Agent-based Simulation of 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
Department of Command and Control
Oude Waalsdorperweg 63
2509 JG The Hague, NL
maartje.spoelstra@tno.nl
[3] 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 construct groups of *simulated learners* that model the behaviour of humans acting in various learning environments, with the aims of studying *group learning* and focusing on the effects of different *goal structures* on individuals and groups of learners. Three sets of research objectives are investigated: (1) imitating the behaviour of human learners with a multiagent simulation by modeling characteristics outlined in pedagogical literature; (2) comparing the outcomes of simulated learners operating with different goal structures; and (3) exploring factors that influence the behaviours of simulated learners acting in groups, such as group size and composition, as well as the inclusion of team rewards. We ran a series of experiments as part of this investigation, which are outlined herein.

1 Introduction

Multiagent simulations based on computational representations of human actors and characteristics of social environments can provide useful approximations of large-scale population studies or fine-grained behavioural studies. Although necessarily abstracted to varying degrees, these types of simulations can be useful either as a pre-cursor to experiments involving real humans or as a means of analyzing previously collected data sets [17]. The work described here examines *group learning* and focuses on the effects of different *goal structures* on individuals and groups of *simulated learners*. Three research objectives are investigated: first, imitating behaviours of human learners in a multiagent simulation by modeling characteristics outlined in pedagogical literature; second, comparing the behaviours of simulated learners responding to different goal structures; and third, exploring factors that influence the behaviours of simulated learners acting in groups, such as group size and composition and the inclusion of team rewards.

Earlier related work describes “SimEd”, an environment that emulates interactions between simple artificial learners and abstract knowledge domains [12]. 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 [9]. Here we expand upon this line of work in two main ways. First, we model peer-to-peer interactions—whereas the previous work only models the results of student-teacher interactions. Second, we base the details of the present simulation on the large body of existing research on "group learning" that has been conducted by developmental psychologists, education researchers and cognitive scientists. Thus our models of human learners are grounded in empirical and controlled experimental studies well-documented in the literature—whereas the previous work abstracted many of the details of the human "agents" and was based on canonical views of classroom activity. 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.* [18] 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 [11]. Uses for systems that simulate students can be grouped into three categories [18]: *teacher training* [4, 3], *peer tutoring* (where the peer is a simulated student) [20], and *instructional design* [19]. Peer tutoring is the most closely related to the work described here.

In the work presented in this paper, we examine aspects of group learning, comparing a range of different reward mechanisms for individuals and groups, as well as various heterogeneous (vs homogeneous) group compositions. First, we provide an overview of relevant pedagogical literature describing the characteristics of individuals and goal structures in group learning situations. In section 3, we describe our *group learning model* and the design of a simulator which we constructed for experimenting with the model. Section 4 presents some results, and we close with a brief discussion.

2 Background

Our group learning model is based on several important pedagogical theories of human learning and skill acquisition as well as applications of these theories to implementations of instructional and learning processes in a classroom. Fitts [6] describes a theory, involving three phases for physical skill learning in adults, which has influenced many others studying skill acquisition. Fitts' theory claims that when learning a skill, human development goes through an "early" phase, an "intermediate" phase and a "late" phase. In the early phase, the emphasis of the learning task 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

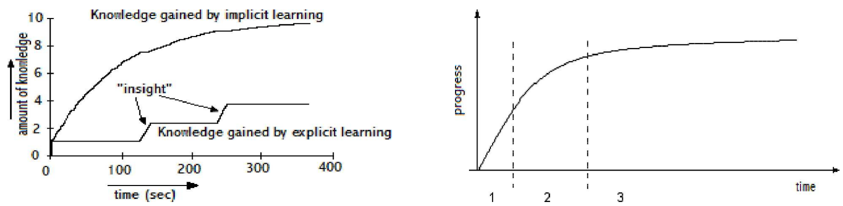

(a) implicit and explicit learning models, taken from [16], based on [1]

(b) stages of cognitive development combined with the stage-theory of [6] and [1], inspired by [16]

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.

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 [16] 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 in "explicit learning", goals and intentions determine what is learned. In an educational system, we can say that explicit learning gives rise to the cognitive outcomes of goal structures and implicit learning gives rise to the affective outcomes.

In many pedagogical studies, researchers distinguish between several levels of *ability* because some learners progress more quickly than others: some studies mention three levels ("high", "medium" and "low") [10], but the most common are two levels of ability ("high" and "low") [2]. In our study we chose to focus on two levels of ability.

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" [21]. In order for a learner to process new information optimally, the level of the information should be such that it can be grasped by 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" [9]. 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 for a large part on the learning environment for a learner and on interaction with others while learning.

The design and implementation of one or more *goal structures* is a part of the educational process within the classroom and can focus 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 personal goals 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, 13, 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.

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 [13]. 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.

3 Group learning model

Our *group learning model* is designed based on the pedagogical theories highlighted above. In this section, we first outline the parameters that define *agents* acting in the simulator; each agent represents an individual human learner. The model of cognitive development—i.e., progress made by individual learners—which underlies our simulation is illustrated in figure 1b. In the initial stage of learning (labelled "1" in the figure), 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 slope of the curve is quite steep. In the associative stage ("2" in the figure), instructions are formalized and made part of the learner's own skills; the slope of the curve 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 autonomous stage ("3"), the learner does not learn new things but constantly elaborates the present knowledge.

The second part of this section discusses the *learning environment* (or *instructional model*) in which the agents interact. The instructional model that we simulate is based on the notion that first a student is exposed to some new knowledge (we use the term "concept" to indicate a unit of knowledge)—this could be like reading about it in a textbook chapter or hearing a teacher give a lecture on the topic—then the student has a chance to practice working with the knowledge, such as, for example, doing a homework assignment or lab work, writing a computer program, answering question at the end of a textbook chapter; and finally the student's new knowledge is assessed. We label this evaluation a "quiz", but it could also mean the teacher marking a homework or lab assignment—the term "quiz" here refers to the stage at which the teacher (or instructional module in an intelligent tutoring system) gains feedback on whether the student is acquiring the new knowledge or not.

The final part of this section describes the simulator and how it demonstrates the group learning model.

3.1 Agents

Each agent is defined by a number of parameters, as detailed below.

– **ability**—indicates whether the agent has "high" or "low" aptitude. This value does not change during the simulation and can be thought of like IQ (intelligence quotient), i.e., a value that indicates a learner's innate aptitude and remains constant over their lifetime.

– **improvement**—reflects the general increase in knowledge throughout the learning of a new concept.

– **progress**—is the cumulative value of improvement (shown in figure 1b).

– **base_score**—represents the score of the quiz the student took before the first concept and after each concept.

– **improvement_score**—is the outcome of the quiz taken by the learner in the evaluation phase of each concept cycle (see below) and is the value used to increment the **base_score** after completing a concept. It is a combination of the improvement of the learner during the concept presentation and the gained **understanding** (see below).

– **understanding**—is gained by the learner when learning a concept of which the **difficulty** (explained in the next subsection) falls within her **zone** (see below). Another way in which a learner gains extra understanding is when explaining things to peers in the cooperative goal structure. The understanding gained by a learner depends on the current improvement of that learner and the help provided to others. The gain in a learner's understanding may be at the cost of that learner's improvement; therefore, it does not always pay off for a learner to help others. At the beginning of each concept, understanding is set to 0 again, indicating that the subsequent concepts are independent.

– **zone**—resembles the center of a frame, bounded by **zone** $\pm\ \epsilon$, and represents the "zone of proximal development". This is a cumulative variable, to which the **learning_rate** (see below) is added after each concept. Note that the size of the frame stays the same throughout the development of the learner; as the value of **zone** increases when the learner improves, the entire frame shifts accordingly.

– **learning_rate**—is calculated as the average change in **improvement_score**, per **tick** (one time unit in the simulation). This variable is used to indicate the overall development of the learner (added to **zone** after each concept), because students take different time spans to learn a concept, as can be seen in the simulation from the way they progress.

– **motivation**—attempts to capture in an abstract way whether a learner is motivated to do well or not, i.e., if a student has the ability to learn a concept, does she actually acquire it? The value of **motivation** changes as the simulation runs and depends on whether the **difficulty** of the current concept lies within the learner's **zone** and whether or not the learner passed the quiz at the end of the previous concept presentation. If the learner "fails", she becomes motivated to do better next time if she failed by a little (**motivation** increases), but demoralized (**motivation** decreases) if she failed by a lot. In the case of cooperative learning, motivation is also influenced by the motivation of the teammates. In case of competitive learning, a learner's motivation increases if both the learner and her opponent have a competitiveness factor above a certain threshold.

– **emotion**—attempts to capture in an abstract way whether a learner is paying attention to the lesson and able to absorb all the input given during the initial presentation phase, i.e., if a learner is unhappy or depressed, she may not listen to everything her teacher says. In our simulation, the value of **emotion** changes over time and depends on how well the student performs on the quiz after progressing through a concept. For cooperative learning, **emotion** also depends on teammates' **emotion** and the rank of the team after the learners all complete the quiz. In the pedagogical literature, researchers often remark on the fact that in a competitive setting, students tend to prefer that others do not get benefits if they themselves do not receive any [7, 8]. This tendency led us to implement an increase in emotion when learners compete; if learners are close together in **zone**, they form a threat to each other and competing gives them a means to try and get ahead of each other.

– **target**—is the individual target for each concept. It is a goal only for the learners in the individual the cooperative goal structures (as explained in [13]); in the competitive goal structure everyone strives for the same goal.

– **likeliness_to_help**—is the help others can give to a learner in a cooperative context and depends on how likely they are to help their peers. The help provided to other learners is calculated as the product of **likeliness_to_help** and **improvement** and represents the amount that is subtracted from the helper's **improvement**, as that learner "stays behind" to help a peer. However, the lost **improvement** is invested in **understanding**. An important factor to note is that the help provided by one learner and the help received by another learner are not necessarily the same. The amount of help given by a high ability learner to a peer depends on that learner's improvement and the **likeliness_to_help** of the learner. The receiver of the help is also responsible for the cooperation: the effort invested in the learner by the other is received according to the receiving learner's motivation; if the receiving learner is unmotivated, the help is not fully effective. Another factor influencing the learner's 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 **zone** of the learner helped, only a fraction of the help provided reaches the learner.

– **competitiveness**—is similar to **likeliness_to_help**, but it applies to interactions in a competitive goal structure. Very competitive learners might become motivated because of this competitiveness, in which case competitiveness has a positive influence on the learner. When a competitive and a non-competitive learner compete, this might have negative influences on the motivation of the non-competitive learner.

3.2 Environment

When implementing each goal structure in a human classroom, the setting of instructional targets for individuals or groups of learners and the design of the evaluation phase are all tasks that belong to the teacher. In our simulation, the role of the teacher is part of the environment: concept difficulties are set randomly, individual and group targets are set according to the previous outcomes of the learners, and evaluation is done according to individuals' targets and progress. The teacher is represented implicitly in the simulation as an agent bearing a unilateral dependence relation with her students; learners' behaviours depend on their teacher, but the teacher's behaviour does not change in response to the learners. This is an example of the *lecture model* of teacher behaviour described in [12].

The performative structure [5] (i.e., a description of the sequence(s) of activity in the model) of the classroom environment is divided into three phases: **initialization**, **learning** and **evaluation**. This three-phase "concept cycle" executes for each concept in the simulated curriculum. Following the STAD learning method, in the **initialization** phase, teachers present an introduction to a new concept and defines individual targets. In the simulation, this translates into setting values for each agent such as **base_score** and **zone**.

Then, agents enter the **learning** phase where they progress through the concept; this is the phase in which students are given the chance to acquire new knowledge—whether they actually progress or not (and how much) depends on their **motivation**, **emotion**, **ability**, **zone** and the **difficulty** of the concept. In a competitive goal structure, this is the phase in which learners interact by competing. In a cooperative goal structure, the learners interact by helping each other to progress and by influencing each other's motivation and emotion. The learning phase is the longest phase in the concept cycle.

In the **evaluation** phase, students' **progress** is measured and combined with **understanding**, which together add up to a learner's **improvement_score**. This value is compared to the learner's target for the concept. The **improvement_score** forms the learner's new **base_score** in the initialization phase for the next concept. If the improvement score is equal to or higher than the target, the learner has passed the quiz. This results in an increase in **motivation** and **emotion** (the learner becomes "happy"), which will have a positive influence on the performance of the learner during the next learning phase. With the cooperative goal structure, the evaluation phase is used for calculating the ranks of the participating groups: in a situation in which five groups compete, the two groups that contain the learners who scored best on their individual quizzes are rewarded, having a positive influence on their motivation and emotion, and the lower scoring groups become disappointed, resulting in a negative effect on their motivation and emotion.

Two variables that drive the learning process are the **difficulty** of each concept to learn and the number of **ticks** spent on each concept. The notion of concept difficulty is based on [12], in which a set of concepts is represented as a graph of nodes, and each concept has a real-valued **difficulty** between 0 (easy) and 1 (hard). In that work, a **concept** is a small bit of knowledge, such as the spelling or meaning of a word or an arithmetic equation. Here, the notion of a concept could be broader and represent a larger amount of information, such as "addition" instead of simply, for example: $5 + 2$.

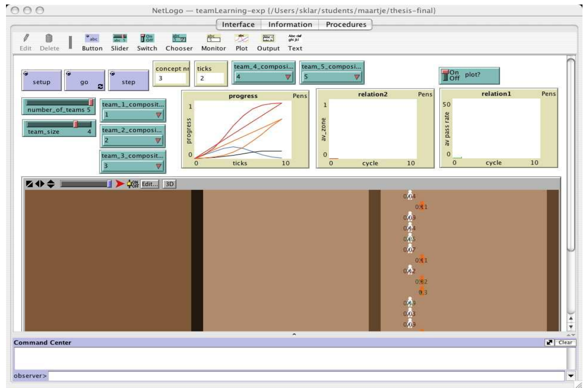

Fig. 2. Screenshot of the simulator. In the bottom portion of the screen, rows of simulated learners are shown; they begin on the left edge of the window and move to the right as they progress through each concept (there are three concepts per screen width, displayed in a horizontal scrolling window). Each dark vertical stripe indicates the evaluation phase of the previous concept and the initialization phase of the next concept. Between the dark stripes is a learning phase, shaded according to the **difficulty** of that concept: the darker the color of the concept, the harder it is. All agents start the learning phase at the same time. Interactions between agents in the cooperative and competitive goal structures are indicated by drawing lines between collaborating agents or competing agents. The interface, at the top of the screen, can be run interactively and allows the user to modify parameter settings before each run: the number of concepts to be learned and the goal structure. In the case of a cooperative goal structure, the user can choose the number of groups participating, the group size and composition for each group, and whether or not team rewards are present.

We use the term "concept" to refer to the amount of curricular material that a teacher chooses to cover in one lesson. In the simulation, learners have a certain amount of time to master each concept, measured in **tick**s; 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.

3.3 Simulator

The group learning model was simulated using NetLogo[4] [22], depicted in figure 2 [14, 15]. Programming in NetLogo is inherently agent-based, and its robust and easy-to-use graphical user interface makes it an ideal environment for prototyping and running relatively small scale experiments.

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,

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

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:

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

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**:

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

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**.

4 Experiments and Results

We conducted a series of experiments designed to monitor the development (i.e., change in **zone**) of individuals within each of the three goal structures. The behaviour of the simulated learners in the individual goal structure was used as a reference for learner behaviour in the cooperative and competitive goal structures. For the cooperative goal structure, we experimented with different settings of the following parameters: 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 behaviour of high and low ability learners. The experiments involved 10 runs of 99 concepts each, for each goal structure. Table 1 contains the change in **zone** for both high and low ability learners in all group compositions, averaged over all runs. Values within the sections of the table 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.

We compare the results for the three goal structures simulated.

Individual and Competitive goal structures produce similar results. When we compare the development of individual learners in a competitive goal structure with the development in an individual goal structure, in table 1a, 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 table 1b, all values of the development of high ability learners are higher than the value for individual learning. We can therefore say that a cooperative learning environment tends to be beneficial for the development for 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 we can see 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.

4.2 Group composition and size.

We compare the results for group composition within 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 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 table 1b, 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.

(a) **high** and **low** ability learners, for both **individual** and **competitive** goal structures:

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

(b) **high** ability learners, **cooperative** goal structure:

	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)

(c) **low** 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)

Table 1. Experimental results: goal structures, group composition and size. 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.

4.3 Team rewards.

We conducted an experiment examining the influence of team rewards in the cooperative goal structure.

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 3(a), team rewards work especially well for high ability learners in large homogeneous groups and small 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 3(b) 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 [13]. Based on this motivational aspect, our

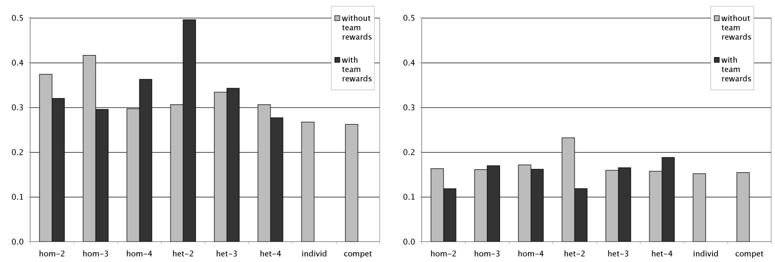

(a) high ability learners (b) low ability learners

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

prediction was that team rewards would have a positive effect on the learning behaviour 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".

5 Discussion and Summary

Our experimental results show differences in learning, measured by a model of each student's zone of proximal development. Additional experiments and details of this work can be found in [14]. Summarizing the results presented here, we can say 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.

We have presented the background for and design of a *group learning model* and simulation system in which theoretical human learners are modeled as artificial agents whose behaviours are influenced by a wide range of individual and environmental parameters. Using this simulator, we have investigated three different goal structures 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 behaviour 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.

Acknowledgments

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

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. M. Esteva, J. A. Rodriguez, C. Sierra, P. Garcia, and J. L. Arcos. On the formal specifications of electronic institutions. *Agent-mediated Electronic commerce (The European AgentLink Perspective), LNAI 1991*, pages 126–147, 2001.
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. A. Kohn. *No Contest: The Case Against Competition*. Houghton Mifflin, Boston, 1986.
9. R. S. Lazarus. Cognition and Motivation in Emotion. *American Psychologist*, 45(4):352–367, 1991.
10. G. Mugny and W. Doise. Socio-cognitive conflict and structure of individual and collective performances. *European Journal of Psychology*, 8:181–192, 1978.
11. S. Ohlsson. Learning from performance errors. *Psychological Review*, 103:241–262, 1996.
12. E. Sklar and M. Davies. Multiagent simulation of learning environments. In *Fourth Int'l Conf on Autonomous Agents and Multi Agent Systems (AAMAS-2005)*, 2005.
13. R. Slavin. *Cooperative Learning: Theory, Research and Practice*. Alley & Bacon, Boston, 1990.
14. 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.

15. M. Spoelstra and E. Sklar. Using simulation to model and understand group learning. In *Agent-Based Systems for Human Learning and Entertainment (ABSHLE) Workshop at Autonomous Agents and MultiAgent Systems (AAMAS-2007)*, 2007.
16. N. A. Taatgen. *Explicit Learning in ACT-R*, chapter 3.1, pages 233–252. Pabst, Lengerich, 1999.
17. Y. Tang, S. Parsons, and E. Sklar. Modeling human education data: From equation-based modeling to agent-based modeling. In *Multi-agent Based Simulation (MABS) Workshop at Autonomous Agents and MultiAgent Systems (AAMAS-2006)*, 2006.
18. K. VanLehn, S. Ohlsson, and R. Nason. Applications of simulated students: An exploration. *Journal of Artificial Intelligence in Education*, 5(2):135–175, 1996.
19. 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 Int'l Conf on Advanced Learning Technologies (ICALT'03)*, 2003.
20. A. Vizcaino and B. du Boulay. Using a simulated student to repair difficulties in collaborative learning. In *Int'l Conf on Computers in Education*, 2002.
21. L. S. Vygotsky. *Mind in society: The development of higher psychological processes*. Harvard University Press, Cambridge, MA, 1978.
22. U. Wilensky. Modeling nature's emergent patterns with multi-agent languages. In *Proceedings of EuroLogo 2002*, Linz, Austria, 2002.