An Evolving Multi-Agent Scenario Generation Framework for Simulations in Preventive Medicine Education

Manan Gupta, Jeffrey W. Bertrand, Sabarish V. Babu
School of Computing
Clemson University
100 McAdams Hall
Clemson, SC – 29634-0974
Tel: 001/864-656-3444
{mgupta, jbertra, sbabu}@clemson.edu

Philip Polgreen
Department of Internal Medicine
University of Iowa
200 CMAB
Iowa City, IA – 52242-2600
Tel: 001/319-335-6707
philip-polgreen@uiowa.edu

Alberto M. Segre
Department of Computer Science
University of Iowa
14D MacLean Hall
Iowa City, IA – 52242-1419
Tel: 001/319-335-1713
alberto-segre@uiowa.edu

ABSTRACT
We describe the design, implementation and evaluation of a novel multi-agent scenario generation framework for interactive virtual reality simulations towards preventive medicine education. Our scenario generation framework is based on recordings of human movements from a distributed sensor networks deployed in a real-world physical setting. The components of our framework include the generation of unique virtual agent behaviors from the sensor data, and algorithms for the generation of low level or gross movement behaviors such as path determination, directional traffic flows, collision avoidance and overtaking. The framework also includes the generation of high level fine actions for multi-agents such as techniques for interactive activities in pedagogical scenarios based on environment and temporal triggers. We applied our multi-agent scenario generation framework in an interactive simulation for hand hygiene education, and conduct an initial usability study to assess the educational benefits of the simulation to nursing students and evaluated the performance characteristics of our framework. Results of our quantitative and qualitative evaluations suggest that our framework was robust in creating engaging, compelling, and realistic interactive training scenarios with multiple virtual agents in simulated hospital situations.

Categories and Subject Descriptors
H.1.2 [**Information Systems**]: User/Machine Systems – *Human Factors*, H.3.4 [**Information Systems**]: Systems and Software – *Performance evaluation (efficiency and effectiveness)*, I.3.7 [**Computer Graphics**]: Three Dimensional Graphics and Realism – *Animation, Virtual reality*, J.3 [**Computer Applications**]: Life and Medical Sciences – *Health*, *Medical information systems*, K.3.1 [**Computers and Education**]: Computer Uses in Education – *Computer-assisted instruction (CAI)*.

General Terms
Design, Experimentation, and Human Factors.

Keywords
Virtual Humans, Embodied Conversational Agents, Medical Virtual Reality, 3D Human-Computer Interaction, and Preventive Medicine.

IHI'12, January 28–30, 2012, Miami, Florida, USA.

1. INTRODUCTION
Multi-agent simulations are of great interest in applications such as historical reconstructions [21], architectural walkthroughs [16], pedagogical environments [20] and entertainment systems [10]. In these interactive simulations, autonomous agents need to not only locomote from one place to another, they must also perform complex actions that are a combination of low level motion behaviors and high level actions that are part of the interactive scenario. Low level motion behaviors include collision avoidance, overtaking other agents, maintaining traffic rules, and observing natural gait and stride. High level actions include natural multimodal verbal and non-verbal interactions with other agents and users, interactions with the environment, and capabilities of perceiving and reacting to their surroundings.

It is indeed a challenging computational problem in integrating these capabilities for agents, and generating realistic and high-fidelity interactive scenarios for simulation and training. Pedestrian simulations and crowd simulations are mostly capable of producing realistic and believable low level human movement behaviors, but lack the high level features of human activities and interactive scenarios with users. Our goal is the creation of a novel and evolving framework towards the generation of realistic activities for multiple agents in interactive simulations based on recorded human movements and activities in a similar real world setting. Our framework models realistic human activities and movements that combine low level gross motion behaviors and high level interactive activities in a compelling and engaging manner for multiple agents in complex pedagogical simulations towards enhanced medical simulations.

Pervasive distributed sensor networks are becoming more commonplace in monitoring human traffic for the purpose of general surveillance and public safety [6]. The use of recorded movements of human traffic in a real world setting has the potential to provide a basis for the realistic modeling of movements and activities in complex multi-agent simulations. We applied our framework for automatic activity and interactive scenario generation for multi-agents, based on data gathered from a distributed sensor network to record human movements in a hospital environment [7], towards a pedagogical simulation to interactively teach nursing students and practitioners in preventive medical hand hygiene procedures.

According to the Centers for Disease Control (CDC), healthcare-associated infections affect about two million patients in US hospitals each year [11]. Tragically, many of these hospital acquired infections are preventable. Despite the fact that hand hygiene is one of the most important measures for preventing healthcare-associated infections [2], hand hygiene rates among healthcare workers remain unacceptably low [5]. Interventions which include feeding hand hygiene rates back to healthcare workers can lead to improvements in hand hygiene practices [2], and improved rates can decrease healthcare-associated infections [2, 17]. Thus, measuring hand hygiene is an important component of infection control programs, and one that is recommended by both the Centers for Disease Control and Prevention (CDC) and the World Health Organization (WHO) [18].

To help educate healthcare workers about when to practice hand hygiene the WHO have defined a set of circumstances when hand-hygiene should be practiced using an alcohol rub or soap and water. These are referred to as the Five Moments. While there are short videos, posters and simple flash based demonstrations to describe the Five Moments, to our knowledge there are no interactive simulations for healthcare worker education that stress the importance of how observations should be done. It is believed that training infection control professionals to appropriately observe, record, and practice hand hygiene is critical. This work proposes that these observation skills would more effectively be learned in an immersive, interactive, engaging environment. Indeed, most of the studies reviewed in the recent Agency for Healthcare Research and Quality report on the effectiveness of continuing medical education suggest that interactive techniques are more effective than non-interactive ones [14].

Previously, we designed an agent based simulation to teach and train users in hand hygiene using pre-determined scenarios [1]. In an initial user study, we found that although users were able to learn the task of hygiene education, they found the simulation boring, non-engaging and non-responsive.

The key contributions of our work are as follows. We describe a novel multi-agent interactive scenario generation framework for training in clinical best practices, and present the results of an initial study evaluating an enhanced simulation on learner's usability criteria such as interest, engagement, interaction, responsivity, realism, believability and usability. Results of our user evaluation suggest that learners on average found our enhanced system for hand hygiene education based on our multi-agent scenario generation framework to be engaging, interactive, and believable. The performance characteristics of our framework revealed that we were successful in generating a rich and complex interactive multi-agent scenario in real-time.

2. RELATED WORK

In this section, we have summarized related methods for automatic generation of movements for agents in simulations. Sud et al. use multi-agent navigation graphs by finding second degree Voronoi diagrams for maximum clearances from the obstacles in the environment [22]. Ming et al., have proposed a multi-agent navigation graph along with Adaptive Elastic ROadmap (AERO) method for the environment to determine global path planning information for the simulation [15]. They also employ global collision avoidance methods and achieve high framerates towards simulation optimization. Karamouzas et al. presented a novel collision avoidance method using predicted time to collision in simulations of large number of agents [13].

Edward et al. developed a training simulation whereby the trainee observed virtual operators who performed risky maneuvers [4]. Cognitive models were developed for realistic virtual human behavior such as stress, tiredness, hunger, etc. Multiple agents exhibited autonomous behavior based on their surroundings, which in turn simulated a more realistic training experience. Human perception models for intelligent virtual agents (IVAs) have been developed by Herrero et al. that employ multiple levels of psychological coherence that maps to how real humans perceive the world [8]. Jur van den Berg et al. created large crowd simulations with algorithms for local collision avoidance and dynamic obstacle planning behaviors [23].

To our knowledge, most of the related work has focused on automatic generation of low level or gross motion behaviors for multiple agents, without any interactive behaviors with users, the environment or other agents. Our novel contributions are in the automatic generation of activities and interactive scenario for virtual agents based on recorded human movements and activities from distributed sensors in a hospital environment. The components of our framework include; movement space behaviors such as agent planning of paths in environments based on sensor logs, cognitive model of global and local paths, collision avoidance and overtaking, realistic motion behaviors, and interaction space behaviors with a model of integrating gross motion with interactive actions towards agent-agent and agent-environment activities in interactive scenarios. We apply our framework towards the creation of a rich and compelling multi-agent interactive training application for hand hygiene education, and have conducted an initial usability study to evaluate the user response and performance characteristics of our framework.

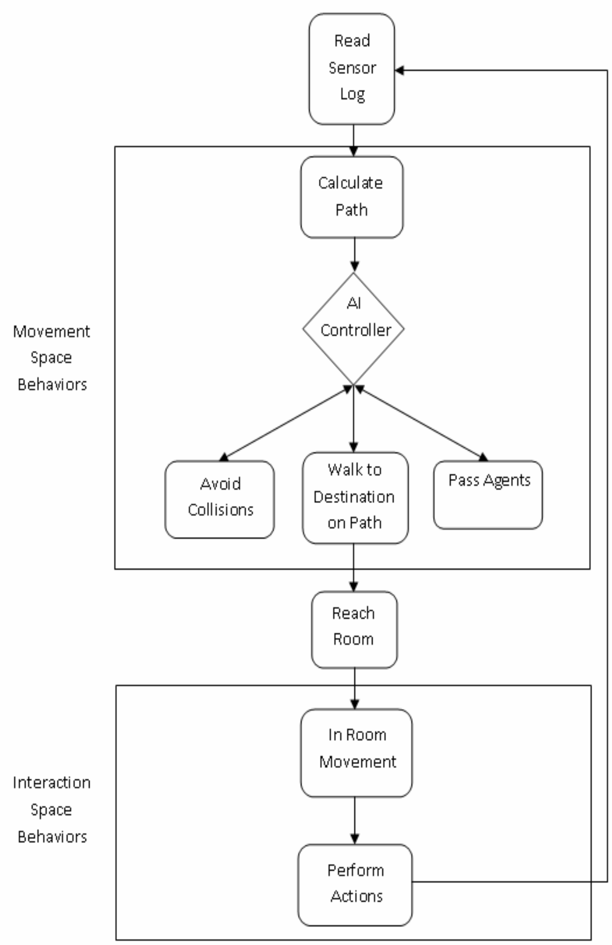

Figure 1. Flowchart depicts the various components of our framework, and the computational process from sensor log input to activity and scenario generation.

3. SYSTEM FRAMEWORK

Our framework was developed to be extensible to a variety of applications, in which agents' activities and movements need to be modeled in an engaging, compelling, and interactive manner from human movements recorded by embedded sensors in a real world setting. The various components of the framework are shown in Figure 1. In general the basic components of the framework work together in the following manner. Initially, sensor log data are read by the framework to determine the starting location and end location of all agents in a simulation instance. Then, the framework finds a path for the first agent from his current location to a destination. An AI Controller first initiates a set of *movement space* behaviors that correspond to walking, collision avoidance and overtaking in navigating the agent to his final destination. At the destination, the *interaction space* behavioral actions are initiated with seamless transitions in behaviors from one scenario space to another. The agent performs the activities in interaction space and then moves to the next action point via natural walking.

3.1 Tracking Human Movement Data Using Distributed Sensor Networks

The sensor logs of human movement were captured using distributed sensor networks at a hospital ward of the University of Iowa [7]. The distributed sensor system utilizes a network of non-RFID proximity sensors that log when a healthcare worker is within 5 feet of the sensor. Sensors were placed in the doorways of each patient room and the sensor logs consist of enter and exit events. Healthcare workers were equipped with a badge that triggered an event whenever she was in range of the sensors. The first application for the sensor network was to monitor hand hygiene compliance when a healthcare worker enters the patient zone defined as the area approximately 3 meters surrounding a patient. Healthcare workers should wash their hands before and after entering a patient zone. The effectiveness of the wireless network was determined by comparing the ground truth from a human observer to the sensor logs and the results indicate a 100% positive predictive value and a negative predictive value of 97% [19].

Our 3D model of the hospital environment was built to scale of the 4th floor of the Carver Pavilion at the University of Iowa Hospital using CAD diagrams. Photographs were taken of the ward to match equipment placement and textures. Realistic lighting and ambient occlusion were baked into the texture maps for faster rendering. The 3D model matched the real ward as closely as possible to create an engaging and accurate simulation.

3.2 Reading In Human Movement Data

Using the tracking system in [7], human movement data was logged over a period of two weeks. The log data was read in from a flat file, and each file consisted of human movement data captured over a period of 12 hours. The tracking log data allows for the use of real world data towards generating events in a virtual environment. A line of sensor log data is shown in Figure 2.

```
265,33,Day Doctors,ENTER,RCP.4104,PT BEDROOM
399,20,Day Nursing Staff,ENTER,RCP.4106,PT BEDROOM
496,27,Day Doctors,ENTER,RCP.4082,PT BEDROOM
592,33,Day Doctors,EXIT,RCP.4104,PT BEDROOM
624,33,Day Doctors,ENTER,RCP.4045,PT BEDROOM
704,27,Day Doctors,EXIT,RCP.4082,PT BEDROOM
```

Figure 2. Sample of the log data used to generate the activities in movement space.

The fields in the log file consist of a timestamp in seconds from start of simulation when an event occurs involving human movement, a unique ID for the individual, the status of the worker that is used to generate unique agents (Doctor vs. Nurse), the type of action performed (Enter vs. Exit), the location where the action occurred, and description of that location. The sensor log data is used to generate the virtual agents' traffic behaviors in movement space, and in the end locations that lead to patient rooms the agents then initiate interactive scenarios and behaviors in activity space.

3.3 Path Determination

A waypoint graph was adopted to facilitate an agent towards finding an optimal path from one location to another point. The waypoint graph allows an agent's path planner to identify paths that are traversable from paths that are obstructed. The waypoint graph nodes and edges are systematically placed within the environment. The nodes of the graph denote the intersections of paths, and the edges determine distance between the nodes. In our virtual hospital ward that was designed to resemble the real hospital ward where the data was collected, we used heuristics in specifying the nodes in the environment. They are as follows; at the center of every intersection between two or more corridors; in each alcove at a common point between entrances to rooms from the path, in other words, to treat this node as an intersection between paths to any number of doorways; at the center of each doorway to a room, so as to serve as a transition point into interaction space from movement space and vice versa.

A weighted directed cyclic graph was constructed of all the paths possible by linking waypoints that have a clear, unobstructed, straight path between them (Figure 3). Each of the edges in this graph was weighted by the physical distance in the 3D model between the two waypoints that were being connected by the edge. This data structure allowed for various algorithms for path determination from one room to the next. Virtual agents were assumed to be taking the shortest path between any two destinations. Therefore, Dijkstra's shortest path algorithm was used to compute an ordered list of waypoints to be reached in traversing the distances between them [3].

The entire environment has a waypoint based network of nodes and each agent has a planner directing them. The planner uses the agent's current position and heading direction in the environment to find the closest waypoint in the waypoint network that is in the general direction that the agent is facing in order to start the agent on the shortest path. Each agent walks in a straight line to the start waypoint of the path based on the assumption that the closest waypoint is in the line of sight of the agent.

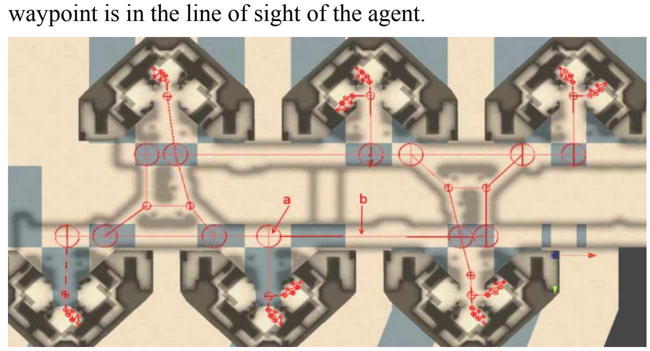

Figure 3. A complete waypoint network where the waypoints are placed at intersections of corridors. a) the position of the waypoints. b) the edges between waypoints that are weighted by the distance between them.

3.4 Directional Traffic

Once agents have the ordered list of waypoints to reach and the distances between them based on the shortest path, the agents begin to locomote to their destinations. The speed of walking is determined by the sensor log data. The logs contain the timestamps for a specific agent in exiting one interaction space and entering another. Hence, the speed of walking can be computed for the agent by using the shortest distance from the Dijkstra path between the end locations, and the time taken by the agent as indicated in the sensor logs to traverse the path. Based on the computed speed of walking, the walking animation of the agent is automatically manipulated in order to prevent foot slippage and to produce realistic and high fidelity motion behaviors. A linear relationship was determined between the walking speed of the agent and the playback speed of the walk animation that prevents foot slippage for the range of natural speed of human locomotion. This model was used to trigger the appropriate walk cycle animations of the agents, from the computed speed of walking.

Generally, the described method for waypoint node determination does not explicitly allow for directional traffic in movement spaces, since waypoints are placed at the center of each intersection. In order to not unnecessarily duplicate waypoints in corridors for producing directional traffic, a method was developed to allow virtual agents to locomote on an offset from the global path. The offset is determined to be half the shoulder width of the virtual agent. Each waypoint contains a property that determines whether it is placed in a corridor that allows for directional traffic or in between corridors as an environmental cue to the agent.

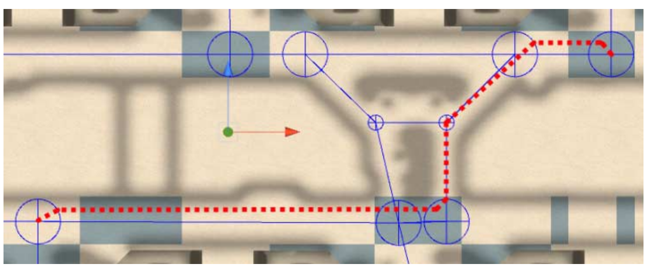

Figure 4. A virtual agent plans his path from a starting location to destination. The continuous lines are the planned paths while the dotted lines are the ideal path. The difference between the two is due to the directionality of traffic in corridors.

The directional walking algorithm for traffic flow consists of three phases. The agents have to transition from the normal path onto the offset path in an initial phase. In order to achieve this, the agent was considered to be walking at the center of a plane formed by the previous waypoint and the approaching waypoint that the agent is walking towards with the normal vector of the plane facing orthogonal to the direction of traffic flow. The offset path can be computed by translating the ideal path plane in the direction of its normal. In order to get on the offset path, the virtual agent walks in the direction of the offset path at a maximum angle of 45 degrees until they are on the offset path as shown by the algorithm below:

```
while SideOf(walkPlane, agentPosition) < 0:
      walkDirection = Normalize(o * d + p)
```

SideOf is a function that returns 0 when the agentPosition is on the walkPlane, a positive number if the agentPosition is on the same side as the normal of the walkPlane, and a negative number otherwise. 'o' is normal to the walkPlane facing the offset path, p is a vector parallel to the walking plane (heading direction), and d is the distance to the walk plane from the position of the agent (magnitude of the offset).

The second phase of directional walking for traffic flow generation is the agent's locomotion on the offset path. In order to determine the walk direction for the offset path, a vector is needed that is parallel to the walking plane. This vector can be computed with simple math since two intermediate waypoint positions are known which form a vector that is parallel to the walk plane; it can be computed in a trivial step.

An enhancement of the shortest path algorithm for agents was implemented in order to account for the offset paths in large corridors. To simulate start and end of smooth turns from the polyline path, each waypoint was given a hit radius. If the distance of the agent to the waypoint is less than the hit radius of the waypoint, the agent is considered to have reached the waypoint. Thus, by not traversing the waypoint before proceeding onwards, the agent can avoid unnatural merging behaviors from the offset path onto the global path (waypoint graph) and vice-versa (see Figure 4). The agent simply begins the turn early due to the large hit radius. The hit radii were determined such that the hit zone for each waypoint was the entire width of the corridor.

In the third phase of locomotion on the offset path, the agent then merges from the offset path onto the regular path where a similar algorithm to the first phase of merging onto the offset path was used. This time merging in the opposite direction from the offset path onto the regular polyline path.

3.5 Collision Avoidance and Overtaking

Collision avoidance algorithms were developed in order to prevent collisions of virtual agents on turns and in narrow corridors or corridors that do not allow for directional traffic. Collision avoidance was accomplished by simulating conical bounding volume in front of the virtual agent. If two agent's conical bounding volumes intersected, the virtual agent modifies its trajectory by turning to the right (Figure 5).

$$\boldsymbol{w} = \begin{cases} Normalize(\boldsymbol{n}+\boldsymbol{p}), Intersects(radarVolume, \boldsymbol{x}_{agent}) \\ Normalize(-\boldsymbol{n}*d+\boldsymbol{p}), otherwise \end{cases} \tag{1}$$

In this component, Intersects is a function that returns whether two bounding volumes intersect, $\boldsymbol{n}$ is the direction of traffic, d is the distance to the ideal path and $\boldsymbol{p}$ is the direction that the agent should walk to stay on the path. radarVolume is the bounding volume around the conical shape and x_{agent} is the position of the agent.

Since the sensor log data determines the speed of locomotion for an agent, it is possible that one agent moves faster than another agent and needs to overtake the agent. Passing behavioral component has three phases: approach, pass, and merge. In the approach phase, a fast moving agent comes up on a slow moving agent from behind. The collision avoidance algorithm prevents a collision between the two characters. At this point, the passing phase begins. The fast virtual agent increases speed to make a swift passing maneuver and moves further to the right from the path due to the collision avoidance algorithm. Once the fast agent is on the passing path, he continues to move parallel to the path until the fast agent is beyond the bounding volume of the slow agent. Then, the merge phase begins. The fast agent merges back on the path at a shallow angle to prevent a collision with the slow agent.

$$\boldsymbol{w} = \begin{cases} Normalize(\boldsymbol{n} + \boldsymbol{p}), Intersects(radarVolume, \boldsymbol{x_s}) \\ \boldsymbol{p}, SideOf(slowBoundingPlane, \boldsymbol{x_f}) \leq 0 \\ Normalize(-\boldsymbol{n} * d + \boldsymbol{p}), otherwise \end{cases} \tag{2}$$

The first branch of the procedure mimics the collision avoidance behavior. The second branch allows for the second phase of passing, the parallel walking maneuver, and the last branch allows for merging back onto the path. In the equation, $\boldsymbol{x_s}$ is the position of the slow agent, slowBoundingPlane is a plane that is formed at the front most point of the bounding volume of the slow agent with the normal to the plane being in the direction of walking dictated by the path, and $\boldsymbol{x_f}$ is the position of the fast agent.

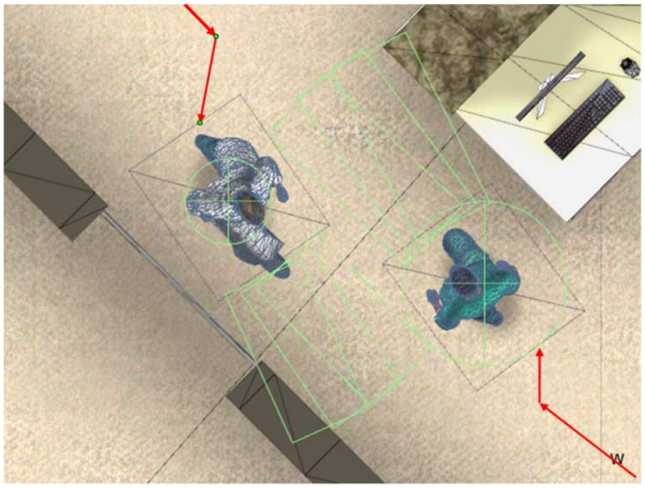

Figure 5. Two virtual agents avoid collision with each other in a narrow corridor that does not allow directional traffic. Their conical bounding volumes for intersection testing and collision avoidance are shown.

3.6 Interaction Space Behavioral Scenarios

Once the virtual agent reaches interaction space such as patient rooms, a different group of behaviors guide the agent's activities towards creating scenarios of the agent's interaction with the environment and other agents. The transition from movement space to interaction space is based on location triggers that are coupled to waypoints. Once an agent passes the location trigger that can be placed at any waypoint node, the agent enters interaction space. At the transition waypoints (W_T), the virtual agent starts moving at a speed that is more suitable for movement in activity space, i.e. natural walking speed. The same graph based algorithm used to convey agents' movement behaviors was leveraged to generate movements of virtual agents in interaction space. The waypoints inside the interaction space were placed systematically as follows (Figure 6): a waypoint was placed at every location, where a verbal and non-verbal action sequence could take place (W_A); a waypoint was placed at the center of all the activity waypoints as a means to get from one activity location to the next (W_I); other waypoint heuristics were applied from movement space as applicable for interaction space.

An agent's interaction space behaviors are initiated once they pass the transition waypoint W_T in the doorway of a patient room from movement space. The set of activities a healthcare worker agent can perform in the patient rooms are: checking the pulse of the patient, shaking the patient's hands, performing an injection, or typing on the computer by the patient's bed. The intermediate waypoints W_I in between the activity waypoints W_A are used as transit waypoints to guide the agents to the location of activities.

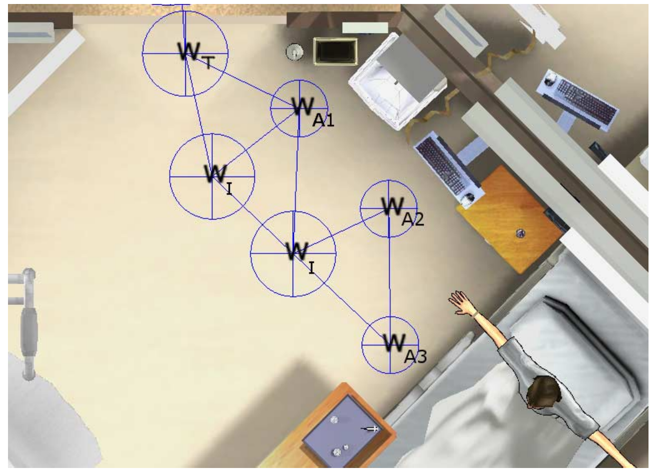

Figure 6. The graph of waypoint locations in interaction space showing waypoints at locations between transition from movement space to interaction space (W_T); at locations where verbal and non-verbal activities could be initiated (W_A); and at intermediate locations to the activity waypoints (W_I).

The set of activities were chosen to simulate each of the 5 Moments of Hand Hygiene [18]. Moments 1 and 4 are represented by a healthcare worker checking the pulse or shaking hands with a patient. The second and third moment requires the healthcare worker to wash hands before and after body fluid exposure risk or a procedure. A healthcare worker giving an injection to the patient represents moments 2 and 3. The 5th moment is to wash hands after interacting with the environment within a patient zone, and is represented by a healthcare worker typing on the computer by the patient's bed at location W_{A2}. The healthcare worker may or may not wash their hands upon entering or exiting the room at W_{A1}. The hand wash event is triggered 50% of the time at location W_{A1} for both entry and exit events. Only one in-room activity was generated for each enter event so trainees do not spend a lot of time observing multiple activities in a single patient room.

In an agent-environment interaction scenario, the virtual agent reaches a waypoint node in proximity to an object (W_{A1}, W_{A2}) where a sequence of corresponding animation actions pertaining to the interaction scenario (washing hands, or using the computer) are location triggered either synchronously or asynchronously. Ambient audio are also initiated to provide aural feedback of the agent's actions.

With interactions involving multiple agents, the synchronization of the inter-agent interactive scenario were executed by temporal triggers in addition to location and proximity triggers that initiated the verbal and non-verbal behaviors for the agents in a synchronous or asynchronous manner. An action state data structure stored the exact frames from the beginning of the action sequence when animations for the various agents should be triggered. This action state data structure can be populated from a state machine of actions pertaining to an interaction scenario for the agent. Our animation engine handles the execution, interpolation and transition of verbal and non-verbal behaviors for each of the agents from one action state to the next.

There are three inter-agent activities that require synchrony of animations between the virtual patients and healthcare workers. In order to properly synchronize the animations, agent animation actions were keyframed together in the same animation sequence, where the offset from the start of both sets of animation actions are identical. In this case, if a patient's handshake animation is called at the same exact time as the healthcare worker's animation, then they will appear in synchrony. The other problem that arises is the issue of world space orientation when the agents are in an interaction space. Once an agent arrives at W_{A3} the position and orientation of the waypoint is copied to the agent in order to correct their heading direction. An event manager calls the next activity at the same time for the interacting agents. The patient agents also exhibit a breathing animation that is interpolated smoothly along with the inter-agent activity that co-occurs as part of the interaction scenario.

When a virtual agent has completed performing a sequence of actions pertaining to an interactive scenario as defined in its queue of activities in the action state data structure, the agent is triggered, based on an end scenario activity trigger, to walk back to the transition waypoint (W_T) between the interaction and movement space. At this point, the location based trigger at the transition waypoint seamlessly enables the agent to transition from interaction space action scenarios to movement space behaviors. At the transition waypoint, the agent plans the route to the next destination based on the sensor log information, then proceeds to enact movement space behaviors accordingly.

4. Initial Evaluation of the Scenario Generation Framework

We conducted an initial high level performance evaluation of the interactive scenario generation framework for preventive medicine education, via a framerate analysis of the framework in multi-agent activity and scenario generation from recorded human movement data. The performance curve is shown in figure 7. The framework initial evaluation was performed on an Intel Core i7 CPU 920 @ 2.67 GHz with Windows 7 x64 for the operating system, a standard workstation in our laboratory. The system has 12 GB of RAM and a GeForce GTX 260 video card. The simulation was run with the specified number of agents and the framerate was measured for one and a half minutes at half second intervals.

A top-down (bird's eye) view of the entire environment was rendered with multiple agents **without any level of detail and culling implementations**. None of the graphics entities were abstracted. Agent appearance, clinical environment, and verbal non-verbal interactive agent behaviors were all rendered in rich details and both low level (movement space) and high level (interaction space) behaviors were generated by our scenario generation framework for an increasing number of agents. Although during interactive training learners only experience the simulation in first person perspective and may only encounter a few agents and other aspects of the simulation they do not experience are usually not rendered, we wanted to perform our evaluation with all aspects *turned on* so that we may objectively evaluate the limits of performance of the framework in extreme situations on a desktop workstation.

The graph in figure 7 shows the mean framerate for a minute of simulation time. The first 30 seconds of the simulation were not considered in the performance analysis to allow for the virtual agents to spread out in the simulation so that collision avoidance was not necessarily predominant over the other activities of the agents, and the measurements taken would be about an average case with the number of agents. The rendered frame rate was measured from a top down view of the entire scene. One factor contributing to the sheer drop in performance is due to the number of draw calls made to render all the polygons in the scene. The top down view prevented standard graphics optimizations from taking place such as visibility culling of polygons and level of detail techniques. In standard first person perspective experienced by trainees, the performance of this framework scales far better as agents are added. Rendered framerate in first person perspective were over 100 FPS with 10 virtual agents in the scene, in rich interactive verbal and non-verbal scenarios with other agents and the environment.

However, in future enhancements to the framework we are keenly interested in performing optimizations such as multi-threading where each agent's movement space and interaction space behaviors are computed in parallel on dedicated threads. We are also investigating optimizations via distributed processing on multiple graphics processor cores in order to leverage our framework to large scale real world behavior visualizations for feedback to hospital employees on preventive medicine compliance.

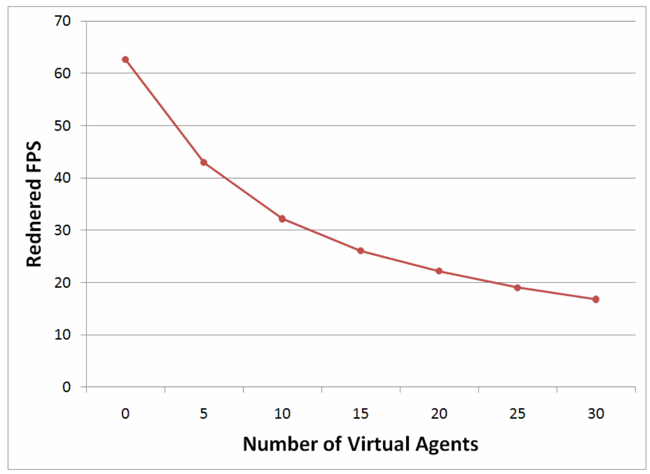

Figure 7. A graph of the rendering frame rate vs. number of virtual agents introduced in the framework. The data exhibits a power curve as the number of virtual agents in increased.

5. APPLICATION: ENHANCED PREVENTIVE MEDICINE TRAINING SIMULATION

In order to evaluate the effectiveness of the framework in providing realistic, compelling and engaging pedagogical scenarios with multiple agents, an interactive training simulation was developed based on our framework to educate users in patient safety procedures pertaining to hand hygiene.

The training simulation is divided into three phases, which are as follows:

The first phase (*introduction phase*) is an introduction to the 5 Moments of Hand Hygiene procedure [18], aspects of the game play controls, and directions regarding how the trainees will be scored on their performance. The game play scenario is designed such that the learner is told to follow agents in the virtual hospital ward, observe and record as many infringements of the hand hygiene procedure in the interactive scenarios between virtual healthcare workers and virtual patients, and the patient surroundings in five minutes. In this pedagogical scenario the trainee plays the role of a health inspector. Upon receiving instructions regarding the game play and the five moments of hand hygiene, the learner can start the training simulation. The initial instruction menu in the interactive simulation is illustrated in Figure 8.

Figure 8. Screenshot showing the game play menu, which provides instruction to the trainee regarding the training goals (5 Moments of Hand Hygiene), navigation and interaction controls, as well as the start interactive training button.

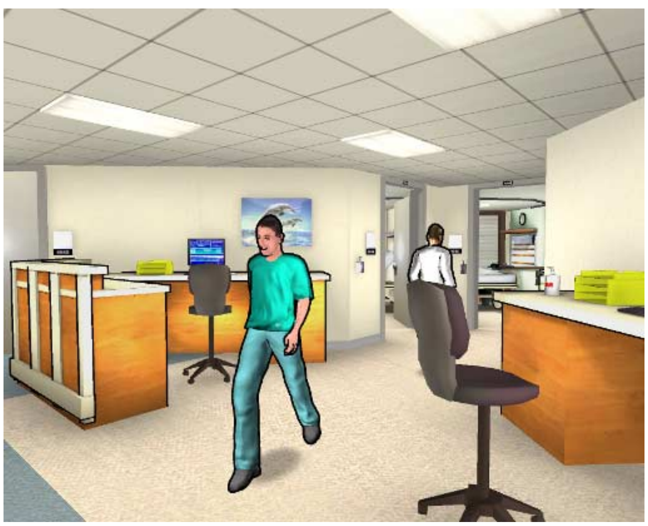

Figure 9. Upon receiving instructions regarding the simulation controls, in the Interaction Phase, the trainee is then allowed to follow any health care worker in the simulation in order to observe compliance with hand hygiene protocols.

In the second phase (*interactive phase*), the trainee navigates through the virtual hospital following virtual healthcare worker agents and attempting to catch as many non-compliant activities of the virtual healthcare workers in five minutes (Figure 9 and 10). This is done via interactively recording the observations on a virtual iPhone interface similar to an application used by real health inspectors at the University of Iowa Hospital [9] (Figure 10).

Figure 10. The trainee observes the interactive scenario of the HCW agent and the virtual patient, and records his observations regarding the agent's activities in the virtual iScrub interface [9].

The last phase (*feedback phase*) is the feedback screen where trainees can review a detailed performance record of how many observations that she correctly recorded. Upon receiving feedback, the trainee may prefer to revisit the interactive training phase for another training opportunity. In this case, the simulation engine automatically generates another instance of the interactive scenario by creating activities for virtual agents based on a different record (log data) of human movements collected on a different day.

Upon completing the training simulation, trainees were provided by a survey where they answered questions regarding their experience and impressions of the training simulation in preventive medicine education.

5.1 Usability Study and Results

The training simulation was implemented in Unity3D game engine, and was deployed online. Post training experience, students and faculty from the School of Nursing at Clemson University were requested to provide subjective rating of their perceived engagement, realism, believability, interactivity, ease of use, and effectiveness on a scale of 1(not at all) to 10(extremely). We also collected comments and suggestions regarding their impressions on each of the user experience issues and computational aspects listed above. We discuss our findings of this usability study in this section. The questions asked in the usability survey are listed in table 1.

Six full-time nurses and eighteen nursing students volunteered to evaluate the online training simulation and completed the usability survey. The quantitative results of the survey are show in figure 11. In all categories such as engagement, realism, believability, interactivity, ease of use, and effectiveness, users rated the system as above average on these qualitative aspects. In user feedback, it was encouraging to find that respondents generally agreed that the believability and realism was sufficient for training purposes.

Table 1. Questions posed after the simulation

1.	On a scale of 1(not engaging) to 10(highly engaging), please rate how engaging the hand hygiene training simulation is
2.	On a scale of 1(not realistic) to 10(realistic), please rate the realism of the training scenarios in the simulation
3.	On a scale of 1(not believable) to 10(highly believable), please rate the believability of the scenarios consisting of the virtual healthcare workers and patients
4.	On a scale of 1(not interactive) to 10(highly interactive), please rate the interactivity of the training simulation
5.	On a scale of 1(very easy to use) to 10(very hard to use), please rate the usability of the training simulation
6.	On a scale of 1(strongly disagree) to 10(strongly agree), please rate your impressions of how effective this training simulations of how effective this training simulation was in teaching proper hand hygiene protocols
7.	What are some medical procedures and best practices besides hand hygiene that can be taught via such interactive virtual reality training simulation?
8.	What did you like most about this training simulation?
9.	What did you like least about this training simulation?
10.	What are some ways in which this training simulation can be improved?

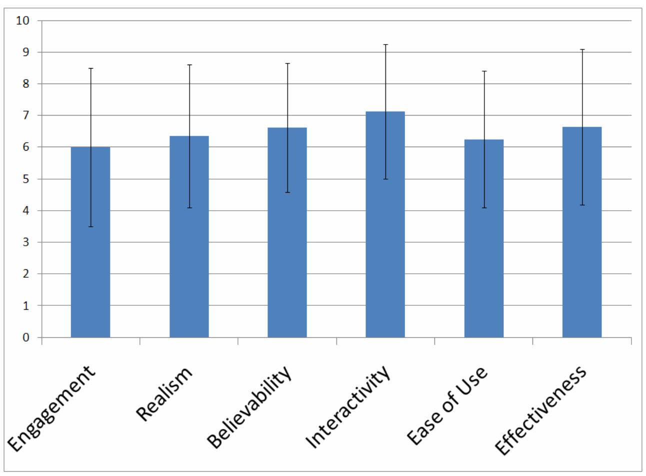

Figure 11. The results of the user study show the mean and SD of scores for the various usability aspects of the interactive preventive medicine education simulation.

5.2 Qualitative Assessment

In this section, we summarize the comments that we received in the initial usability study from nursing faculty and nursing students that participated in evaluating the enhanced hand hygiene training simulation. Most people found the simulation engaging and had suggestions on how to improve the realism and effectiveness of the simulation. The users suggested adding more activities for the healthcare workers to perform in the simulation. Some sample responses include:

"If it were able to perform tasks like removing dressings and isolation scenario that would make it even better."

"I think it would be cool to see more procedures and other times to wash hands. All I saw was before and after patient contact."

When asked about the realism of the training scenario, people commented that there needed to be more activity in the hospital ward. Even though we faithfully simulated an accurate replica of agent movements based on real world data, users noticed that there weren't as many people in the hospital as they would have liked to encounter. They wanted more workers, patients, and visitors to increase their immersion. Sample responses include:

"Need more people in the nursing station, halls, and activity…perhaps a group of professionals making rounds, secretaries in nursing station, etc."

"The basic hospital setup seemed pretty real, but the simulation needed more people present. Many of the rooms were empty, there were no visitors, and there were few employees present."

You would not find a nursing station without folks milling around, secretaries, and a variety of staff…nutrition, housekeeping, laboratory…depending on time of day. Also speed of activity needs to be increased."

Generally, we found that the subjective responses showed that the simulation was effective in engaging the learners and that increasing the activity and the number of agents in the hospital along with the variety of procedures performed can make the simulation highly realistic and engaging. As our scenario generation framework evolves and as our collaborators work on improving the tracking of the movements of activities of not only healthcare workers but also visitors and other staff members, we believe our simulation quality will significantly improve to include these rich and compelling scenarios with diverse agents. Finding other key preventive medicine tasks beyond hand hygiene training will also enable the application of this system towards the training of various other tasks in a clinical environment.

Users were also concerned about simulation and navigation controls in the interactive user experience. Volunteers found it difficult to navigate around corners using the key based navigation controls. The difficult controls coupled with the unfamiliarity of the complex layout of the hospital made it difficult for the users to form a mental map of the hospital and find healthcare workers at various times. Thus, our current work involves the addition of a dynamic map of the environment as a wayfinding aid in the simulation that could be rendered upon request (via a key press) as shown in figure 12. Dynamic maps provide an abstract overhead view of the environment along with the current position of the user in the environment via a dynamic "you are here" pointer, and abstract locations of patients and healthcare workers. Based on wayfinding research, we believe that the dynamic map will enhance user's spatial awareness and will enable them to quickly spot where the healthcare workers and patients are in their immediate surroundings.

In future work, we plan to introduce a training phase in the simulation where the trainee can gain interactive practice with the navigation controls prior to the interactive training phase where they actively apply and learn the preventive medical protocols pertaining to hand hygiene. This feature will allow trainee to acclimatize to the training environment and practice the controls prior to the training simulation.

Overall, the responses from the users suggest that the simulation was engaging, realistic and believable. The responses also suggest that the simulation can be an effective means to train healthcare workers in hand hygiene protocols once the concerns raised by the users have been addressed.

Figure 12. A dynamic map shows the trainee her current location and heading direction via the "you are here" pointer, and depicts the relative position of the nurses (green) and patients (brown) in the surrounding hospital environment.

6. CONCLUSIONS

In this research, we describe a novel evolving framework for the automatic generation of activities and scenarios for multiple agents in interactive simulations based on sensor logs of tracked human movements. The use of real world human movement data has the potential to enhance multi-agent simulations towards making them realistic, believable, and engaging. The key components of the multi-agent framework include algorithms and techniques for not only low level movement (gross) behaviors, such as global and local path planning, traffic generation, realistic human locomotion behaviors, overtaking, and collision avoidance. But also high level (fine) verbal and non-verbal activity generation for multiple agents, based on a model of location, proximity, and time based triggers of animation actions.

Our framework is extensible for developing multi-agent simulations for visualization of hand hygiene and other preventive medical procedure compliance in real hospitals, and creating pedagogical situations. Our scenario generation framework is only limited by an action sequence database of activities pertaining to the content of the scenario and the need to model the simulated environment to resemble the structure of the real world counterpart it simulates. We created an online application using our multi-agent simulation and sensor logs of human movements from a real hospital setting, gathered over a period of two weeks, to educate users in patient safety practices pertaining to hand hygiene. The results of our quantitative and qualitative evaluation suggest that this framework can be successfully applied for the creation of engaging, realistic, and compelling multi-agent interactive simulations and visualizations, in faithful virtual replications of real hospital environments.

As our collaborator's distributed sensor network systems evolves at the University of Iowa in sensing activities of real people in the clinical settings, we hope to leverage that data in our framework to automatically generate a database of realistic action sequences in interaction space for virtual agents.

In future work, we plan to automate the global path generation component of our framework, by integrating algorithms for the generation of paths automatically using environment geometry, by leveraging techniques proposed by Kallmann et al [12]. We also plan to conduct a comparative evaluation between our initial minimally interactive system of hand hygiene training vs. our current interactive simulation. We plan to conduct a second user study towards evaluating the effectiveness of wayfinding aids in trainee's spatial cognition and training task performance and learning. Lastly, we plan to extend our enhanced hand hygiene training simulation with a module that involves teaching healthcare workers in identifying patient safety infringements such as the presence of open needles, and lowered bed rails (patient fall risk) in the patients surrounding.

7. ACKNOWLEDGMENTS

We wish to gratefully acknowledge the assistance of our colleagues Dr. Tracy Fasolino and Dr. Nancy Meehan, and the students in the School of Nursing at Clemson University for their participation in our usability study.

8. REFERENCES

[1] Bertrand, J., Babu, S.V., Polgreen, P., and Segre, A. Virtual Agents Based Simulation for Training Healthcare Workers in Hand Hygiene Procedures. in *Lecture Notes in Computer Science: Proceedings of the 10th International Conference on Intelligent Virtual Agents (IVA 2010)*, Springer Verlag Berlin/Heidelberg, Vol. 6395, pp. 125-131.

[2] Boyce, J., and Pittet, D. Guidelines for hand hygiene in health-care settings: Recommendations of the healthcare infection control practices advisory committee and the hicpac/shea/apic/idsa hand hygiene task force. *Infection Control and Hospital Epidemiology* 23, 3–41.

[3] Dijkstra, E. W. A Note on Two Problems in Connexion with Graphs. *Numerische Mathematik*, vol. 1, 269-271.

[4] Edward, L., Lourdeaux, D., Barthes, J.P., Lenne, D., Burkhardt, J.M. Modeling Autonomous Virtual Agent Behaviors in a Virtual Environment for Risk. *The International Journal of Virtual Reality*, vol. 7, 13-22.

[5] Erasmus, V., Daha, T., Brug, H., Richardus, H., Behrendt, M., Vos, M., and Van Beeck, E. Systematic review of studies on compliance with hand hygiene guidelines in hospital care. *Infection Control and Hospital Epidemiology* 31, 283–294.

[6] Gomez, L., Laube, A., and Ulmer C. Secure Sensor Networks for Public Safety Command and Control System. *IEEE International Conference on Technologies for Homeland Security*, pp. 128-136.

[7] Herman, T., Pemmaraju, S., Segre, A., Polgreen, P., Curtis, D., Fries, J., Hlady C., and Severson, M. Wireless applications for hospital epidemiology. *Annual Scientific Meeting of the Society for Healthcare Epidemiology of America*.

[8] Herrero, P., Antonio, A. A Human Based Perception Model for Cooperative Intelligent Virtual Agents. *On the Move to Meaningful Internet Systems*. 2002, 195-212.

[9] Hlady, C., Fries, J., Curtis, D.E., Segre, A.M., Severson, M., Polgreen, P. iScrub Online: A System for Standardized Hand-Hygiene Observation and Feedback. *5th Decennial*

International Conference on Healthcare-Associated Infections.

[10] http://assassinscreed.us.ubi.com

[11] Jarvis, W.R. Selected aspects of the socioeconomic impact of nosocomial infections: Morbidity, mortality, cost, and prevention. *Infection Control and Hospital Epidemiology* 7, 552–557.

[12] Kallmann M. Shortest Paths with Arbitrary Clearance from Navigation Meshes. *Symposium on Computer Animation (SCA)*. 2010.

[13] Karamouzas, I., Overmars, M. A Verlocity-Based Approach for Simulating Human Collision Avoidance. *IVA 2010, LNCS 6356*. 180-186.

[14] Marinopoulos, S., Dorman, T., Ratanawongsa, W., Ashar, B., Magaziner, J., Miller, R., Thomas, P., Prokopowicz, G., Qayyum, R., and Bass, E. Effectiveness of continuing medical education. evidence report/technology assessment. *Publication 07-E006 Agency for Healthcare Research and Quality*, 149.

[15] Ming, L. C., Sud A., Van den Ber, J., Gayle, R., Curtis, S., Yeh, H., Guy, S., Andersen, E., Patil, S., Sewall, J., Manocha, D. Real-time Path Planning and Navigation for Multi-agent and Crowd Simulations. *Motion in Games, LNCS 5277*. 23-32.

[16] Narain, R., Golas, A., Curtis, S., and Lin, M. C. Aggregate Dynamics for Dense Crowd Simulation. *ACM Transactions on Graphics*. vol. 28, 122:1-122:8.

[17] Pessoa-Silva, C., Hugonnet, S., Pfister, R., Touvenaeau, D., Posfay-Barbe, K., and Pittet, D. 2007. Reduction of health care associated infection risk in neonates by successful hand hygiene promotion. *Pediatrics* 120:2, 382–390.

[18] Pittet, D., Allegranzi, B., and Boyce, J. 2009. The world health organization guidelines on hand hygiene in health care and their consensus recommendations. *Infection Control and Hospital Epidemiology World Health Organization World Alliance for Patient Safety First Global Patient Safety Challenge Core Group of Experts* 30, 611–622.

[19] Polgreen, P., Hlady, C., Severson, M., Segre, A., Herman, T. Method of Automated Monitoring of Hand Hygiene Adherence without Radio-Frequency Identification. *Infection Control and Hospital Epidemiology*, 31:12 (December 2010), 1294-1297

[20] Rizzo, A., Parsons, T.J., Buckwalter, J. G., and Kenny, P. A New Generation of Intelligent Virtual Patients for Clinical Training. *Proceedings of the IEEE Virtual Reality Conference*. 2010.

[21] Shao, W., Terzopoulos, D. Populating Reconstructed Archaeological Sites with Autonomous Virtual Humans. *Proceedings of the 6th International Conference on Intelligent Virtual Agents*. 420-432.

[22] Sud, A., Andersen, E., Curtis, S., Lin, M., Manocha, D. Real-time Path Planning for Virtual Agents in Dynamic Environments. *2007 IEEE Virtual Reality Conference*. 91-98.

[23] Van den Berg, J., Patil, S., Sewall, J., Manocha, D., Lin, M. Interactive Navigation on Multiple Agents in Crowded Environments. *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. 200.