

Agent-Based Modeling of Urban Society and Interactions

Creating a Realistic City Simulation in Order to Model Infections

Andrew Imm

TJHSST Computer Systems Research Lab

2009-2010

Abstract

Current systems used to model the spread of disease treat populations as single entities, and neglect the actions of individuals. By developing an agent-based simulation focused upon the accurate modeling of social interactions seen in an urban environment, a testing bed that resembles a modern city arises. This testing bed --- with its accurate modeling of day-to-day interactions within a city --- provides a far better system to use when developing epidemiology simulations. Using an implementation of goal-oriented agents who are guided by a number of variables that make up their "personality," this program attempts to create this urban model and use the system to run a number of epidemiological studies. Once the virtual society is established within its routines and a network of social relationships has developed within the city's inhabitants, the simulation will reach an autonomous running state, where it can develop on its own. At this point, the introduction of a simulated influenza virus will determine how an urban population reacts to such an infection, and how the disease is likely to spread through the city. Various quarantine methods will also be tested in order to measure their effectiveness.

Background


In the field of epidemiology, most models used to predict the outcomes of plagues and epidemics are math-based. They treat the entire population of a region or nation as a single entity. This take on the problem of studying the spread of disease has one major downfall --- it assumes that all members of the population have similar behaviors. If any stratification is done to divide the population into subgroups, these are generally only related to susceptibility to the disease in the study. In other words, the unique characteristics of individuals are lost. An agent-based model, while more processor-intensive than a strict mathematical model, brings into play this individuality. However, past models that took an agent-based approach were very simplistic. For instance, viral modeling has been popular in the TJHSST Computer Systems Research Lab for years, but nearly every project has involved agents moving randomly within a closed, featureless environment. Effectively, these simulations resembled nothing more than an experiment of specialized bacteria moving around in a petri dish --- hardly an experiment that can be used to make generalizations or conclusions about a human population. For such conclusions, the agents in the model must act as humans do; this necessity provides the reason for developing an accurate simulation of an urban society.

The in-depth qualities of this simulation will provide an effective testing environment for epidemiology studies. Because the simulation is designed with a focus on individual interactions, the program works well for simulating the spread of disease from one individual to another. Using agent-based models to analyze the spread of disease is something that has been explored before by a few scientists, but the fact is that it is not a mainstream method of epidemiology modeling. Dr. Stephen Eubank from Virginia Polytechnic Institute is one of the leaders in the field of agent-based epidemiology modeling, and his projects explore the spread of many diseases in a variety of environments. For instance, his "Modelling Disease Outbreaks in Realistic Urban Social Networks" takes on a similar problem as my project does --- exploring the spread of disease through a social network. However, his program does not take into account all of the aspects of simulating a city that mine does. With the extra features found in my simulation model, I hope to create a platform that can accurately assess the quality of various quarantine methods when dealing with infectious diseases.

Description of Simulation

The simulation makes up the majority of the code written for this project. Initially, it loads a map file that tells the simulation how to construct the city. It then loads an agent file which tells the program how to configure the virtual city's population. Each agent is assigned a name, a schedule, and a "personality" --- a set of preferences that dictate how likely the agent is to perform various actions. Once the world and its inhabitants have been built, the program initializes its internal clock to 12:00 midnight on day 0. As the model runs, the virtual clock updates, and eventually agents wake up. As time progresses in the simulation, the agents go about the daily routines dictated in their schedules, navigating the city using the simulation's path-finding algorithm. Inherently, the agents encounter others throughout the day, and begin to remember other agents whom they often see. These memories of acquaintances are the beginning of the agent's social network: a stored list of friends and colleagues that allows the agent to keep track of people it has already met. The agent's list of acquaintances also keeps track of how well the agent knows others; this data is used by the agent to decide whom to interact with. As the simulation ages, the virtual city begins to resemble its real counterpart. Agents become established in their routines, and have dependable networks of friends that keep them socially active. At this point, a range of tests can begin in the simulation. Manipulation or addition of variables --- such as a virus --- at this stage ensures that the results resemble a real-world reaction as best as possible.

The map builder, featuring a graphic user interface.



Agents moving towards the green square in the center of the map.

Discussion

Currently, the simulation is at a stage where it can load a map and a file full of agent definitions. The simulation creates the agents as they are specified within the file, each with its own name, schedule, and personality attributes. The simulation keeps track of virtual time, and uses this clock to time and control the actions of agents. As it is, the agents within the simulation can continue to navigate the map indefinitely, moving to the various destinations indicated in their schedule. Although agents encounter each other during the day, the ability to interact with each other has not yet been implemented. At this point, there is still no data to be collected, since the simulation is not at a point where a virtual virus can be introduced. The other large piece of code is the map builder, which currently creates maps with far more features than those that are used in the simulation at this point. The map builder features a graphical user interface that makes creation of the map much easier than editing a text file by hand. Dialogs for creating buildings on the map, as well as defining new types of terrain allow for a feature-filled map creation environment.

Additional Programs

This project requires the creation of other programs that speed up the process of development. For instance, the simulation uses complex files to store maps, and the easiest way to create these maps is with a secondary program. The map builder allows the user to create maps with a graphical interface that displays the map as it will appear when the simulation is run. The program can also be used to create buildings on the map, and such buildings are used by the simulation in order to determine where to send agents. The map builder also features a variety of other features that can be useful for development, including distance calculations and map-printing abilities. While programs such as the map builder are not used in the final simulation, the products they create enable the experiments of this project, and these programs are therefore crucial to the completion of this project.

Tests

In order to improve the efficiency of the program and determine the optimal scale of the simulation, various tests will be used to analyze the program's internal algorithms. These tests will import methods from the simulation and run them on large sets of data to determine their practical limits. One algorithm that is very important to test is the path-finding method. This is one of the most frequently-called methods in the simulation, and it needs to be tested to determine how many times it can be run per program cycle before a noticeable lag occurs. Testing it again and again with large sets of data will help to determine this number. These tests are like the additional programs in that their code does not appear in the final project. Instead, they are used to develop and refine the simulation so that its final state is the optimal version.

