in Agent Based Systems for Human Learning, International Transactions on Systems Science and Applications, 4(1), 2008.

Using Simulation to Model and Understand Group Learning

Maartje Spoelstra[1,2] and Elizabeth Sklar[3]

[1]Dept of Computer Science, Vrije Universiteit Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, NL
[2]TNO Defence, Security and Safety, Oude Waalsdorperweg 63, 2509 JG The Hague, NL
[3]Dept of Computer and Information Science, Brooklyn College, and Dept of Computer Science, The Graduate Center, City University of New York, 2900 Bedford Ave, Brooklyn, NY, 11210 USA
Email: maartje.spoelstra@tno.nl,sklar@sci.brooklyn.cuny.edu

Abstract: In this research, simulation is used to model individuals participating in various group learning scenarios. By reviewing the pedagogical literature for key themes found in studies of skill acquisition, the design of learning environments and their effects on individual learners, a set of characteristic factors are identified that can describe a human learner. These characteristics are used 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 a simple environment in which to conduct low-cost, non-invasive experiments on the design of learning environments.

Keywords: multi-agent simulation, education, group learning

1. Introduction

The work described in this article 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 eventually to use this knowledge to help design more effective environments for learning. This long term goal is predicated on the ability to build software agents whose behaviors are controlled based on a set of parameterized "human" characteristics. The aim of this research is to define these characteristics by modeling computationally those factors that are considered important by pedagogical researchers. The method used in the work described here 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. The hypothesis is, that by carefully examining this body of research, computational models can be developed that represent different types of human learners, and these models can be instantiated as agents in a simulation that can be used to gain insights into the design of effective learning environments.

Earlier related work describes "SimEd", an environment that emulates interactions between simple artificial learners and abstract knowledge domains [21]. Students and teachers are modeled as agents acting within a complex social system, namely the education system; and their behaviors are controlled by features such as emotion, motivation and ability [13]. Here, this work is expanded upon in two main ways: first, peer-to-peer interactions are modeled (the earlier work only modeled student-teacher interactions) and second, the details of the simulation on existing pedagogical research are based 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. The work described in this article 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 with which the interplay between various characteristics of learners and the environments in which they progress can be explored.

The approach of this research differs from other work that describes "simulated students". VanLehn *et al*. [26] 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 behavior 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 [26]: *teacher training* [4, 3], *peer tutoring* (where the peer is a simulated student) [28], and *instructional design* [27]. 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* [29]. Another approach to peer tutoring is the use of *peer learning agents* [22], 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.

In this article, results are presented of simulated *group learning*—situations where students are placed in groups and given problems to address as a team. A range of different groupings are explored, based on student ability, as well as the effects of providing group-based (or only individual) rewards for progress. The article is organized as follows. First, some background is provided by discussing some of the pedagogical literature on human learning. Then, in section 3, the implementation details of the group simulation environment are presented. Section 4 describes the experiments that were run using the simulation and provides results. The last section provides a discussion and directions for future work.

2. Background

There is a large pedagogical literature on the topic of group learning. This literature was reviewed to create a basis for the construction of the simulation. The major themes are outlined below, beginning with some background on theories of individual learners and then continuing with discussion of group learning scenarios.

2.1 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. In the late phase, the task learned is perfected.

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 towards an asymptote. 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 [25] expands on Anderson's learning model and describes the outcomes of learning in terms of "explicit" and "implicit" learning (see figure 1a). He uses the term "implicit learning" for unconscious and unintentional learning, whereas

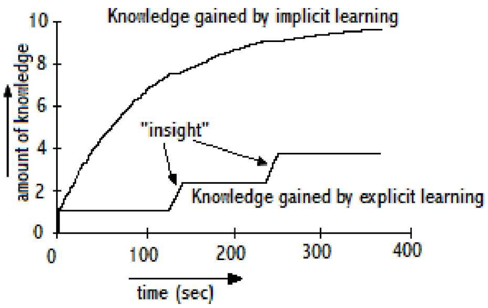

(a) implicit and explicit learning models, taken from [25], based on [1]

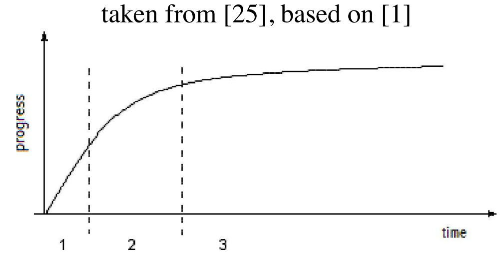

(b) stages of cognitive development combined with the stage-theory of [6] and [1], inspired by [25]

Fig. 1. Models of knowledge acquisition during learning, i.e., "progress". The horizontal axes represent the passage of time; the vertical axes represent the amount of knowledge acquired by the learner. In figure (b), "1" represents the initial stage of learning; "2" is the associative stage; and "3" is the autonomous stage.

in "explicit learning", goals and intentions determine what is learned. Some sources make a distinction between cognitive and affective outcomes of learning [7]. In an educational system, one can say that explicit learning gives rise to the cognitive outcomes of goal structures and implicit learning gives rise to the affective outcomes.

The model presented here (see figure 1b), does not distinguish between the two kinds of learning but rather uses 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 (or 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.

2.2 Characteristics of human 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 categorize students into two levels of ability ("high" and "low") [2], but some studies mention three levels ("high", "medium" and “low”) [15]. In the current study, the focus lies on two levels of ability.

Another factor influencing learning behavior is the level of difficulty of the information being processed in comparison to the level of development of the learner. The current level of development of a 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” [29]. 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 among 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 [21], *motivation* is modeled 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 are labeled here as “motivation”. Some students 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., experienced as difficult, near the upper bound of the zone of proximal development) concepts to learn, then motivation increases; otherwise it declines.

Again, from [21], *emotion* is modeled 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. High levels of emotion could be equated 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.

In a real classroom, there exists a wide variety of factors inside and outside of the learning environment that influence the motivation and emotion of the learner. In the group learning simulation, these factors are approached in an abstract way by initializing the variables *emotion* and *motivation* as random numbers, thus creating a difference between the mood and eagerness of simulated learners from the start.

2.3 Design of learning environments

The design of an effective learning environment revolves 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, 23, 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. Competitive goal structures do not necessarily have a negative influence on learners; they 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.

Goal structures can be implemented in different ways, according to how an instructor wishes to use them to help teach concepts and motivate her students. One teaching methodology that can implement all of the aforementioned goal structures is the STAD learning method [23]. The STAD (Student Teams Achievement Divisions) method has five major characteristics, which can be implemented as 4-5 sequential phases in the learning process that, collectively, are performed iteratively:

1. *teacher presentations*—the initial phase of the learning process in which a teacher explains the concept to be acquired;
2. *student teamwork or individual work*—the phase in which activities designed to facilitate learning are undertaken by one or more students, working alone or in groups;
3. *quizzes*—the phase in which the teacher evaluates the progress made by each student;
4. *individual improvement*—the phase in which individuals receive recognition (from the teacher and/or their peers) for any progress they have made; and, optionally,
5. *team recognition*—the phase in which teams are ranked and “prizes” (or some other form of recognition) are bestowed upon team members—this phase is only relevant when the “cooperative goal structure” is in place and students are working in teams.

A typical feature of the STAD learning method is that before learning a concept, students are each given individual “targets” to reach, customized according to their ability. Because these targets are personalized, every student has as much chance of performing well on her quiz as her peers do on theirs. Team recognition is based on collective performance as well as individual performance relative to personalized targets. This means of assessing progress and determining rewards was used in the current research to simulate the learning of a series of concepts by groups of students in an environment with various reward structures. The STAD learning method was used as the basis for the process of learning implemented in the group learning simulation.

The remainder of this article describes the simulation environment in detail and presents the results of experiments conducted to explore the elements listed above.

3. Implementation

A simulation environment was constructed which demonstrates varied outcomes for learners experiencing different environments. The behaviors of these simulated learners—software agents—are controlled by a set of parameters representing factors considered important by pedagogical researchers, some of which were highlighted in the previous section. Several environmental elements can be explored with the group learning 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 group learning simulation was built using NetLogo[1] [31]. This is a visually-oriented, agent-based simulation tool, written in Java, and is a powerful environment for developing experimental and/or prototype systems [20]. It has its basis in Logo [5], a tool designed to teach 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 behavior of individuals and the macro-level patterns that emerge from the interaction of many individuals.

3.1 Learner parameters

Figure 2 shows an overview of the parameters that represent the dynamics within and between two learners. Each of these are described in this section. 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 have proven in base line experiments not to influence the emergent properties of the learning behavior of the simulated learners that are the objective of this research.

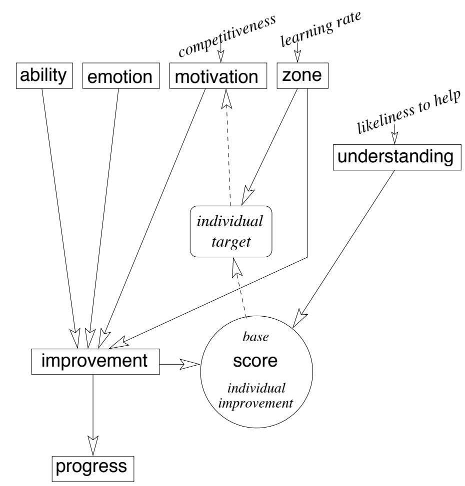

Fig. 2. Overview of learner parameters. Solid lines indicate direct influence within a single timestep; dashed lines indicate influence from one timestep to the next.

Ability. As mentioned earlier, many pedagogical studies distinguish between learners of different levels of ability. Most common are two, three and four levels. In the group learning simulation, two levels of ability were implemented: $1 = low$ and $2 = high$. Ability is an independent characteristic of the learner and therefore never changes while the simulation runs. In the 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 3.4).

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 3a, 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.

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*

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

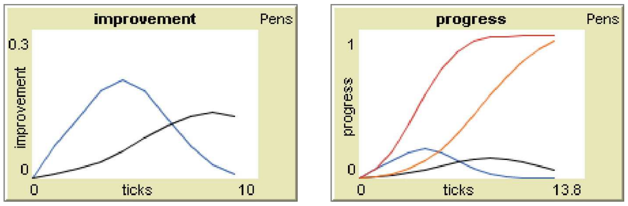

(a) improvement (b) progress

Fig. 3. 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.

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 3b. 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 the simulation (described in section 3.3). 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 the group learning simulation, the increase in understanding is at the cost of that learner's rate of improvement; therefore, it does not always pay off for a learner to help others. Within the fixed lesson time of the scenario used in the simulation, by helping a peer, a learner will spend some "ticks" helping their peer and have fewer ticks available to improve their own score; however, this type of peer tutoring will deepen the learner's own understanding in the long run. 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:

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

In the simulation experiments described here, $zone\_width = 0.15$. 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 [23]); 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 the 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 behavior of the learner.

3.2 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 [21], 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.

3.3 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 the 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' behavior depends on their teacher, but the teacher's behavior does not change in response to the learners. This is also an example of the lecture model of teacher behavior described by [21], though this could be expanded to encompass other types of teacher behaviors as well, like the lecture-feedback model and the tutorial model described in [21]. The activity in the classroom environment is divided into different phases, as shown in figure 4, through which the learners progress. For each concept presented, there are three phases: initialization, learning and evaluation.

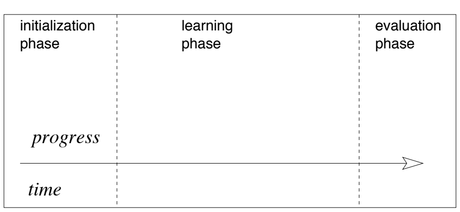

Fig. 4. Phases of activity in the simulated classroom.

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 4, 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.

The objective of the conducted experiments was to measure rates of development of different kinds of learners in different learning situations. To this end, the change in zone was monitored for each kind of learner, within each kind of goal structure, group size and composition. The average change in zone depends on an agent's learning_rate. In the evaluation phase, at the end of a concept, the learner's zone is incremented by the learning_rate, which incorporates improvement_score, which, in turn, encompasses understanding, motivation, emotion and the value of zone after the previous concept to be learned. In this way, all variables that are mentioned in the pedagogical literature influence the learning behavior of each agent. The differences in ability are implemented by using a different mean and standard deviation for a Gaussian curve describing possible improvement for each level of ability.

The algorithm implemented for simulating learning centers around the variable improvement, which is modeled as a curve (following figure 1b), using a normal distribution, with a different mean and standard deviation for learners with high or low ability, indicating the general increase in knowledge throughout the learning of a new concept.

Cognitive development is measured as the change in zone, which is calculated by adding the learning rate to the present value of zone. The value of learning_rate is calculated as improvement per tick, and improvement is determined mathematically as:

$$f(\mathsf{tick}, \mu, \sigma) = \frac{1}{\sigma \cdot \sqrt{2} \cdot \pi} \; exp \left(-\frac{(\mathsf{tick} - \mu)^2}{2 \cdot \sigma^2} \right)$$

where μ and σ are the mean and standard deviation, respectively, of the improvement curve.

As suggested by the literature, improvement depends on the trilogy of mind: cognition, motivation and emotion. Cognition is represented by zone; the rest of the trilogy is represented directly as motivation and emotion. Another variable that influences improvement is concept difficulty. These factors are combined and used to modulate improvement:

$$\mathsf{improvement} \cdot = (\mathsf{motivation} \cdot \mathsf{emotion} \cdot \mathsf{zone}/\mathsf{difficulty})$$

This indicates that if motivation, emotion and zone are optimal, then improvement is maximized. For the individual goal structure, no other variables contribute to improvement. For cooperative and competitive goal structures, improvement is also influenced by help (given and received) and competition, respectively, through the variables likeliness_to_help, understanding and competitiveness.

3.4 User Interface

The interface of the NetLogo application is shown in figure 5. 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.

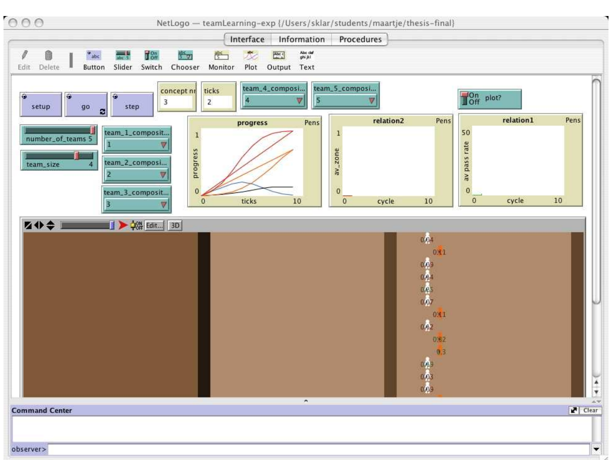

Fig. 5. User interface of the simulation as implemented in NetLogo.

Before the simulation begins, the user can initialize the variables to the necessary values. As can be seen in figure 6a, 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 6b. Each possible team composition is numbered and the legend is shown in the interface (see figure 6c). 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.

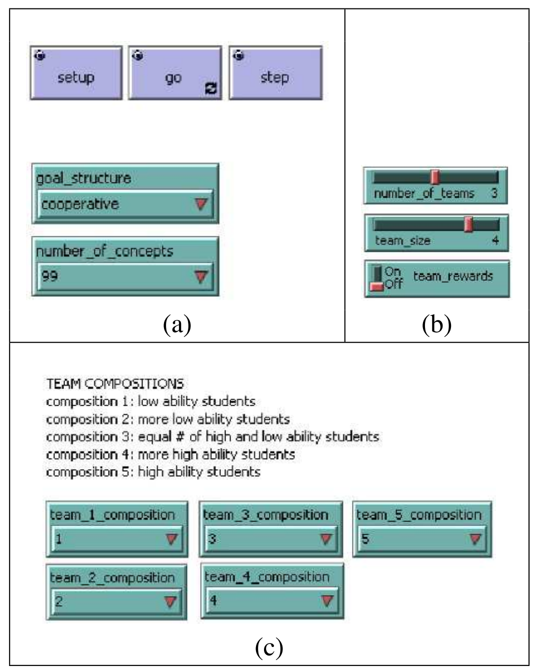

Fig. 6. The user's controls in the simulator.

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. Figure 7 shows all the different heterogeneous team compositions for each possible team size implemented in the simulation.

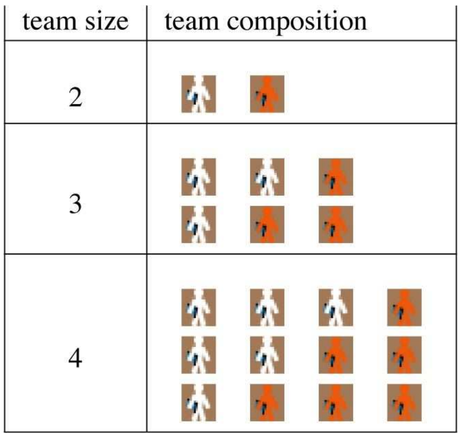

Fig. 7. Team sizes and their possible heterogeneous team compositions

The lower portion of the screen is subdivided vertically, into three sections, each separated by dark vertical bars. This is illustrated in figure 5. 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 rightmost 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. Experiments and results

A series of experiments was conducted, designed to monitor the development (i.e., change in zone) of individuals within each of the three goal structures (for complete details, see [24]). The behavior of the simulated learners in the individual goal structure was used as a reference for learner behavior in the cooperative and competitive goal structures. For the cooperative goal structure, the following parameters were subject to experimentation: group size (number of learners in each group), group composition (homogeneous and heterogeneous with different mixes of high and low ability students), and the influence of team rewards on the learning behavior of high and low ability learners. The experiments involved 10 runs of 99 concepts each, for each goal structure.

Figures 8-10 contain the change in zone for both high and low ability learners in all group compositions, averaged over all runs. Mean change in zone and standard deviation are shown. Different group compositions are illustrated by combinations of (H) and low (L) ability learners. Groups sizes (2, 3 and 4) are represented implicitly in the number of learners denoted in each group composition. Values within the three figures can be compared, but note that it is not meaningful to contrast the change in zone between high and low ability learners; by definition, high ability learners will progress more quickly due to the different implementation of their improvement.

4.1 Goal structures

In this subsection, the results for the three goal structures simulated in the group learning simulation are compared.

Individual and Competitive goal structures produce similar results. When the development of individual learners in a competitive goal structure is compared with the development in an individual goal structure, in figure 8, the values lie too close together to point out significant differences between the behaviors. The only difference that can be pointed out is that the standard deviations of the learners in a competitive goal structure are smaller than in an individual goel structure, which might indicate that competitiveness creates more coherence among the learners. But generally, for both high and low ability simulated learners, it can be said that individual and competitive goal structures give rise to similar learning behavior.

Cooperative goal structures benefit high ability learners. As can be seen from figure 9, all values of the development of high ability learners are higher than the value for individual learning. It can therefore be stated that a cooperative learning environment tends to be beneficial for the development of high ability learners.

Cooperative goal structures only benefit low ability learners some of the time. On the other hand, when comparing the results of low ability learners in an individual versus a cooperative environment, it can be seen that learners in some group compositions do not seem to benefit from working in groups. Some of the team compositions result in the learners performing worse than in an individual goal structure. It cannot be said that the cooperative environment would therefore not be beneficial for low ability learners; it does however become clear that other factors might influence the success or failure of the cooperative goal structure for low ability learners.

The difference in learning behavior of low and high ability learners in a cooperative environment can be explained using the analogy of a race in team sports. In such a race, if the team is rewarded for everyone crossing the finish line, it makes sense for the faster runners to hang back and help the slower runners, to make sure the slower runners make it across the finish line. Instead of focusing on individual gain, the team members now have to help each other and, in order to do that, concentrate on the approach that would be best for the others to make them win. This difference between the individual or cooperative context has been explained in section 3.1 as a development of *understanding* instead of the increase in the *individual improvement score*. In the analogy used above, this translates into the fact that, for the faster runners, hanging back to help slower team members will not increase the faster runner's individual speed (rate of improvement) but it will deepen their understanding.

	individual goal structure	competitive goal structure
H	**0.2670** (0.0253)	0.2620 (0.0273)
L	**0.1517** (0.0233)	0.1542 (0.0185)

Fig. 8. Experimental results: **high** and **low** ability learners, for both **individual** and **competitive** goal structures.

4.2 Group composition and size

This section contains a comparison between the results for group composition within the cooperative goal structure.

High ability learners benefit most from working in small groups of homogeneous composition, while low ability learners benefit most from heterogeneous groups. The reason for the latter result could be that the low ability learners benefit from cooperation with high ability learners, since the high ability learners can help them progress. One result that illustrates this very well are the results for learners of both abilities in a group

	without team rewards	with team rewards
HH	0.3738 (0.1509)	0.3201 (0.0576)
HHH	**0.4162** (0.2085)	0.2955 (0.0958)
HHHH	0.2972 (0.0360)	0.3627 (0.0891)
HL	0.3062 (0.1011)	**0.4959** (0.3448)
HHL	0.2916 (0.0744)	0.3621 (0.2213)
HLL	0.3514 (0.1015)	0.3210 (0.1084)
HHHL	0.3082 (0.0535)	0.2899 (0.0632)
HHLL	0.3295 (0.0703)	0.2610 (0.0763)
HLLL	0.2807 (0.1153)	0.2792 (0.1790)

Fig. 9. Experimental results: **high** ability learners, **cooperative** goal structure.

	without team rewards	with team rewards
LL	0.1631 (0.0713)	0.1182 (0.0646)
LLL	0.1611 (0.0526)	0.1698 (0.0509)
LLLL	0.1714 (0.0455)	0.1619 (0.0429)
HL	**0.2321** (0.1279)	0.1185 (0.0986)
HHL	0.1679 (0.1060)	0.1806 (0.0933)
HLL	0.1546 (0.0695)	0.1568 (0.0653)
HHHL	0.1819 (0.1236)	**0.2247** (0.0956)
HHLL	0.1441 (0.0840)	0.1478 (0.0708)
HLLL	0.1464 (0.0574)	0.1921 (0.0660)

Fig. 10. Experimental results: **low** ability learners, **cooperative** goal structure.

with a composition with a small number of low ability learners and a large number of high ability learners (like **HHHL**). This composition is most fruitful for low ability learners; working together with only high ability learners will provide the low ability learner with a lot of help. The high ability learners benefit in turn from cooperating with low ability learners because they gain understanding. This trade-off can be seen in figure 9, for the same team composition: where the low ability learner scored relatively very well, the high ability learner does not develop much more than in an individual goal structure. The results do show, however, that both high ability learners and low ability learners can thrive in heterogeneous groups, whereas homogeneous group compositions only pay off for high ability learners. A possible explanation for this could be that high ability learners only benefit from helping low ability learners in certain circumstances; low ability learners, on the other hand, are always helped by high ability learners.

4.3 Team rewards

In the next experiment, the influence of team rewards in the cooperative goal structure was examined.

Team rewards do not always have the intended effect of improved development; very often, both high and low ability learners perform worse than without team rewards. As can be seen from figure 11a, team rewards work especially well for high ability learners in large homogeneous groups and small

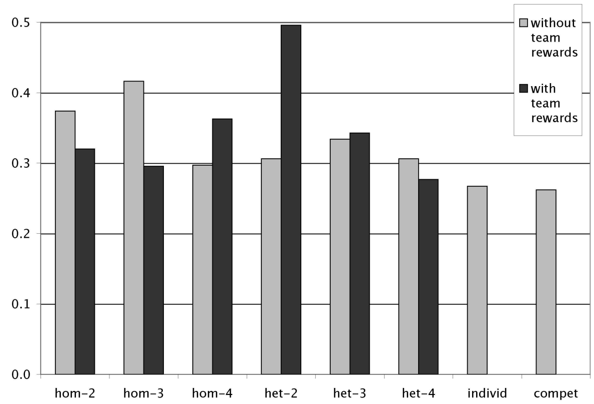

(a) high ability learners

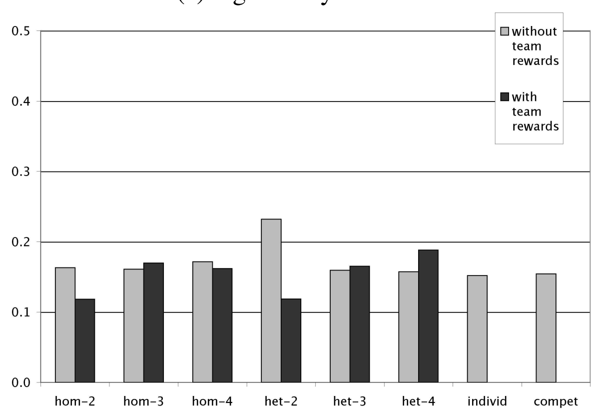

(b) low ability learners

Fig. 11. Experimental results: team rewards. Results are averaged over each group size (2, 3 or 4), for each type of group (homogeneous vs heterogeneous).

heterogeneous groups. This can be explained by the increased chances of high ability learners to rank highly in a learning environment where group performance is compared. Low ability learners, especially in small groups, cannot “outrank” the groups with more high ability learners and will therefore lose motivation. An interesting result shown in figure 11b is therefore the development of low ability learners in a homogeneous group of three; they seem to benefit from team rewards, while many other groups would seem to be better cognitively. The influence of team rewards on the simulated learners is closely related to group size.

By introducing team rewards, the pedagogical literature predicts that group members are more responsible for their group members’ progress. Team rewards can therefore be an important motivator for group members, and can be compared to “team spirit” amongst members of a sports team [23]. Based on this motivational aspect, the prediction was that team rewards would have a positive effect on the learning behavior of the simulated learners in the cooperative goal structure. In the simulation, team rewards have an influence on the motivation and emotion of the group members: when a cooperative group improves a lot compared to the other groups, the motivation of all its members will increase; the motivation will decrease if a group is ranked last. As a result of the increased motivation, the emotion will also increase: the learners become “happier”.

4.4 Learner interactions

In addition, the level of interaction that occurred among the learners was observed. In a cooperative goal structure, this refers to members of the same team helping each other. In a competitive goal structure, this refers to individuals competing against each other. It is assumed that, in the individual goal structure, there is no interaction between learners. A representative sample of one run within each goal structure is shown below.

Figure 12 illustrates results of running the simulator for different combinations of settings selected by the user, for individual, competitive and cooperative goal structures. In figure 12a, it can be seen 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.

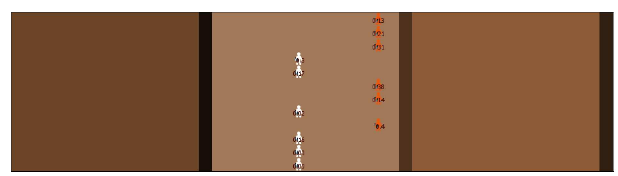

(a) individual goal structure

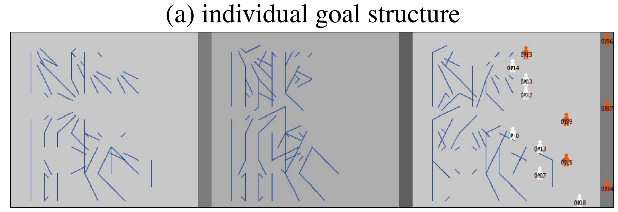

(b) competitive goal structure

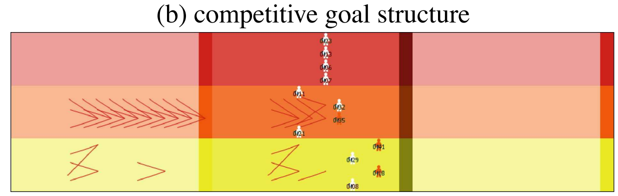

(c) cooperative goal structure

Fig. 12. Simulation experiments

In figure 12b, 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 behavior compared to the behavior 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, it can be seen that some low ability learners progress faster than some high ability learners. This shows that cognition is not the only factor influencing progress: the factors of motivation and emotion play an equally large role in the learning process of all learners.

In figure 12c, 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 improvement score in the evaluation phase. There are only three teams depicted in the figure, but there is a clear difference between the behavior 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, a heterogeneous team can be seen 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 to progress. 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.

5. Discussion

The background for and design of a *group learning model* and simulation system have been presented, in which theoretical human learners are modeled as artificial agents whose behaviors are influenced by a wide range of individual and environmental parameters. Using this simulator, three different goal structures were investigated in groups of simulated learners, characterized by features such as size, homogeneity and reward structures. A number of the parameters defined in the simulation have significant effects on learning outcomes, corresponding to trends observed in empirical studies of human learners described in the pedagogical literature. Even though computational modeling will always be an abstraction of the behavior of human subjects, agent-based simulation can be a powerful tool for examining aspects that are difficult to study *in situ* and can provide better understanding of individual and environmental characteristics that influence the progress of human learners. The approach described in this article is informed by work in the areas of pedagogical and peer learning agents, but an abstract approach to knowledge is taken, since the long term goal of this research 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.

The experimental results show differences in learning, measured by a model of each student's zone of proximal development. Summarizing the results presented here, it can be said that group composition, team rewards and team size have clear influences on the development of simulated learners in a cooperative environment. Different variable settings may help to overcome the apparent negative influences of this goal structure for low ability learners. This can be compared to a real-life situation, in which a teacher implements a goal structure in such a way that it enables her students to develop optimally.

The results also show that there appears to be no single optimal group size for either high or low ability learners; however group size is a very powerful factor in combination with other variables, like group composition or the presence of team rewards. The hypothesis that a larger group would give rise to more development is proven to hold only for homogeneous groups with team rewards, or for high ability learners in heterogeneous groups without team rewards. One observation that can be made from watching the visualization of the learners in the simulation is that team rewards have a positive effect on group *coherence*, although this was not measured formally. The learners seem to progress more "together" in a situation with team rewards (than without). This is related to the helping principle, which enables a high ability learner to gain understanding by helping a low ability learner.

Webb and Palinscar [30] 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. This range of 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, a deeper understanding of the characteristics of learning environments and their effects on students may be gained.

Acknowledgements

This work was partially supported by grants from the U.S. National Science Foundation (NSF #ITR-05-52294) and the City University of New York (PSC-CUNY #68525-00-37).

References

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

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

[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), Vol. 32, No. 2, 1999, pp. 307–318.

[4] S C Daum, Using simulated students to improve residents' teaching. Acad Med., Vol. 72, No. 5, 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, Vol. 4, No. 2, 1970, pp. 13–17.

[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. International Conference on Computers in Education 1995.

[9] W L Johnson, Pedagogical Agent Research at CARTE. AI Magazine, Vol. 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, Vol. 11, 2000.

[11] Y Kim, Agent learning companions: Learners' expectations of the desirable characteristics. 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, Vol. 45, No. 4, 1991, pp. 352–367.

[14] S Marsella, W L Johnson and K LaBore, Interactive Pedagogical Drama. 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, Vol. 8, 1978, pp. 181–192.

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

[17] K Sehaba and P Estraillier, A multi-agent system for rehabilitation of children with autism. 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. 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, Software review: Netlogo, a multiagent simulation environment. Journal of Artificial Life, Vol. 13, No. 3, 2007, pp. 303–311.

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

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

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

[24] 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.

[25] N A Taatgen, Explicit Learning in ACT-R, Pabst, Lengerich, 1999, ch. 3.1, pp. 233–252.

[26] K VanLehn, S Ohlsson and R Nason, Applications of simulated students: An exploration. Journal of Artificial Intelligence in Education, Vol. 5, No. 2, 1996, pp. 135–175.

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

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

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

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

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

Author Bios

Maartje Spoelstra studied Artificial Intelligence at the Vrije Universiteit in Amsterdam, The Netherlands. During her studies, she did a six month internship at the Brooklyn College of the City University of New York (USA), under supervision of Elizabeth Sklar. Since she graduated for her MSc degree in 2006, she has been working as a defence scientist at TNO, the Netherlands Organisation for Applied Scientific Research.

Elizabeth Sklar is an Associate Professor in the Computer and Information Science department at Brooklyn College and at the Graduate Center, both of the City University of New York (CUNY). She received her PhD from Brandeis University in 2000. Dr Sklar's research interests include interactive learning systems, multiagent simulation of learning environments, and educational robotics. She has received several grants from the US National Science Foundation (NSF) and Dept of Education to support this work. Dr Sklar is on the Board of Trustees for the RoboCup Federation and has served on the program committees for AAMAS, IAT, AAAI, RoboCup, and on many other international workshop and conference committees.