problems**.
- ▶ **Difficulty can be modified** while the simulation runs.

Exercise/Certification mode

- ▶ Targeted at more **experienced personnel**.
- ▶ Focuses on applying **already acquired skills**.




A project at the intersection of many computing areas



Software architecture: simulation as a modular platform

Unity 3D was extended through specialized software packages.

The architecture separates concerns ranging from AI and terrain processing to events, networking, and user interaction.




AI: navigation graph and autonomous displacement |
Terrain Engine: roads, rivers, vegetation, elevation, physics | Networking: synchronization

Why SIS-ASTROS became a research platform

Real training requirements generated research problems

- ▶ Large-scale **terrain** visualization and processing.
- ▶ High-resolution **imagery and multi-resolution** data management.
- ▶ Autonomous and **semi-autonomous** vehicle navigation.
- ▶ Hierarchical **pathfinding** over realistic terrain.
- ▶ Real-time **vegetation** and **physics** simulation.
- ▶ **Synchronization** across multiple simulator nodes.
- ▶ Human **interaction** through 2D and 3D interfaces.
- ▶ **Interoperability** among different simulation levels.




Part II

AI research generated by simulation requirements


Recurring problems

- ▶ 1. Scalable pathfinding:
Large terrains, real-time response
- ▶ 2. Terrain-aware planning:
Relief, slope, smoothing, plausibility
- ▶ 3. Learning-based heuristics: DNN-guided search with risk/topography



AI problem 1: scalable pathfinding in large realistic terrains

- ▶ **Very large terrain representations (e.g., 700 km x 350 km)**
- ▶ **GIS-derived features and obstacles**
- ▶ **Interactive simulator response time**
- ▶ **Simulated agents and routes**



1.1 Fast Pathfinding at Massive Scale

- ▶ **Realistic simulation environments** may cover about $55 \times 55 \text{ km}^2$.
- ▶ Global routes must be **computed fast enough to preserve interaction and immersion**; the project targets response times below 1 s.
- ▶ **Fine regular grids create very large search spaces**, while GIS-based terrains make some alternative representations difficult to construct.

How can the terrain representation reduce pathfinding cost without losing the detail needed for navigation?





Fig. 3. The irregular (quadtree) representation of a virtual terrain used in the experiments.

1.2 Hierarchical + Irregular Representation




- ▶ We propose a **quadtree-based navigation map** for large realistic terrains.
- ▶ **Only regions containing relevant obstacles are recursively subdivided**; open areas remain represented by larger cells.
- ▶ **The leaf nodes form an irregular grid**, reducing the number of nodes that the search algorithm must explore.

1.3 Two-Phase Hierarchical Pathfinding

- ▶ **Phase 1 — Coarse search:** A^* searches the initial coarse quadtree over the whole terrain.
- ▶ **Phase 2 — Local refinement:** only nodes belonging to the coarse path are subdivided to a finer resolution.
- ▶ **Neighbors are reconnected and A^* searches again** inside the reduced refined search space.




Refine only where the candidate route actually needs more terrain detail.




sample of the different phases of the hierarchical pathfinding algorithm. (A) The first phase: a coarse path is computed. (B) The second phase: a

1.4 Evidence on Small Realistic Terrains

- ▶ Experiments with 40%, 50%, and 60% blocked areas.
- ▶ The proposed PHQ was compared with non-hierarchical/hierarchical and homogeneous/quadtrees representations.
- ▶ PHQ kept almost all A^* searches within the 1 s target and no PHQ execution exceeded 10 s.
- ▶ PHQ also opens far fewer nodes as route distance increases.




1.5 Scaling to Larger Terrain



- ▶ PHQ scaled with route distance and kept the mean response below 1 s for paths up to about 40 km.
- ▶ Overall, **36.89%** of executions exceeded 1 s and 11.78% exceeded 10 s in the massive terrain.
- ▶ The main trade-off is some **loss of shortest-path quality** in sparse regions, but the speed gains compensate for it.

AI problem 2: paths must be plausible, not only short




- ▶ In training systems, realism is not decoration.
- ▶ It affects the trainee's interpretation of the scenario and the credibility of the simulated behavior.

2.1 Large-scale Topographic Pathfinding


- ▶ Agent-based simulations need routes that are compatible with real-world terrain topography.
- ▶ Shortest-only paths may cross risky slopes; **safest-only paths may become unnecessarily winding.**
- ▶ The challenge is to combine **response time, safety, and visual realism.**

How can we plan smooth, relief-aware routes without making search infeasible in large-scale maps?




2.2 Topography-aware $\text{HPA}\Theta^*$

- ▶ We propose the $\text{HPA}\Theta^*$, combining **hierarchical search with path smoothing during search.**
- ▶ The terrain is represented with a **QuadTree.**
- ▶ **Search first operates over clusters and then refines the concrete path inside the relevant regions.**
- ▶ The path cost combines **distance and navigation difficulty** associated with terrain inclination.



2.3 How Relief Enters the Path Cost


- ▶ Each QuadTree leaf node stores minimum/maximum height, walkability, and four normal vectors.
- ▶ The *normalNorth*, *normalSouth*, *normalEast*, and *normalWest* vectors represent slope by movement direction.
- ▶ The cost of crossing a node increases as the slope approaches the agent's safe movement limit.
- ▶ Thus, the route avoids steep regions without treating every elevation change as an absolute obstacle.



2.4 Smoothing: Myopic Line of Sight

- ▶ Θ^* smooths the route by **connecting the current node's parent to a visible neighbor**.
- ▶ We adapt this idea: **line of sight** also **checks obstructions and terrain slopes**.
- ▶ A direct connection is accepted **when its accumulated topographic cost is lower than the local alternative**.
- ▶ The result is **fewer artificial turns**, without ignoring relief-related risks.

Smoothing is guided by relief-aware path cost.




2.5 Low Cost with Acceptable Runtime


- Evaluation used **real-world SIS-ASTROS** terrains with different levels of topographic complexity.
- **HPA Θ^*** achieved the lowest path costs among the tested algorithms.
- The hierarchy strongly **reduced runtime** compared with the non-hierarchical topographic Θ^* .

Smooth, topographically safe paths with performance compatible with interactive simulation.

Terrain	HPA Θ^*	Topo. Θ^*	speedup
A, 27 km	16 ms	240 ms	15.0×
B, 27 km	137 ms	896 ms	6.5×
C, 27 km	200 ms	1117 ms	5.6×






AI problem 3: learning to guide the search



- ▶ Learn where search should go.
- ▶ Encode terrain and risk.


3.1 Risk-aware Pathfinding in Topographic Terrains

- ▶ Tactical agent simulations require routes that account for both **terrain elevation** and **risk exposure**.
- ▶ **Conventional A*** heuristics mainly estimate geometric distance and may ignore complex risk distributions.
- ▶ Multiple risk sources increase the difficulty of estimating the true **cost-to-go** and can enlarge the search space.
- ▶ The target is therefore **to compute tactically meaningful paths** while keeping pathfinding fast enough for simulation use.



3.2 Risk Representation

- ▶ **View Point Configuration (VPCs)**
- ▶ **Visibility** defines node-level risk: high visibility means **higher exposure in dangerous scenarios** and **lower cost in safe scenarios**.
- ▶ A **VPC** encodes the **active combination of viewpoints** using collision-free power-of-two identifiers.



Example: VPC 140 activates risk locations 4, 8, and 128.


- ▶ Path cost combines topographic distance and risk:

$$g^*(\pi) = \sum_i [d_{Euc}(n_i, n_{i+1}) + \psi r(n_i, n_{i+1})] .$$

3.3 Learned Heuristics + Iterative Heuristic Map


DNN heuristic correction

- ▶ Inputs: **current node** $S = (x, y, z)$, **goal** $G = (x, y, z)$, and **VPC information**.
- ▶ Two encodings are evaluated: $\mathbf{BI}_{input}$ (one binary input per risk location) and $\mathbf{VPC}_{input}$ (single encoded value).
- ▶ The network predicts two correction factors: CF_{α} for distance and CF_{β} for risk.
- ▶ These factors adjust the Euclidean heuristic.




Iterative Heuristic Map (IHM)

- ▶ DNN calls are expensive when executed **node-by-node** during A^* .
- ▶ IHM queries a **local neighborhood in batches**, stores the resulting heuristic factors, and reuses them through a cache.
- ▶ Only **spatially relevant nodes are evaluated**, reducing redundant GPU inference.



3.4 The Learned Heuristic Changes the Search


- ▶ In **dangerous** scenarios, the search is **guided away from visible/high-risk regions** while still considering path length.
- ▶ In **safe** scenarios, **paths are attracted to lower-risk regions** and the learned heuristic can restrict node exploration.
- ▶ The goal is not merely shorter paths, but a better **distance–risk trade-off** with fewer unnecessary expansions.



- ▶ **Learning improves the true topographic and risk-aware cost-to-go**, reducing the effective search space.

3.5 Main Results

- ▶ **Learned heuristics** reduce **opened nodes** and **execution time** relative to the risk-aware topographic A* baseline.
- ▶ In **unseen mixed VPCs**, the safe scenario achieves about **82.4% fewer opened nodes** and **80.1% lower execution time**.
- ▶ With **multiple simultaneous VPCs**, the safe scenario reaches about **84.6% fewer opened nodes** and **81.8% lower execution time**.
- ▶ Even in the harder dangerous multi-VPC case, **execution time is reduced by about 33%**.



- ▶ We propose a scalable framework that **combines risk modeling, DNN-guided heuristics, and cached batch inference for tactical topographic pathfinding.**

Part III

Concluding remarks

Case-study summary

Long-term applied collaboration

University + domain experts
+ real-world training
environment



Navigation, learning, autonomy

Pathfinding, terrain-aware
planning, learned heuristics,
etc for agent-based simulation



Simulation capability


Repeatable training and
knowledge capture



Industrial AI and simulation-based training

Reusable capability

- ▶ Real-world applications generate relevant research questions.
- ▶ Training AI is not game AI: educational validity matters.
- ▶ Build a simulation-based training ecosystem, not only a single prototype.



Applied collaboration and AI research can continuously reinforce each other.

Simulation-based training as a laboratory for industrial AI: scalable, interactive, human-centered, and mission-aware.

Contact: luisalvaro@inf.ufsm.br