Using simulation to model and understand group learning

Maartje Spoelstra
Dept of Computer Science
Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081 HV Amsterdam, NL
maartje.spoelstra@gmail.com

Elizabeth Sklar
Dept of Computer and Information Science,
Brooklyn College, City University of New York
2900 Bedford Ave, Brooklyn, NY, USA 11210
sklar@sci.brooklyn.cuny.edu

Abstract. We use simulation to model individuals participating in various group learning scenarios. By reviewing the pedagogical literature for key themes found in studies of skill acquisition and the design of learning environments, and their effects on individual learners, we have identified a set of characteristic factors that can describe a human learner. We use these characteristics to construct computational models that act as controllers for agents acting in a simulated learning environment. Varying parameter values can change the learning environment, as well as control some of the "human" factors that describe the population of agent learners instantiated in the simulation. The simulation can emulate the expected effects based on empirical and experimental results of education and developmental psychology research; and also gives us a simple environment in which to conduct low-cost, non-invasive experiments on the design of learning environments.

1. INTRODUCTION

The work described in this paper addresses the question of whether a simulation environment can be constructed which demonstrates different outcomes for learners who experience different learning environments. The primary purpose for constructing such a simulation environment is to be able, through simulation, to gain better understanding of learning environments and subsequently to use this knowledge to help design more effective environments for learning in the future. This goal is predicated on the ability to build software agents whose behaviours are controlled based on a set of parameterized "human" characteristics. Our aim is to define these characteristics by modeling computationally those factors that are considered important by pedagogical researchers. The method we have used in the work described herein is to explore the pedagogical literature, selecting key factors highlighted by empirical and experimental studies that have been conducted by researchers in the fields of education and developmental psychology. Our hypothesis is, that by carefully examining these prior works, we can develop computational models that represent different types of human learners, and we can instantiate these models as agents in a simulation that we can use to gain insights into the design of effective learning environments.

Here, we investigate these questions within the context of *group learning*, i.e., situations where students are placed in groups and given problems to address as a team. There is a large pedagogical literature on the topic of group learning, and prior to constructing our simulation, this literature was reviewed. The major themes are outlined below, beginning with some background on theories of individual learners and then continuing with discussion of group learning scenarios.

1.1 The process of learning.

Some of the most often cited work on theories of individual learners comes from Fitts, who studied the process of skill acquisition in adults who were learning to perform physical tasks. Fitts' [6] theory involves three phases of learning: an "early" phase, an "intermediate" phase and a "late" phase. In the early phase, the emphasis is on understanding instructions and on establishing the proper cognitive set for the task, resulting in a better grasp of the task at hand. The latter is done by performing a series of short, simple tasks and trials, like an introduction to the task to be learned. In the intermediate phase, people learn to associate parts of the skill they are acquiring with different stimuli. The late phase involves the perfection of the task learned.

Anderson [1] describes three similar stages in the context of the acquisition of cognitive skill. He names and explains the three phases slightly differently: the first phase is called the "cognitive" stage. A characteristic of this phase is verbal mediation, which enables the learner to clarify instructions for herself. The second stage is the "associative" stage, in which skill performance is "smoothed out": errors in the initial understanding are detected and overcome. In this phase, no verbal mediation is necessary anymore. The last phase is the "autonomous" stage, in which the learner gradually improves in performance of the skill. As a part of this stage, Anderson mentions the "procedural stage" which applies purely to the increase in speed with which the skills are performed. Taatgen [24] expands on Anderson's learning model and describes the outcomes of learning in terms of "explicit" and "implicit" learning. He uses the term "implicit learning" for unconscious and unintentional learning, whereas in "explicit learning", goals and intentions determine what is learned.

In our model, we do not distinguish between the two kinds of learning but rather use a combination of them. In the initial stage of learning, a large amount of new knowledge is introduced to the learner in a short amount of time, mostly in the form of instructions; hence the rate of a learner's progress is quite steep. In the second stage, the associative stage, instructions are formalized and made part of the learner's own skills; the learner's rate of progress decreases because it takes more time to formalize and associate actions with the new information and because the amount of new information that is presented also decreases. In the third stage, the autonomous stage, the learner does not learn new things but constantly elaborates the present knowledge.

1.2 Characteristics of learners.

In order to model the process of learning, it is important to take into consideration the human characteristics that contribute to that process. In many pedagogical studies, researchers distinguish between several levels of *ability* because some learners advance more quickly than others. Most common is to divide students into two levels of ability ("high" and "low") [2], but some studies mention three levels ("high", "medium" and "low") [15]. Another factor influencing learning behaviour is the level of difficulty of the information being processed in comparison to the level of development of the learner. The *zone of proximal development* is defined by Vygotsky as "the distance between the actual developmental level as determined by independent problem solving and the level of potential development as determined through problem solving under adult guidance or in collaboration with more capable peers" [28]. In order for a learner to process new information optimally, the level of the information should be such that it can be grasped within the learner's present zone of proximal development. *Collaborative* activity amongst learning peers promotes growth because peers are likely to operate within each other's zones of proximal development and interactions can help reinforce knowledge and smooth learners' transitions from the early to later stages of skill acquisition.

Throughout much of the pedagogical literature, three factors are cited as influencing individual human learning: *cognition*, *motivation*, and *emotion*. These are often referred to as the "trilogy of mind" [13]. All three elements influence the learning process equally. The zone of proximal development can be seen as the cognitive component of this trilogy. Motivation and emotion are factors that depend largely on one's learning environment and on interaction with others while learning. While these are both abstract values and difficult to measure quantitatively in the same way that cognitive progress can be measured by performance on a test, the notion of representing affect computationally has become popular within the agents and artificial intelligence communities over the last decade.

Following from [20], we model *motivation* as a value that indicates how much a learner tries to acquire new knowledge. If the learner has the cognitive ability to acquire a skill, s/he may choose not to for any number of reasons—all of which we label here as "motivation". We have all had students who do not perform to their capacity; the value *motivation* can be thought of as the noise that detracts from a student learning a skill that s/he should be able to acquire. From one timestep to the next, the value of motivation changes based on the level of challenge felt by the student. If a student is presented with challenging (i.e., difficult) concepts to learn, then motivation increases; otherwise it declines.

Again, from [20], we model *emotion* as a value that results from success in learning. If a student acquires a skill, her emotion value increases; if she does not, then emotion decreases. We could equate high levels of emotion with being "happy" and low levels with being "sad." In any case, a high value of emotion has a positive influence on learning; a low value has a negative influence.

1.3 Design of learning environments.

The design of an effective learning environment needs to revolve around one or more *goal structures*, which are a key part of the educational process within a classroom and can be focused on (1) individual, (2) cooperative and/or (3) competitive aspects. With *individual goal structures*, each student can set his or her own learning goals, regardless of the goals of others. With *cooperative goal structures*, students work together on a task. One inherent feature of this cooperation is that students only obtain their goal if the students with whom they work also obtain their own goals. If implemented correctly, the cooperative goal structure is generally believed to be beneficial for students' learning processes [7, 22, 2] because they not only learn the concept that is in fact the objective of their cooperation, but also the interactive skills necessary to cooperate. With *competitive goal structures*, students working individually can obtain their goal by scoring well in relation to others, even if others fail to achieve their goals and even if students block others' successes. Not always negative, competitive goal structures can be very motivating for some students [7]. The three goal structures all vary in the amount and type of interaction that takes place among learners: with an individual goal structure, there is no interaction; with a competitive goal structure, there are only competitive interactions; with a cooperative goal structure, interactions are designed to help all participants. One of the few methods that claims to be useful for all kinds of learning is Slavin's *STAD learning method* [22]. The STAD (Student Teams Achievement Divisions) method has five major characteristics: (1) teacher presentations, (2) student teamwork, (3) quizzes, (4) individual improvement, and (5) team recognition. We used the STAD learning method as the basis for the process of learning implemented in our simulation.

1.4 Simulation of group learning.

We have constructed a simulation environment which demonstrates varied outcomes for learners experiencing different environments. The behaviours of these simulated learners—software agents—are controlled by a set of parameters representing factors considered important by pedagogical researchers. Several environmental elements can be explored with our simulator:

- *Group Composition:* Within the cooperative goal structure, learners of both high and low abilities learn best in heterogeneous teams. Low ability learners are helped by high ability learners, and high ability learners gain understanding by helping other learners.

- *Group Size:* Larger teams provide more opportunities for simulated learners to learn from others; however, there are negatives, such as more opportunities for learners to dissent or compete.

- *Team Rewards:* Team rewards, shared equally by all members of a team, may have a positive influence on the motivation of learners within one team.

The remainder of this paper describes the simulation environment in detail and outlines experiments conducted to explore the elements listed above.

2. IMPLEMENTATION

We constructed our simulation using NetLogo[1] [30]. This is a visually-oriented, agent-based simulation tool, written in Java. It has its basis in Logo [5], a tool designed to teach

[1] `http://ccl.northwestern.edu/netlogo/`

young students about programming, but it has been developed into a multi-agent modeling environment for social and natural phenomena. NetLogo implements a parallel version of Logo in which instructions can be given to hundreds or thousands of independent agents all operating and interacting concurrently, making it possible to explore the connection between the micro-level behaviour of individuals and the macro-level patterns that emerge from the interaction of many individuals.

2.1 Learner parameters

Figure 1 shows an overview of the parameters that represent the dynamics within and between two learners. Each of these are described below. Note that for each, the range of possible values and the initial values are included in order to give the reader an idea of the scope of each parameter. These values were determined through trial and error while developing the simulation so that the combination of parameters used produces the desired effects reflecting the real life scenarios examined and discussed in the pedagogical literature. The exact values are arbitrary and only their relationships to each other are important.

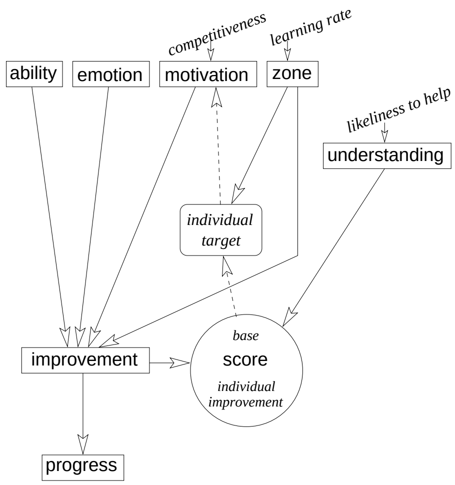

Figure 1: Overview of learner parameters. Solid lines indicate direct influence within a single timestep; dashed lines indicate influence from one timestep to the next.

Ability. As mentioned earlier, many pedagogical studies distinguish between learners of different levels of ability. Most common are two, three and four levels. In our simulation, we chose to implement two levels of ability: $1 = low$ and $2 = high$. Ability is an independent characteristic of the learner and therefore never changes while the simulation runs. In our simulation of individual and competitive goal structures, 50% of the learners have an ability of 1, whereas the other 50% of the learners have an ability of 2. When simulating the cooperative goal structure, these percentages depend on the compositions of the teams, which are set experimentally by the user (see section 3).

Improvement. The variable improvement is based on a general increase in knowledge throughout the learning of a new concept, in which the first phase of skill acquisition is an introduction to the concept; not very much new information is given. In the second phase, there is a peak in improvement per time unit, where the learner grasps the material and works hard to understand its details. In the third phase, the learner can elaborate on the concept, for example, by performing exercises that have to do with the concept. The rate of improvement per time unit declines, since nothing new is introduced. The shape of the improvement function, shown in figure 2a, is based on a normal distribution curve, with a different mean and standard deviation for students with high or low ability. The reason for this is that high ability students grasp information more quickly and may process more information at once, so their peak (the top of the curve) in improvement comes earlier in time and is higher than the peak in improvement of low ability students. Low ability students need a longer introduction to a new concept; they learn more gradually and their peak in improvement happens later in the process of grasping a concept.

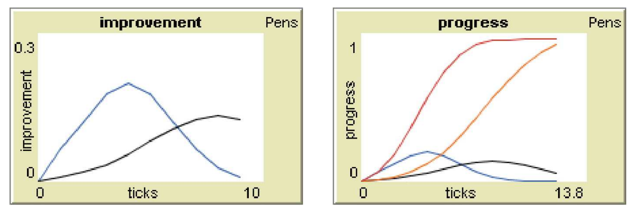

(a) improvement (b) progress

Figure 2: Improvement and Progress of simulated learners, over time. (a) The improvement of high ability learners are shown in blue (early peak); low ability learners are shown in black (later peak). (b) Progress is shown by the two top lines and improvement is shown underneath. Again, high ability learners peak earlier than low ability learners.

As discussed in the pedagogical literature, the rate of improvement of a learner while acquiring a new concept also depends on the learner's motivation, emotion and zone of proximal development. Therefore, in the simulation, *improvement* is multiplied by the values of $motivation$, $emotion$ and $zone$, and divided by $concept\_difficulty$. If the motivation, emotion and zone are optimal for a given learner, then improvement is maximal. For simulating the individual goal structure, this is the complete definition of improvement. For simulating cooperative and competitive goal structures, improvement is also influenced by help (given and received) and competition. These characteritstics are represented by the variables $likeliness\_to\_help$ and $competitiveness$ (explained below).

Progress. Individual progress is calculated as the total improvement of a learner per concept, as illustrated in figure 2b. Ideally for a learner, the cumulative value of improvement is 1, but since a learner's improvement depends on other factors (motivation, emotion and zone of proximal development), this is seldom the case.

Base Score. This variable represents the score of the quiz the student took before the first concept and after each concept, as defined in the STAD learning method and implemented as part of the *evaluation* phase in our simulation (described in section 2.3). Its initialization is the same for each goal structure. For high ability learners, this is a random value with range $[0.1, 0.4]$ and for low ability learners this is a random value with range $[0.0, 0.3]$. After completing a concept, the base score is replaced by the individual improvement score of the student; this is to simulate the quiz in the STAD learning model that each student takes individually after learners have finished learning a concept, i.e., the learner has been given time to grasp at least some portion of the concept that has been presented at the beginning of the learning process of the concept.

Individual Improvement Score. The individual improvement score reflects the outcome of the quiz taken by the learners at the end of each concept. This is a combination of the improvement of the learner during the concept and the gained understanding. The value of the individual improvement score ranges from 0 to 1. It is initialized to 0. The individual improvement score (iis) is calculated as:

$$iis = MAX(improvement) * u$$

where $MAX(improvement)$ is the maximum value of improvement during the previously learned concept and u is the understanding gained by the learner.

Understanding. A learner gains understanding when learning a concept that falls within the learner's zone of proximal development ($zone$). Each concept has a $difficulty$ value associated with it, that has the same numeric range as $zone$. If $zone - zone\_width < difficulty < zone + zone\_width$, then understanding increases. Another way in which understanding can increase is if a learner explains things to fellow team members in the cooperative goal structure. The amount of understanding gained depends on the current improvement of that learner and the help provided to others. Because one learner can provide help to more than one peer, the total help provided by the learner is divided by the number of learners assisted:

$$u+ = u + (help\_provided/num\_helped)$$

In our simulation, the increase in understanding is at the cost of that learner's improvement; therefore, it does not always pay off for a learner to help others. When each new concept is presented, understanding (u) is initialized to 0, indicating that the concepts are independent of each other. In other words, the understanding gained during the learning of one concept has no influence on the understanding of the next concept. Future work will explore dependent relationships between concepts and the effects of carrying understanding from one lesson to another.

Zone. This variable resembles the center of the zone of proximal development, which is described as a "frame". In the simulation, the variable $zone$ is a cumulative variable; its value increases after each concept is presented to a learner based on the value of $learning\_rate$. The zone of proximal development itself is defined by:

$$(zone - zone\_width) < zone < (zone + zone\_width)$$

We used $zone\_width = 0.15$ in our simulation experiments. The size of the zone frame stays the same throughout the learner's development; the entire frame shifts up with improvement, indicated by in increase in the value of the variable $zone$. Zone is initialized at the setup of the simulation for learners of both abilities as equal to the $base\_score$. (Because the simulation focuses on relative differences between students, the initial value of zone has proven experimentally not to make a difference in the outcomes of experiments.)

Learning rate. The learning rate expresses the average rate at which the student has learned per time unit and is calculated as:

$$learning\_rate = iis/ticks$$

This variable is used to indicate the overall development of the learner and, as above, is added to $zone$ after each concept is presented. Different students take different amounts of time to learn a concept. Typically, high ability learners "finish" learning sooner than low ability learners, but they can still progress and improve their understanding during the time units in which the slowest learners are finishing by helping those peers. That is why all students' learning behavior should be viewed in the context of how long it took them to grasp the concept.

Motivation. Motivation has a range of $[0.1, 1.0]$. It is initialized randomly, according to a normal distribution with a mean which is set to 0.5. (This mean is not variable, but gives the experimenter the chance to influence the general motivation more). In general, motivation depends on: (1) whether the difficulty of the current concept to learn lies within the learner's zone of proximal development; (2) whether or not the learner passed the quiz at the end of a concept; and (3) the motivation of a learner's teammates (in a cooperative learning scenario). In the case of a "failed" quiz, the learner becomes motivated to do better next time if the failure was only small, but demoralized (or a lot less motivated) if she failed by a lot. In the case of cooperative learning, if the motivation of a teammate is higher than the student's own motivation, the motivation of the teammate is decreased by 0.01; otherwise it is increased by 0.01. In addition, in a cooperative learning scenario, motivation is effected by the rank of the learner's team after taking their quizzes. If everyone on the team has passed the quiz, motivation increases by 0.1 for all team members. In the case of competitive learning, a learner's motivation increases if both the learner himself and her opponent have a competitiveness factor of more than 0.75 (see below).

Emotion. The range of emotion is $[0, 1]$. It is initialized randomly according to a normal distribution with a mean of 0.5. In general, emotion depends on how well the student performs on the quiz after progressing through a concept. For cooperative learning, apart from the individual performance on the quiz, emotion depends on the following: (1) emotion of the teammates: if the emotion of a teammate is higher than the emotion of the learner, the emotion of the teammate is decreased by 0.01; otherwise the emotion of the teammate is increased by 0.01; and (2) the rank of the team of the student after the quiz; if the team scores relatively well, the team members become happy (resulting in an increase of emotion); otherwise, they become sad (resulting in a decrease in emotion) or remain indifferent. In pedagogical literature, researchers have often remarked on the fact that in a competitive system, students tend to prefer others not to get benefits if they do not receive any benefit themselves, even if the students themselves would not benefit [7, 12]. This tendency of learners to try and stop others from achieving what they themselves cannot achieve led to the implementation of an increase in emotion when learners compete. If learners are close together in terms of zone of proximal development, they form a threat to each other and competing gives them a means to try and get ahead of each other.

Individual Target. This is the target for each individual for each concept. It is a personal goal only for the learners in the individual the cooperative goal structures (as explained in [22]); in the competitive goal structure everyone strives for the same goal. The individual target for the next concept is calculated according to the zone of the learner, the base score of the learner (or the grade the learner scored after the previous concept) and the concept difficulty of the next concept. It consists of a number that is slightly higher than the current value of zone of the learner. If, after progressing through the concept, the learner has not achieved a zone that is as high as the individual target, the learner fails the quiz. If the learner has achieved that level of zone, however, the learner has passed the test.

Likeliness To Help. The help others can give to a learner depends on their likeliness to help and is calculated as:

$$help\_provided = likeliness\_to\_help * improvement$$

This represents the amount of improvement that is subtracted from the helper's improvement, as that learner "stays behind" to help a peer. The lost improvement is invested in understanding. An important fact in our simulation (and in real life) is that the help provided by one learner and the help received by another learner is not necessarily the same. This is represented computationally as follows. The amount of help given by a high ability learner to a peer depends on that learner's improvement at that moment and the likeliness to help of the learner. The understanding gained from this interaction by the high ability learner depends on the actual help provided to others and the number of other learners helped, i.e., by explaining something to three other learners, the amount of understanding gained does not become three times higher than if the high ability learner would only help one other learner. The receiver of the help is also responsible for the cooperation: the effort invested in the learner by the other is equal to help provided, but is received according to the receiving learner's motivation. If the learner is unmotivated, the help cannot be fully appreciated. Another factor influencing the learners' cooperation is whether the help provided falls within the zone of proximal development of the learner receiving the help. If the help does not fall within the learner's zone, only a fraction of the help provided reaches the learner. In the simulation, this fraction is set to the arbitrary value of 0.5.

Competitiveness. Competitiveness is initialized randomly within the range $[0.01, 0.09]$. It is similar to $likeliness\_to\_help$, but it applies to negative interaction in a competitive goal structure, although learners that are very competitive might become motivated because of this competitiveness. In that case, competitiveness has a positive influence on the learning behaviour of the learner.

2.2 Learning environment

The learning environment in the simulation is designed to resemble a classroom in which students have to progress through a certain number of concepts, each with varying difficulty, within a time frame indicated by "ticks".

Concept Difficulty. The notion of concept difficulty is based on the representation of the knowledge domain described by Sklar and Davies [20], in which a number of related concepts—bits of information—are represented as a set of nodes in a graph. Each concept has a difficulty value between 0 and 1, in which 0 indicates the lowest level of difficulty and 1 the highest level of difficulty. Here, the notion of a concept is broader and represents a the amount of information comparable, e.g., to a topic in a geography class or a mathematical principle, in which practicing examples or sums is included in the time used to study the information. Each concept has a concept number; the number of concepts the learners have to progress through per run of the simulation can be initialized by the user of the simulation. The difficulty of each concept is divided into three levels: easy ($difficulty = 0.3$), intermediate ($difficulty = 0.6$), and hard ($difficulty = 0.9$).

Ticks. The learners have a certain amount of time to master each concept, measured in "ticks", i.e., the basic atomic unit of time in the simulation. Each of the phases in the learning process are measured in ticks. If time runs out while learners are still working, they have to stop and move on to the evaluation phase, after which they start a new concept. In practice, this is just to prevent the simulation from stalling. Future work will explore this aspect in more detail, since the phenomenon of students losing interest and their learning process stalling is only too common in real life and it would be useful to be able to emulate this situation in simulation.

2.3 The simulation

In a real classroom, the teacher initiates the learning process. The setting of instructional goals, goals for individuals or for groups of learners, and the evaluation after the learning process are all tasks that belong to the teacher. In our simulation, this role of the teacher is made a part of the environment: concept difficulties are set randomly, individual and group goals are set according to the previous learning outcomes of the learners, and evaluation is done according to the individual goals and progress of a learner. The teacher is represented implicitly in the simulation as an agent bearing a unilateral dependence relation with her students; i.e., the learners' behaviour depends on their teacher, but the teacher's behaviour does not change in response to the learners. This is also an example of the lecture model of teacher behaviour described by [20], though this could be expanded to encompass other types of teacher behaviours as well, like the lecture-feedback model and the tutorial model described in [20]. The activity in the classroom environment is divided into different phases, as shown in figure 3, through which the learners progress. For each concept presented, there are three phases: initialization, learning and evaluation.

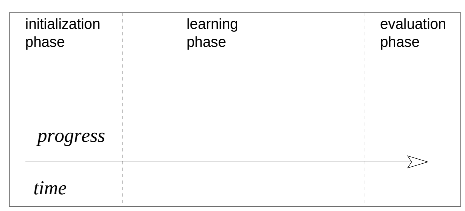

Figure 3: Phases of activity in the simulated classroom.

The initialization phase is the phase in which the necessary variables are prepared for the learning phase of the learner. This represents the phase in the STAD learning method in which the learners are given an introduction to the concept they are about to learn and their individual goals are set. At the very beginning of the simulation, the variable $base\_score$ is set in this phase according to the ability of the learner (after the first concept is learned, the base score is adjusted in the evaluation phase). The main feature of the initialization phase is anticipation to the coming concept to learn, which consists of the setting of the individual target according to the base score and zone of the learner, and the concept difficulty of the concept at hand.

The learners can then start the learning phase. In this phase, the learners progress through the concept for which they were prepared in the initialization phase. In other words, this is the phase in which learners are given the chance to learn; whether they actually learn depends on their motivation, emotion, ability and the other variables that interact with those three factors (as described earlier). The progress of the learners depends on the difficulty of the concept being presented, their own ability and their zone. In a competitive goal structure, this is the phase in which learners interact by competing. In a cooperative goal structure, the learners interact with each other in this phase by helping each other progress and influencing each other's motivation and emotion. As can be seen in figure 3, the learning phase is the largest phase; this indicates that this phase takes the most time to progress through. In the simulation, the horizontal distance travelled on the screen through the different phases is equivalent to the progress made by the learner.

In the evaluation phase, learners are put to the test: their progress after completing the learning phase is measured and, combined with their understanding, their progress is evaluated. Progress and understanding add up to the individual improvement score of the learner, to which their individual target is compared. The individual improvement score forms their new base score in the next initialization phase. If the individual improvement score is equal to or higher than the individual target, then the learner has passed the test. This results in an increase in motivation and emotion (the learner becomes "happy"), which will have a positive influence on the attitude of the learner towards the next concept.

In the cooperative goal structure, the evaluation phase is used for the calculation of "ranks" of the participating teams: in a situation where five teams interact, the two teams whose members score best on their individual quizzes are rewarded, having a positive influence on their motivation and emotion, and the lower scoring teams become disappointed, resulting in a negative effect on their motivation and emotion.

3. USER INTERFACE

The interface of the NetLogo application is shown in figure 4. There are three kinds of controls: sliders, switches and choosers with which the values or settings of several parameters can be adjusted, and the large buttons with which the simulation can be initialized (**setup**) and run (**go** and **step**). There are also two "monitors", which keep track of which concept that learners are currently progressing through and the number of ticks used by the learners per concept, indicating the time necessary for the learners to progress through the concept in the learning phase.

Before the simulation begins, the user can initialize the variables to the necessary values. As can be seen in figure 5a, the user can select the goal structure to use and the number of concepts can be chosen for the learners to progress through. If the user selects either *individual* or *competitive* goal structures, these values are all that need to be specified. If the user selects the *cooperative* goal structure, the number of teams, the team size and the composition of each team must also be set, as well as the use of team rewards (team rewards can be set on or off with the switch in the interface), as shown in figure 5b. Each possible team composition is numbered

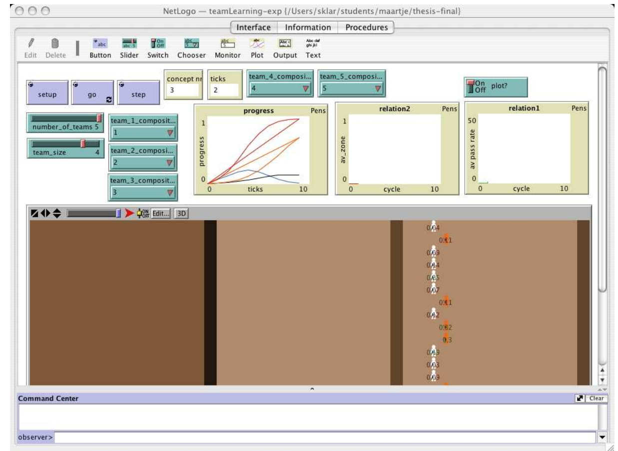

Figure 4: User interface of the simulation as implemented in NetLogo.

and the legend is shown in the interface (see figure 5c). A team with team composition 1 consists only of low ability learners; a team with team composition 2 consists of more low ability learners than high ability learners; a team with team composition 3 consists of an equal number of high ability and low ability learners (when the team size is an even number, otherwise the learners are divided into a team with composition 2 or 4); a team with team composition 4 consists of more high ability learners than low ability learners and a team with team composition 5 consists of only high ability learners.

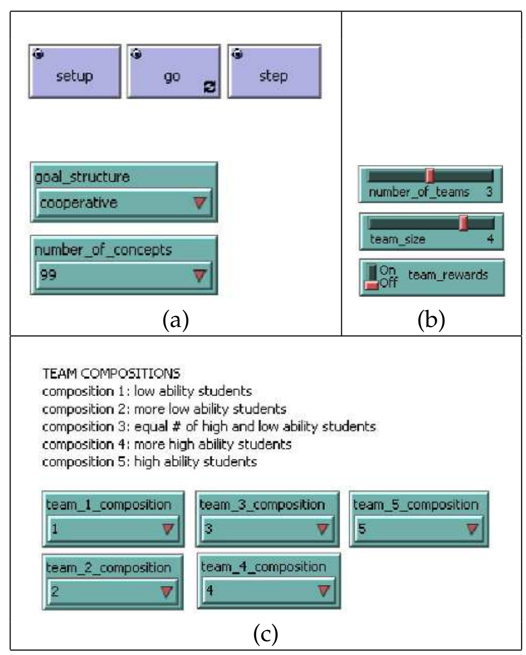

Figure 5: The user's controls in the simulator.

In experiments done with the cooperative goal structure, it is important to remember that in the simulation, team size and the possible team compositions are related to each other.

In any setting, there are two possible homogeneous teams: a team consisting of only low ability learners (team composition 1), or a team consisting of only high ability learners (team composition 5). The possible heterogeneous compositions, on the other hand, become more varied when team size increases. The simulation deals with the changes in team size and composition in the following way: in the cooperative goal structure, there are three possible team sizes: 2, 3 and 4. In a setting in which team size is 2, there is only one possible heterogeneous team composition, namely that consisting of one high ability learner and one low ability learner. For a team size of 3 learners per team, there are two possible heterogeneous team compositions: two high ability learners and one low ability learner, or two low ability learners and one high ability learner. For a team size of 4, there are three possible team compositions for a heterogeneous team. Table 6 shows all the different heterogeneous team compositions for each possible team size implemented in the simulation.

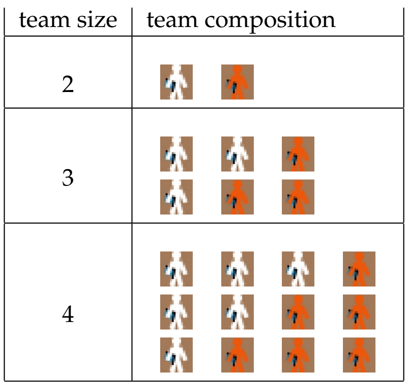

Figure 6: Team sizes and their possible heterogeneous team compositions

The lower portion of the screen is subdivided vertically, into three sections, each separated by dark vertical bars. This is illustrated in figure 7. Each of the three sections represents a concept to learn; the hues of the concepts indicate their difficulty: the darker the hue, the more difficult the concept. The order of concepts is initialized randomly, so students are not necessarily presented with easier concepts before seeing harder ones. The total number of concepts presented in the simulation varies, as set by the user. The three vertical sections represent three concepts and a scrolling window passes over all concepts in the simuation, three at a time. Each dark vertical bar separating two concepts stands for a combination of the evaluation phase for the current concept and the initialization phase for the next concept; within each dark bar, the progress of the learners achieved in the previous concept is evaluated and the variables for the next concept are initialized.

The lesson begins with all the students lined up on the far left side of the screen, ready to study the first concept presented. High ability learners are colored orange; low ability learners are colored white. All learners carry a blue book (but this is only decoration). As the simulation runs, the students move to the right, each arriving at the end of a concept at different times. After the learners have progressed through the right-most concept, the scrolling window shifts further to show the next three concepts and the learners continue on the left-most side of the screen. Each student travels horizontally within a "lane". With individual and competitive goal structures, equal numbers of high and low ability learners are distributed randomly over the vertical space in the learning environment. For the cooperative goal structure, the learners are divided into the teams and the lanes are shaded the same for all individuals in the team.

When the learners progress, they move towards the right with a speed that is proportional to their progress. When the learners have finished a concept, they stand still in the evaluation phase, indicating that the concept is stopped (i.e., the teacher interferes and says it is time to quit) and the learners are evaluated according to their individual progress and/or their team progress. They receive their individual improvement score, reflecting their individual progress relative to their individual target which was set at the beginning of the concept. If they managed to realise their individual target, that number is shown (on their "shirts") in dark green; if they did not succeed in making it to their individual target, this number is shown in dark red. The consequences of this failing are mostly motivational. A learner that had almost made it to her individual target will become more motivated to do well next time, since the learner only failed by a little. Only if the learner has not reached the individual target by far, will the learner lose motivation. Subsequent concepts are started by all students at the same time: they take a step forward into the next concept and the learning process starts again.

4. SIMULATION EXPERIMENTS

Figure 7 illustrates results of running the simulator for different combinations of settings selected by the user, for individual, competitive and cooperative goal structures. In figure 7a, one can see that for the individual goal structure, there is only a visible difference in progress between the high ability learners in general and the low ability learners in general. The only differences that can be noticed from the figure are the individual targets, which differ per learner. All their individual targets are written in red; this means that none of the learners progressed to their individual target after the previous concept. This is due to the level of difficulty of the concept they just learned; the dark shade of color indicates the highest level of difficulty, which, in this stage of their development (the screenshot was taken during the beginning of the progress of the learners) lies very much above the zones of both the high ability learners and the low ability learners.

In figure 7b, blue lines are shown when two learners compete. In the competitive goal structure, competition is the only mechanism of interaction added to the individualistic goal structure, and a clear difference can be observed in learning behaviour compared to the behaviour displayed by learners which learn individually; even though the learning capacity of high ability learners is defined to be higher than that of a low ability learner, we see that some low ability learners progress faster than some high ability learners. This shows that cognition isn't everything in progress: the factors of motivation and emotion play an equally large role in the learning process of all learners.

In figure 7c, the interaction among learners is most complex of all goal structures. Red lines are drawn from learners who help others to progress in their own team. The helping learner slows down in her own progress to help a team member to progress. The helping learner gains understanding by doing this, which will count towards the individual

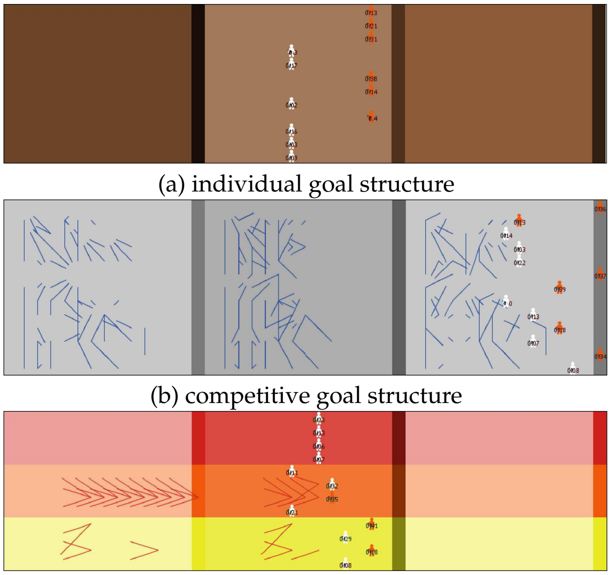

(a) individual goal structure

(b) competitive goal structure

(c) cooperative goal structure

Figure 7: Simulation experiments

improvement score in the evaluation phase. There are only three teams depicted in the figure, but there is a clear difference between the behaviour of the team members in different teams: in team 1, in the red lane, a homogeneous team of low ability learners is shown. All learners progress the same, since none of the learners has a higher ability than the others and none of them can help the others. In team 2, however, we see a heterogeneous team with one high ability learner. From the red lines in the first concept, it can be seen that the high ability learner has been helping all her teammates progressing. In the second concept, one of the low ability learners has been helped enough and can progress without help; in the second concept, this learner is helping the two other teammates that lag behind. Team 3 is a heterogeneous team with a different composition. The low ability learners in this team are progressing better than all the other low ability learners in the other two teams. As can be seen from the green numbers on the learners, many of the cooperative learners have developed to their individual targets after the previous concept was learned.

A series of experiments were conducted with the simulator (and fully described in [23]). From a comparison of experiments done with all three goal structures, it can be concluded that both high and low ability simulated learners display similar learning behaviours in the individual and competitive goal structures. When comparing these with the cooperative goal structure, it can be observed that most cooperative learning situations lead to an increase in the overall development of both high and low ability simulated learners. According to our results, the only situations in which the individual goal structure is more beneficial for high ability simulated learners are the heterogeneous teams of size 2 without team rewards, and for low ability simulated learners the homogeneous teams of size 2 with team rewards.

Due to the possibility of providing help to other learners, our results show that for both kinds of simulated learners, the effect of team composition is related to team size. When no team rewards are given, high ability simulated learners in homogeneous teams appear to learn best in groups of 2 or 3, whereas for heterogeneous teams, the larger teams seem to be more effective for high ability learners. For low ability simulated learners in an environment without team rewards, small heterogeneous teams appear to be best. Team rewards have a different effect on teams of different team sizes and compositions. High ability simulated learners benefit from team rewards in heterogeneous teams of size 2 or 3 or in homogeneous teams of size 4; low ability simulated learners benefit from team rewards only in a heterogeneous team.

5. DISCUSSION

We have presented the design of our environment in which simulated learners display different outcomes in varied learning scenarios. Earlier related work describes "SimEd", an environment that emulates interactions between simple artificial learners and abstract knowledge domains [20]. Students and teachers are modeled as agents acting within a complex social system, namely the education system; and their behaviours are controlled by features such as emotion, motivation and ability [13]. Here we have expanded upon this line of work in two main ways: we have modeled peer-to-peer interactions (the earlier work only modeled student-teacher interactions) and we have based the details of the simulation on existing research on "group learning". Thus the models of human learners presented here are grounded in empirical and controlled experimental studies well-documented in the developmental psychology, education research and/or cognitive science literature. Our work is related to the fields of cognitive modeling and user modeling; however the goal here is not to build or augment an intelligent tutoring system but rather to build a simulation system in which we can explore the interplay between various characteristics of learners and the environments in which they progress.

Our approach differs from other work that describes "simulated students". VanLehn *et al.* [25] present an analysis of machine learning systems that behave like human students, identifying two inputs of such systems (a student's knowledge prior to the learning event that will be simulated and the instructional intervention that led to the learning event) and two outputs (the student's behaviour during and updated knowledge after the learning event has occurred). Subsequent work employs this notion for analyzing skill acquisition, for example emulating learning from error correction [16]. Uses for systems that simulate students can be grouped into three categories [25]: *teacher training* [4, 3], *peer tutoring* (where the peer is a simulated student) [27], and *instructional design* [26]. Peer tutoring is the most closely related to the work described here.

One popular approach to peer tutoring is the use of *pedagogical agents* [8, 10], personalized assistants that interact directly with a learner and explicitly guide her through a domain. Recent work in this realm has focused on *interactive pedagogical drama* [9, 14, 18], where animated pedagogical agents become actors in a pseudo-theatrical environment and learners either become immersed as participants in the drama or act as observers. Typically, pedagogical agents consult a *student model* in order to understand the student and provide feedback that encourages the learner within her appropriate *zone of proximal development* [28]. Another approach to peer tutoring is the use of *peer learning agents* [21], the explicit use of agents as interactive partners in the learning process itself [11, 17, 19]. These agents are built into the user interface and, as with pedagogical agents, have knowledge of the user. While these agents may have teaching capabilities, they are less engineered for overtly guiding learning than pedagogical agents.

Our approach is informed by work in the areas of pedagogical and peer learning agents, but we take an abstract approach to knowledge, since our long term goal is not to build a tutoring system but rather to construct a simulation framework, based on social science research, designed to demonstrate and predict systemic effects caused by various characteristics of learning environments.

Webb and Palinscar [29] wrote the following from a pedagogical perspective:

> "...the research on the effects of group composition on group processes and learning outcomes shows that the makeup of a collaborating group has profound implications for the experiences of the students in it. It also shows that determining the optimal assignment of students to groups is no easy matter. Groups can vary on so many variables simultaneously that it is difficult to unravel the relative impact of each one."

Whereas all experiments performed in the field of pedagogy try to show causal relations between certain factors and learning outcomes, it is important to keep in mind that it is always a human being who has to set up the experiments and interpret the results. A bias is easily implemented, even subconsciously, into an experiment or into the interpretation of empirical results. This can be proven by the outcomes of the experiments of some publications; some results contradict each other, possibly due to a slightly different setup of an experiment, but the competitive nature of the presentation of results published in some papers shows some friction in the personal interpretations as well.

The variety of the different findings from the simulation experiments performed in the current research is proof of many dynamic interactions within the model, reflective of the many interacting factors in a real-life learning scenario. The many interacting variables and variety of experiments also show that in simulation, much more data can be generated than in a pedagogical experiment, and likely more objectively. By grounding the simulation in pedagogical theory and proven educational models, our hope is that better understanding of characteristics of learning environments and their effects on students may be gained in a less socially intrusive way.

Acknowledgments

This work was partially supported by NSF #ITR-02-19347 and by PSC-CUNY #68525-00-37.

6. REFERENCES

[1] J. R. Anderson. Acquisition of cognitive skill. *Psychological Review*, 89:369–406, 1982.

[2] M. Azmitia. Peer interaction and problem solving: When are two heads better than one? *Child Development*, 59:87–96, 1988.

[3] A. H. Brown. Simulated Classrooms and Artificial Students: The Potential Effects of New Technologies on Teacher Education. *Journal of Research on Technology Education (formerly Journal of Research on Computing in Education)*, 32(2):307–318, Winter 1999.

[4] S. C. Daum. Using simulated students to improve residents' teaching. *Acad Med.*, 72(5), May 1997.

[5] W. Feurzeig, S. Papert, M. Bloom, R. Grant, and C. Solomon. Programming-languages as a conceptual framework for teaching mathematics. *ACM SIGCUE Outlook*, 4(2):13–17, April 1970.

[6] P. M. Fitts. Perceptual-motor skill learning. *Categories of human learning*, 1964.

[7] D. W. Johnson and R. T. Johnson. *Learning Together and Alone: Cooperation, Competition and Individualization*. Prentice-Hall, New Jersey, 1975.

[8] W. Johnson. Pedagogical agents for virtual learning environments. In *International Conference on Computers in Education*, 1995.

[9] W. L. Johnson. Pedagogical Agent Research at CARTE. *AI Magazine*, Winter, 2001.

[10] W. L. Johnson, J. W. Rickel, and J. C. Lester. Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *International Journal of Artificial Intelligence in Education*, 11, 2000.

[11] Y. Kim. Agent learning companions: Learners' expectations of the desirable characteristics. In *Agent-based Systems for Human Learning, AAMAS Workshop*, 2005.

[12] A. Kohn. *No Contest: The Case Against Competition*. Houghton Mifflin, Boston, 1986.

[13] R. S. Lazarus. Cognition and Motivation in Emotion. *American Psychologist*, 45(4):352–367, 1991.

[14] S. Marsella, W. L. Johnson, and K. LaBore. Interactive Pedagogical Drama. In *Fourth International Conference on Autonomous Agents*, 2000.

[15] G. Mugny and W. Doise. Socio-cognitive conflict and structure of individual and collective performances. *European Journal of Psychology*, 8:181–192, 1978.

[16] S. Ohlsson. Learning from performance errors. *Psychological Review*, 103:241–262, 1996.

[17] K. Sehaba and P. Estraillier. A multi-agent system for rehabilitation of children with autism. In *Agent-based Systems for Human Learning, AAMAS Workshop*, 2005.

[18] M. Si, S. C. Marsella, and D. V. Pyndath. Thespian: Using multi-agent fitting to craft interactive drama. In *International Conference on Autonomous Agents and Multi-Agent Systems*, 2005.

[19] E. Sklar. *CEL: A Framework for Enabling an Internet Learning Community*. PhD thesis, Department of Computer Science, Brandeis University, 2000.

[20] E. Sklar and M. Davies. Multiagent simulation of learning environments. In *Fourth International Conference on Autonomous Agents and Multi Agent Systems (AAMAS-2005)*, 2005.

[21] E. Sklar and D. Richards. The use of agents in human learning systems. In *Fifth International Conference on Autonomous Agents and Multi Agent Systems (AAMAS-2006)*, 2006.

[22] R. Slavin. *Cooperative Learning: Theory, Research and Practice*. Alley & Bacon, Boston, 1990.

[23] M. Spoelstra. Simulating the Effects of Goal Structures in Human Learning Environments. Master's thesis, Artificial Intelligence Section, Department of Computer

Science, Faculty of Sciences, Vrije Universiteit Amsterdam, The Netherlands, 2006. Advisors: Martijn Schut and Elizabeth Sklar.

[24] N. A. Taatgen. *Explicit Learning in ACT-R*, chapter 3.1, pages 233–252. Pabst, Lengerich, 1999.

[25] K. VanLehn, S. Ohlsson, and R. Nason. Applications of simulated students: An exploration. *Journal of Artificial Intelligence in Education*, 5(2):135–175, 1996.

[26] M. Virvou and K. Manos. A Simulated Student-Player in Support of the Authoring Process in a Knowledge-Based Authoring Tool for Educational Games. In *Third IEEE International Conference on Advanced Learning Technologies (ICALT'03)*, 2003.

[27] A. Vizcaino and B. du Boulay. Using a simulated student to repair difficulties in collaborative learning. In *Proceedings of the International Conference on Computers in Education*, 2002.

[28] L. S. Vygotsky. *Mind in society: The development of higher psychological processes*. Harvard University Press, Cambridge, MA, 1978.

[29] N. M. Webb and A. S. Palinscar. Group Processes in the Classroom. *Handbook of Educational Psychology, 3rd edition*, 1996.

[30] U. Wilensky. Modeling nature's emergent patterns with multi-agent languages. In *Proceedings of EuroLogo 2002*, Linz, Austria, 2002.