Spatial AI

Columbia GSAPP
Fall 2026 A6596-1
William Martin

Final Submissions

The final submission for the course is a web application that contains a spatial simulation utilizing the core concepts introduced in the course.

Your **spatial simulations** should be able to:

- visualize a baseline simulation given a semantic model of a spatial idea
- encode at least two notions of location
    - occupants have a "location" relative to each system (frame of reference)
- encode at least two notions of distance with their own metric
- model various "spatial agents"
    - at least two kinds of spatial agents
    - demonstrate limited knowledge of their environments, but scan
    - should have an idea of spatial memory (they create landmarks they use subsequently to make better decisions later)
    - they should have at least one "goal" they are trying to achieve
    - there should be at least one notion of measurable "performance" for each type of goal
    - agents should demonstrate attempts to "plan ahead", either path planning given their limited knowledge or otherwise
- model an environment with the following characteristics:
    - physical vs abstract (the environment can be a virtual environment, like a game board or a social network)
    - partial observability (agents must not be able to sense the entire environment)
    - multi-agent (multiple agents that consume space)
    - your choice → dynamic or static
    - your choice → deterministic vs non-deterministic environment
- model environments demonstrating using these concepts:
    - obstacles (occupants that block the traversal of your spatial agents)
    - hindrances and affordances (regions of space that have an inherent penalty or bonus to the performance metrics of an agent)

- primary function → enable a user to generate multiple environment options
    - taking on the role of the environment-designer, this simulation tool should enable ways of viewing the dynamics of different environment arrangements
    - manual option → for example, a "choose your own adventure" UI flow that runs two simulations side-by-side and allows a user to select one or the other, then create two more variations from the selection.
    - automated option → automate the evolution of an environment by generating one, looking at the aggregate "performance" scores of its agents, modify the environment, then repeat. give this a number of "epochs" that represents the iterative count, the number of environments that were generated and evaluated.
- secondary function (at least one)
    - should incorporate real-world data, real-time or historic (perhaps served from an MCP server)
    - should represent true measurements of a real place in three physical dimensions + time

Deliverables

A slide deck that includes:

1. a URL to your simulation
2. a markdown file with your semantic model (can be hosted on GitHub as part of your code)
3. a series of diagrams that call out the above requirements, demonstrating that you have fulfilled them
4. an instruction-manual video (show how to run the simulation and generate a new environment)
5. a demo video (show the creation of an environment given a series of simulations)
6. a series of interesting environments you have generated and a visualization of the agent behavior that results from them → note in a short description for each why they are interesting