

# Ultra-Personalization and Decentralization: The Potential of Multi-Agent Systems in Personal and Informal Learning

Ali M. Aseere<sup>1</sup>, David E. Millard, Enrico H. Gerding

<sup>1</sup> School of Electronics and Computer Science,  
University of Southampton, Southampton, UK  
{ama07r,dem,eg}@ecs.soton.ac.uk

**Abstract.** Agents are autonomous software components that work with one another in a decentralized fashion to achieve some end. Agent systems have been used in Technology Enhanced Learning (TEL) before, but these applications seldom take advantage of the fact that each agent may have its own goals and strategies, which makes agent systems an attractive way of providing personalized learning. In particular, since agents can solve problems in a decentralized way, this makes them an attractive way of supporting informal learning. In this paper we use scenarios to examine how common problem solving techniques from the agents world (voting, coalition formation and auction systems) map to significant challenges for personalized and informal learning in the TEL world. Through an agent simulation we then show how an agent system might perform in one of those scenarios and explore how different agent strategies might influence the outcome. Based on this work we argue that agent systems provide a way of providing ultra-personalization of the learning process in a decentralized way and highlight equitability and scrutability as two key challenges for future investigation.

**Keywords:** Agent Systems, Personalized Learning, Informal Learning

## 1 Introduction

Agents are special software components that work together in an agent framework to achieve some end. Their main features include autonomy, reactivity, proactivity and social ability [1]. Multi-agent systems, where several such agents interact, are being used in a wide variety of applications, ranging from comparatively small systems for personal assistance, to open, complex, systems for industrial applications [2]. In Technology Enhanced Learning (TEL) they can provide new models of learning and applications, such as personal assistants, user guides and alternative help systems, which are helpful for both students and teachers [3]. It has also been argued that using multi-agent systems to design educational systems lead to more versatile, faster and lower cost systems [4].

We believe that the major potential in multi-agent systems has yet to be fully explored, and that it relates to the ability of agent systems to support personalized and informal learning. In the e-learning domain we are increasingly seeing a move from a

world of VLEs (Virtual Learning Environments) into a space where students are taking more control of their learning in the form of PLEs (Personal Learning Environments), either as monolithic applications to help students manage their resources and time, or as a collection of online tools (such as Google calendar to manage time, 43 Things to manage goals, etc). In this personalized learning context agent technology becomes even more appropriate because agents are good at representing the requirements of users, and negotiating a more personalized experience. There is also a lot of potential to support informal learning, because in a decentralized agent system there is no need for a central authority (such as a tutor or academic institution) to orchestrate collaborations and learning activities.

In this paper we explore the potential of multi-agent systems for personalized and informal learning. We first present a number of scenarios that show how common problem solving techniques in the agents world (voting systems, coalition formation and auction systems) could map to problems in TEL, and explore how agent technologies could lead to ultra-personalization and decentralization, enabling new scenarios that are not possible with today's technology.

We then present a multi-agent simulation of the first of these scenarios (students making module choices) in order to demonstrate how a multi-agent system can solve problems in a decentralized way. In a comparative experiment we also explore how different student strategies (the algorithms that individual agents use to negotiate with each other) affect the outcome.

Our research demonstrates how agent systems could be applied to TEL to support personalized and informal learning, but it also highlights a number of key issues that may be of concern to educators.

The rest of this paper is organized as follows. Section 2 describes related work focusing on personalized and informal learning and the existing use of agent technology in TEL. Section 3 presents three scenarios that show how solutions from the agent domain might be used to solve problems from the TEL domain in a novel way. Section 4 describes an agent simulation of the first of these scenarios, and an experiment to compare the performance of three different student strategies. Section 5 analyses our findings and argues that while agent systems enable ultra-personalization and decentralization they also present new problems of equitability and scrutability. Section 6 concludes the paper and discusses our future work plans.

## **2 Background**

### **2.1 Personalized and Informal Learning**

Personalization in an educational setting is about working in partnership with the learner to tailor their learning experience and pathways according to their needs and personal objectives. Personalization is perceived as the task of providing every learner with appropriate learning opportunities to meet individual learning needs supported by relevant resources that promote choice and advance learner autonomy [5]. The concept of personalized learning emerged as a result of several developments. Partly, it is a reflection of living and working in a modern society, the developments of new

technologies, and in particular how they can enable learners to break down institutional barriers and become a part of a global society.

There is also a growing recognition that current educational provision may be too narrow and restrictive and is not meeting the individuals' learners or society needs as a whole [6]. Current learners see technology as core to their learning environments in particular computer and mobile devices. They use the Internet usually to support their learning, to find information and to discuss work with other students and teachers. They are comfortable working with multiple representations, are digitally literate, and happy to turn to Internet-based tools to help achieve their learning [5].

Sampson [7] has suggested that e-learning benefits from the advanced information processing and the Internet technologies to provide the following features which could be considered as lineaments of personalized learning:

- *Personalization*, where learning material are customized to individual learners, based on an analysis of the learners' objectives, status and learning preferences.
- *Interactivity*, where learners can experience active and situated learning through simulations of real-world events and on-line collaboration.
- *Media-rich content*, i.e. educational materials presented in different forms and styles.
- *Just-in-time delivery*, i.e. support systems that can facilitate training delivery at the exact time and place that it is needed to complete a certain task.
- *User-centric environments*, where learners take responsibility for their own learning.

Although personalized learning is increasingly recognized as important in formal settings, it is key for informal learning. There are many definitions of formal and informal learning, however the key distinction is that formal learning is typically described as learning that is managed in some manner by an authority (for example, a School or University), while informal learning is less managed, or may be managed by the learner themselves [8, 9]. A survey by Cross showed that 70 percent of adult learning is self-directed learning [10] and informal learning is increasingly recognized as a key domain for TEL,

Based on these definitions we believe that agent technology is a good approach to support personalized and informal learning. This is because of the characteristics of intelligent agents, which are autonomy, social ability, adaptability, and reaction. Agents can easily represent learners, adapting content and acting autonomously on their behalf. In addition, they can interact with multiple students and agents at the same time in order to facilitate collaborative and team learning [1] without the need for a formal centralized authority.

## 2.2 Agent Technologies

The term agent has been in existence in a number of technologies and been widely used, for example, in artificial intelligence, databases, operating systems and the marketplace [2]. Researchers in the agent technology field have proposed a variety of definitions of what comprises an agent but all agree that a key feature is autonomy.

One of the most common definitions is that proposed by Jennings and Wooldridge: “*An agent* is a computer system situated in some environment that is capable of autonomous action in this environment in order to meet its design objectives” [11]. *An intelligent agent* is a flexible agent that is pro-active, reactive, or social and able to learn to improve its own performance [12].

While agents work in the same *agent framework* (and often work together) they are inherently autonomous, and may have different and conflicting goals. This means that in their interactions they are trying to maximize their own benefit (this is more formally described as maximizing a utility function). There is therefore a need to establish the rules by which agents can converse and negotiate with one another; this is called the *agent protocol*. For example, in situations where a group of agents need to come to a common decision a voting protocol can enable them to reach agreement as a group while taking into account individual preferences [13].

Voting theory is an active area of research in multi-agent systems and one example of how agents can make decisions together in a decentralized way. It is part of the general area known as *social choice*, which is concerned with procedures for making collective decisions that maximize the *social welfare* (the sum of utility of individual agents), while at the same recognizing that agents are self-interested and act in a way that maximizes their own individual preferences. In the agent experiment described in Section 4 we use a novel voting procedure that combines two approaches: the *single transferable vote* STV<sup>1</sup> and *cumulative voting*<sup>2</sup>.

### 2.3 Agents in Technology Enhanced Learning

In TEL, multi-agent systems appear to be a promising approach to deal with the challenges of educational environments. A number of researchers have applied agent technology to e-learning, however they often use the agents as advanced components, focusing on the individual agent and/or its relationship with an individual student, rather than looking at a system of agents that works together to achieve some goal.

In this context, De Meo et al. [14] proposed the X-Learn system, an XML-based multi-agent system for adaptive e-learning based on user preferences and requirements. However, rather than the multi-agent aspect, they focus on the adaptation itself and how to exploit XML technology facilities for handling and exchanging information related to e-learning activities.

Shi et al. [15] designed an integrated multi-agent systems for computer science education that focuses on two introductory courses where the learning process is student-centered, self-paced and highly interactive. They use Java RMI, JavaSpace and JATLite to create a web-based system; in this case they use personal agents to manage student’s data and their interactions with course material.

---

<sup>1</sup>This is a multi-agent voting procedure when the alternative that is ranked lowest is removed. Votes that had this alternative ranked first will now have another alternative ranked first. This is repeated until one alternative remains.

<sup>2</sup> Here, each voter receives a number of points (usually the number of points is equal to the number of candidates), and they are free to choose how many points to allocate to each candidate. The candidates with the highest cumulative points are selected as winners

Furthermore, Yang et al. [16] proposed to apply an intelligent system to enhance navigation-training systems that consists of the client portion and server portion using JADE framework. Like most work in this area, this paper focuses on the intelligence of individual agents themselves, rather than communications between agents. One exception to these approaches is Soh et al. [17] who have shown a system called Intelligent Multi-agent infrastructure for Distributed Systems in Education to support student in real time classroom where a buddy group is formed dynamically to support the members to achieve common goals.

Although these papers apply agents to e-learning, only Soh et al. demonstrates any kind of self-organization, and none of them applies any fundamental agent theories, such as mechanism design or social choice theory, to guide their design choices. In contrast, our approach is to explore how agent systems can be used for decentralization as well as personalization. For example in our experiment we examine how voting mechanisms can be used in an e-learning scenario where a University agent represents all the modules available, and where student agents can vote in any way he or she prefers. Thus our work explores, for the first time, voting procedures in an e-learning setting, and the consequences of the resulting system behavior for learners in that scenario.

### 3 Motivational Scenarios

We believe that agents have the potential to transform Technology Enhanced Learning by enabling scenarios that are simply not feasible with today's technology. This is possible because of some of the key features of agent systems such as distributed control and agent autonomy. In this section we illustrate this potential through three different TEL scenarios that show how agent technologies could be used in e-learning to take full advantage of the agent's ability to communicate and negotiate. Each case is composed of a *description* of the scenario, an analysis of the *agent solutions* that make the scenarios possible, and more speculative *variations* of the scenario that would share the same features.

Through the scenarios we hope to show how certain types of problem in Technology Enhanced Learning fit with known agent solutions (Voting systems, Coalition Formation, and Auction systems). We also hope to show how agent systems enable a very high level of personalization, and to start a discussion about the implications for education in the future.

In Section 4 we will take the first of these scenarios and describe a prototype agent system that supports it, along with a suitable voting protocol, and analysis of how potential strategies perform. In Section 5 we reflect on the scenarios and experiment in order to identify key issues and challenges that will arise from the use of agent systems in education.

#### 3.1 Scenario One: Module Selection

**Description.** This scenario concerns a University that wants to support students who are interested in a wider variety of modules than it is possible for the University to

offer. The University must therefore somehow choose which subset of modules to run. This is a common scenario with Higher Education degree courses, where often students are offered a number of modules, and for economic reasons only the most popular modules are run. However, current solutions are centralized, requiring students to hand over their preferences to a central algorithm controlled by the University. In addition students are unable to respond to cancelled modules by changing their preferences. From a personalized learning point of view this is undesirable, as despite the tension between the goals of the institution and the students (the institution really wants to run as few modules as possible whereas each student wants to get the modules in which he or she has most interest) the student must hand over almost all control to an opaque process managed by the University.

**Agent Solutions.** In agent systems this scenario can be characterized as a **voting problem**. It occurs whenever agents are required to invest in or vote for a limited number of options within a greater number of more or less attractive possibilities. There are numerous potential solutions to voting problems where the outcome impacts all the agents (sometimes described as problems of **social choice**) but through transparent protocols they offer fairness, decentralization and independence (as they allow agents to choose their own voting strategies). This distribution of control fits well with personalized learning.

**Variations.** This scenario describes students making choices about modules within a single institution, however because agent solutions are decentralized an agent solution could also work in situations where students were choosing modules from multiple institutions (for example, as part of a personalized degree programme across Bologna compliant Universities). In this case, the factors taken into account in an individual agents voting strategy might also include issues such as institutional reputation, distance from home and student facilities.

### 3.2 Scenario Two: Group Formation

**Description.** In education it is often necessary for students to arrange themselves into groups for learning, for example to share equipment, to help with timetabling, or for pedagogical activities such as discussion. Students can group themselves, or be grouped by a teacher either randomly or based on some criteria. Group formation is important because although all students need to be allocated to a group, the mix of students might be important. For instance, it may be desirable to have a mix of abilities, so that no one group has an advantage over another in assessment.

Current solutions are normally centralized, meaning that students cannot have different criteria for group selection (for example, some students might wish to be in the most effective groups, while others would rather learn with existing friends) – similarly to Scenario One this one-size-fits-all approach is at odds with personalized learning and requirements to consider the learner experience.

An interesting aspect of this scenario is that sometimes the goals of the teachers are at odds with the goals of the students. The students may wish to be placed in groups with their friends or with students that will help them to achieve good marks, while the teacher may want to arrange students in groups that will help them to learn more

material or to learn it more quickly. This means that even non-centralized solutions may need to be mediated by a central authority.

**Agent Solutions.** In agent systems an appropriate metaphor for this scenario is **coalition formation** - a process by which agents form, join and switch groups until a stable set of coalitions is made. There are numerous potential protocols for this, for example by having an initial allocation, perhaps based on criteria set by the teacher, and then for the students to negotiate exchanging their places with students in other groups. The agent framework provides the conversational mechanism for this negotiation, but the agents need some self-organization. For example, each coalition might produce a virtual leader agent to negotiate with the leaders of the other groups. At the same time, each leader agent has to negotiate with the teacher agent because any changes made in group membership still have to conform to the constraints set by the teacher agent.

**Variations.** This scenario envisages group formation occurring under the supervision of a teacher or lecturer, and therefore implies a more formal educational context. However, distributed group formation enabled by agents could enable informal learners to also benefit from group work, by helping them form coalition with other (potentially remote learners) who share similar pedagogical goals. Such distributed agent-based group formation systems could be of great help to life-long learners, and could form the basis of informal group work and peer assessment without the need for a mediating teacher or institution.

### 3.3 Scenario Three: Personalized Learning

**Description.** Different students may have different personal preferences about the way they want to learn or to be assessed. These preferences may be because of preferred learning styles, but could also be for other practical reasons (such as time commitments in their personal lives, or different project requirements). An institution has difficulty catering for these preferences, due to the mixed cost of providing different activities (for example, lectures are cheaper than tutorials), resource restrictions (such as time commitments of staff, or access to specialized equipment or information sources) and their own guidelines and regulations about having a mixed set of assessment styles (for example, many Universities are cautious about having modules assessed totally by course work).

It is therefore rare for an institution to allow much flexibility at an individual level, although there are limited solutions that allow a cohort to make choices about how they will be taught or assessed, but these tend to be managed directly by teachers and are therefore of limited complexity (for example, it might be possible for the students to negotiate with their teacher about the methods of learning or assessment that will be used).

**Agent Solutions.** In this kind of scenario there are a number of limited resources (tutorial slots, lab equipment, seminar places, etc.) and many individuals competing for them. In agent systems this situation is characterized as an **auction**. The institution associates a cost with each type of activity and wants to minimize the total cost, or at least prevent it from rising above an agreed level. This cost need not be only

financial; it could include factors such as value to external assessors or complexity for staff to manage.

There are many different kinds of auction, and therefore different solutions to this problem. But as an example we can define a utility function for each agent that calculates their student's satisfaction with the activities that they have been allocated. Following an initial allocation, agents could then bargain (negotiate) with their institution, exchanging items according to their cost until their utility function is maximized within the constraints of the institution's cost level.

***Variations.*** Using an economic model allows a University to adjust the wealth (and therefore purchasing power) of certain students according to circumstances. For example, students with learning differences, such as dyslexia, could be given more credit, allowing them to tailor their learning experience within the same economic framework as other students. More controversially students might actually be allowed to purchase additional credit, in effect buying themselves more expensive tuition through the University fees system.

## **4 Agent Simulation and Experiment**

In Section 3 we described a number of TEL Scenarios and described how agent technologies could be used to solve them in a decentralized way. But what would be the performance of such a system, how would it behave and what would be the consequences for the learners and the institution? In this section we attempt to answer these questions by presenting an agent simulation of Scenario One (decentralized module selection using a voting protocol). To do this we must describe the context for our simulation (number of students, modules, etc.), the agent protocol (the rules under which the agents negotiate), and the strategies taken by individual agents. We can then use the system to demonstrate how a decentralized approach compares to an optimal centralized solution, and to compare the performance of different student strategies.

### **4.1 Context for the Simulation**

In our experiment we considered three different cases. These differ in terms of the number of students, the number of total modules, and the number of running modules. We choose these cases to reflect the kind of modules typical in UK computer science departments. We consider a large (undergraduate), medium (smaller undergraduate) and small (postgraduate) module. Table 1 shows the settings for these scenarios.

**Table 1.** Different setting to the scenarios.

Case	# modules (m)	#running modules (r)				#students (n)
1	51	10	20	30	40	100
2	33	9	18		27	60
3	15	4	8		12	20

In our system, each agent is autonomous, that is, it is in control of its own actions and responses. The system consists of two types of agents: student agents (SAs) and the university agent (UA). SAs and the UA use a voting procedure to interact with each other and to choose which modules to run. To this end, the UA manages the votes cast by the student agents and decides, based on the voting procedure and the votes received, which modules will be cancelled. Furthermore, after completing the entire process, it will provide the SAs with a final list of running modules.

#### 4.2 Protocol

In general, a protocol is the set of rules that controls the interactions between agents and determines the beginning and end conditions of a given conversation [18]. The protocol we used in our system works in several stages. In each stage, the students can cast their votes for the modules by allocating points to each module. The module that receives the lowest number of cumulative points is cancelled, and the points that were allocated to the cancelled module are refunded. In the next round, the students can use these points (and any points that they did not use in the previous rounds), to vote again. Furthermore, in each round, the students are informed about which module is cancelled and the total number of points that have been allocated to the remaining modules so far. Note that, once allocated, a student cannot retrieve its points, unless the module is cancelled. The advantage of this iterative approach is that votes are not wasted since points allocated to the cancelled module can be reused for the remaining modules. Furthermore, the student can use the information about the current “popularity” (i.e. the current cumulative points) of the modules to guide its voting behavior (we discuss the potential strategies of students in more detail in Section 4).

For example, if there are 40 modules available in total, but the university only has sufficient resources (e.g. staff and lecture rooms) to run 30 modules, then the voting will proceed for 10 iterations or rounds. At the end of each of these rounds, the module with the least number of cumulative points is cancelled.

#### 4.3 Strategies

Abstractly, a strategy determines the agent’s plan of action to achieve a particular goal. It specifies the way in which an agent behaves in a given environment [13]. In our scenario, the strategy determines the number of points to allocate to the modules in each voting round, given the preferences of the agent and the information received

by the UA about the voting process. In this experiment we compared three different strategies for the SAs in order to explore what would happen to students that adopted different strategies. These strategies were: proportional, equal share and intelligent.

**Proportional:** The proportional strategy was included as an example of a simple but sensible strategy. Consequently, it provides a good benchmark that we can use to compare the performance of more sophisticated strategies. The main idea behind a proportional strategy is that, in each round of voting, the student agent distributes its points proportionally to the student's preferences for each module. This strategy is simple in that it does not consider the information received by the UA about the current number of points allocated to the modules.

In more detail, the number of points allocated to module  $j$  is calculated as follows. Let  $RP$  denote the total number of points remaining (in the first round  $IP=RP$ ),  $m$  is the total number of available modules available, and the vector  $\vec{v} = \{v_1, v_2, \dots, v_m\}$  denotes the student preferences. Then, the total number of points to be allocated to module  $j$ ,  $b_j$  is:

$$b_j = \frac{RP}{\sum_{i=1}^m v_i} \cdot v_j \quad (1)$$

**Equal share:** The equal share strategy is included as an example of a very simple and ineffective strategy, and provides a good lower bound on the performance of the system. An equal share strategy is based on the principle that the SA gives all modules an equal number of votes, regardless of the student's preference. The following formula was used to calculate voting points each module:

$$b_j = \frac{RP}{m} \quad (2)$$

**Intelligent:** The intelligent strategy is included as an example of what can be achieved with a more sophisticated strategy that learns as the voting procedure progresses from one round to the next. The main idea behind this strategy is that, in each round, the agent tries to predict the probability that a module will be cancelled based on the number of points currently awarded to each module from previous rounds. Then, based on this probability, it can calculate its *expected satisfaction* for a given allocation of points, and it will allocate the points such that the expected satisfaction is maximized.

In more detail, the probability of a module being cancelled is estimated using a *softmax* function, which is commonly used in discrete choice theory to make decisions in the case of incomplete information [19]. The probability that a module  $i$  is going to be cancelled in the future is given by:

$$ProbCancel_i(\vec{b}) = \frac{e^{\frac{-(cp_i+b_i)}{\sum_{j=1}^m (cp_j+b_j)} \beta}}{\sum_{k=1}^m e^{\frac{-(cp_k+b_k)}{\sum_{j=1}^m (cp_j+b_j)} \beta}} \quad (3)$$

Where  $cp_i$  is the *cumulative* number of points which have so far been allocated to module  $i$ , and  $b_i$  is the number of points that the student agent is planning to allocate

to module  $i$  in the current voting round, and  $\vec{b}$  is the vector of points to be allocated. Furthermore,  $\beta$  is constant which weights the importance of the current point allocation (for example, if  $\beta = 0$ , then each module is equally likely to be cancelled, but as  $\beta \rightarrow \infty$ , the module with the lowest total number of points will be cancelled with probability 1, and all other modules will be cancelled with probability 0).

We can use this probability to calculate the *expected satisfaction*,  $ES$ , of the student. The expected satisfaction is given by:

$$ES(\vec{b}) = \sum_{i=1}^m (1 - ProbCancel_i(\vec{b})) \cdot v_i \quad (4)$$

The next step is then to find the allocation that maximises this expected utility. We estimate this using a search algorithm based on *random sampling*:

1. We randomly generate an allocation vector  $\vec{b}$  subject to the constraint that the total number of points is equal to the maximum number of points that we would like to spend in the current round.
2. The student agent calculates the expected satisfaction.
3. If the current solution has a higher expected satisfaction than any previous solution, then keep the solution. Otherwise, discard it.
4. This process is repeated for 1000 times and the solution with the highest expected utility is kept.

In our experiments, the number of points allocated in any round was 50%, except in the last voting round where we allocate all remaining points. In the first round (where there is no existing allocation of points) we use the proportional strategy (but only use 50% of the available points).

#### 4.4 Comparing Agent Strategies to the Optimal

In the first part of our experiment we compared the agent strategies explained above to the optimal case in which the University Agent has access to all of the preference information and makes the decisions centrally. Our objective is to discover if the decentralized agent solutions can match a centralized approach

In each of the experiments that follow, each scenario was run 30 times with different preferences. Thus, the results shown are the average results over these runs.

Figure 1 shows the results for cases 1, 2 and 3 respectively. Here, the y axis shows the *percentage of student satisfaction*. This is calculated by the total satisfaction of the running modules, as a percentage of the total satisfaction if all the modules would be running. Furthermore, on the x axis we vary the total number of running modules (while keeping the other parameters in the scenarios fixed). The graphs show the differences in the satisfaction of the agents using different strategies and also compares this with the satisfaction of the optimal solution.

These results show that the outcome of the proportional strategy is almost identical to the optimal strategy (although this is not visible in the figure, there is some difference but this is not statistically significant), and the intelligent strategy does slightly less well but is still very close to optimal. On the other hand, we see that the

equal share strategy does significantly worse. This suggests that a decentralized solution using voting results in high quality solutions that are comparable to optimal.

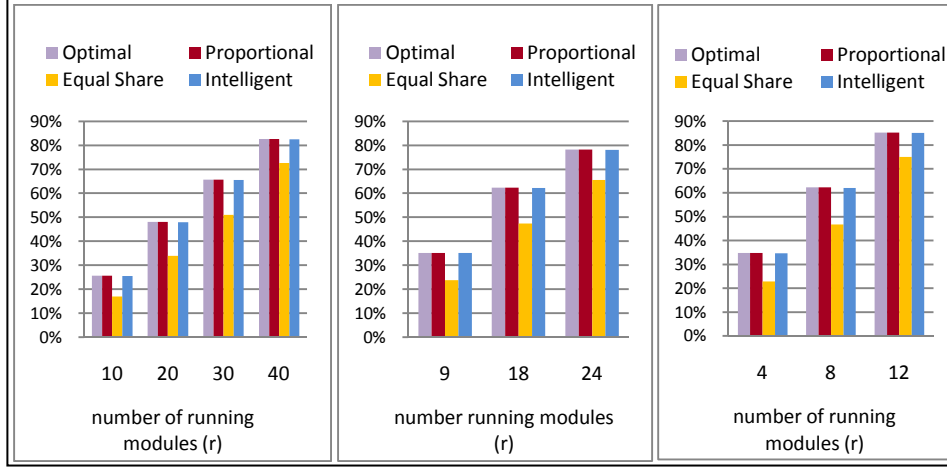


Figure. 1. Results for Case 1 (left), Case 2 (centre), Case 3 (right)

#### 4.4 Comparing Agent Strategies to Each Other

In the next set of experiments we compare the case where a proportion of the students use one strategy, and the remainder of the students uses another strategy. This allows us to see what would happen if students used a mixture of strategies to choice their options. We might expect some strategies to work better than others.

If all the students had random preferences then the agents would cancel each other out regardless of their strategy choice, so in order to show the effect of strategy we biased the preferences in such a way that the students using the same strategy are also likely to have similar preferences. In this way each group of students with the same strategy is pulling in the same direction. We can then measure the comparative success of each strategy by looking at the satisfaction within each group. We can also compare the power of the strategies by varying the size of the group.

In the results that follow, the y axis shows the percentage of satisfaction for each group of agents using a particular strategy. The x axis shows the proportion of students using a particular strategy. For example, in Figure 2, 90-10 means that 90 students use the proportional strategy, and 10 students use the equal share strategy.

The results in Figures 2 show that the intelligent and proportional strategies are both significantly better than the equal share, irrespective of the proportion of students that use this strategy. On average, the improvement is around 8% compared to the equal share strategy. The results shown are for case 1 (see Table 1) with 40 students, but the results for other cases are very similar and not shown to avoid repetition.

Figures 3 show the results of comparing the intelligent strategy and the proportional strategy for the 3 different cases. The result show that, as the number of students allocated to a particular strategy increases, the student satisfaction for these students also increases. However, this is mainly because of the bias that has been

introduced; since students with the same strategy have similar preferences, when more students have these preferences they have greater voting power since they act as a group. The difference in the effectiveness of the strategies can be seen by comparing the number of students needed in a group for it to become the most successful.

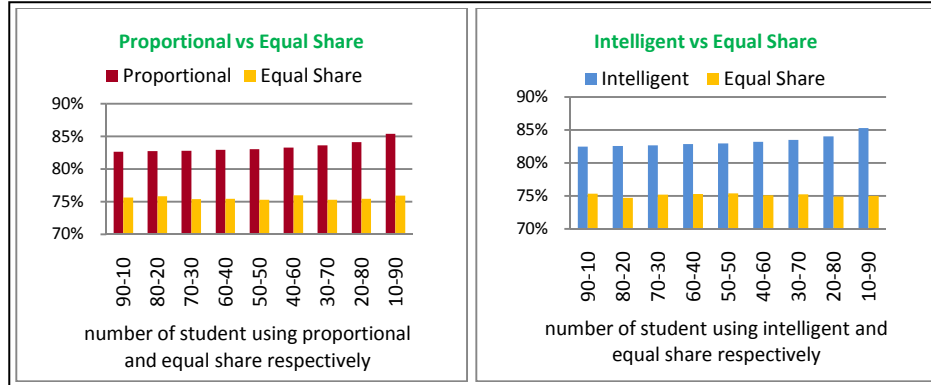


Figure 2. Left: Proportional vs. Equal Share, Right: Intelligent vs. Equal Share

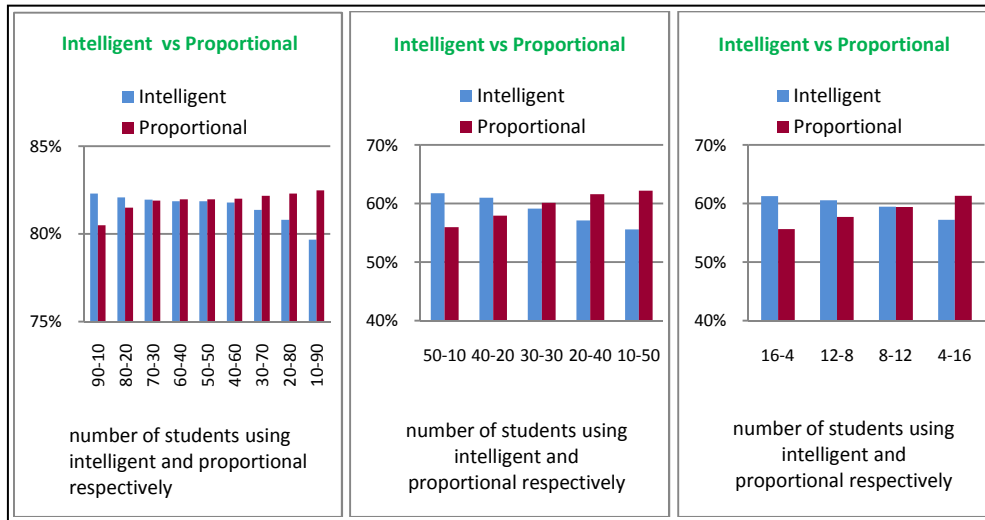


Figure 3. Intelligent vs. Proportional. Case 1 (left), Case 2 (centre), Case 3 (right)

Comparing the intelligent and proportional strategies, it can be seen that there is not much difference between them. Although in Case 3 the intelligent strategy slightly outperforms the proportional strategy (given the same number of students are using that strategy), in the other two cases, the proportional strategy outperforms the intelligent strategy. We have also tried to vary the parameters of the intelligent strategy (such as the beta parameter), but the results do not change significantly. This suggests that a learner who takes a more advanced intelligent strategy cannot easily exploit the system.

## 5 Discussion

Through the Scenarios described in Section 3 and the experimental simulation in Section 4 we have explored how multi-agent systems could be used for TEL. While others have focused on agents as containers of student information and user interface adaptation, we have focused on their ability to act autonomously in a system and to negotiate with one another to reach an agreed outcome.

Our work shows that when used in this way agents provide two key advantages:

- **Ultra-personalization** – In a multi-agent system autonomous agents allow for unprecedented levels of personalization. Not only can students have preferences about any given learning scenario, but by selecting different strategies they can change the way in which their agents negotiate. For example, in the module selection scenario students have different preferences and in our experiment can choose different strategies for negotiating based on those preferences. But it would also be possible to introduce other student agents that had a completely different basis for choosing modules (for example, based on the choices of their friends, or on the requirements of some qualification framework). Multi-agent systems provide the necessary level of abstraction for the tailoring of every aspect of the negotiation, including the basis for making choices (e.g. preferences or some other criteria), the individual's personal data (e.g. the preferences themselves), and the algorithm that uses that data to negotiate (e.g. how to vote according to those preferences).
- **Decentralization** – In Section 3 we argued that agent systems could provide decentralized solutions to a number of key TEL problems. In Section 4 we have demonstrated that not only is this possible, but that if students choose sensible strategies the results tend towards an optimal solution (calculated as the result of a centralized approach).

However, we also believe that our work highlights potential concerns:

- **Equitability** – In situations of ultra-personalization it is very difficult to guarantee that all students will have the same potential for satisfaction. This is because, although the agents are handled equally, the system relies on the agents themselves making sensible choices and selections. Power and responsibility are both transferred from a central authority (the institution) to the individual agents (the students). If an agent makes irrational choices, or chooses a bad strategy, then their student will be disadvantaged when compared to others. In Section 4 we demonstrated this by showing how a foolish equal share strategy penalized students who acted in that way, however we also showed how a well-designed protocol had made it difficult for a more intelligent (or intentionally subversive strategy) to gain advantage over a sensible strategy.

- **Scrutability** – in decentralized systems it can be very difficult for any individual in the system to see and understand the big picture, making accountability and transparency difficult. Although in our experiment the University Agent was in a position to see the voting behavior of all the students, it is possible to imagine situations where no single agent understands the sequence of events that lead the system as a whole reaching a decision (for example, if multiple institutions had been involved in our scenario then no one of them would have seen all the voting behavior). Institutions in particular may find it difficult to engage in a system where they cannot fully account for the outcome.

## 6 Conclusions and Future Work

The main aim of this work was to investigate how the autonomy and negotiation aspects of multi-agent technology might impact on the domain of Technology Enhanced Learning. Using three scenarios we have shown how common multi-agent solutions (voting, coalition formation and auction systems) map to problems in TEL. We argue that multi-agent technologies could allow genuine *decentralization* and *ultra-personalization* allowing these scenarios to be extended to include types of personal and informal learning that are difficult to support with today's systems.

As an illustration we took the first of these scenarios (module selection) and presented a multi-agent simulation that uses a suitable voting protocol to support module selection. Using our simulation we have been able to show that a decentralized agent approach not only works, but that with reasonable agent strategies it approximates an optimal centralized solution. We have also been able to show how different agent strategies compare to one another, revealing that with this particular protocol simplistic strategies are penalized, but that it is difficult to use intelligent (subversive) strategies to significantly improve on a naïve sensible strategy.

Based on the scenarios and experiment we believe that agent systems have a great deal of potential for TEL, but that their use raises concerns about the *equitability* of results (as agents become responsible for their own performance) and the *scrutability* of the process (as no single agent understands the system as a whole). Depending on the context this may impact on the acceptability to stakeholders of using an agent-system in a given scenario.

In our future work we intend to explore more sophisticated versions of the module choice scenario, where agents use different selection criteria and where there are multiple institutions vying for student interests. Our aim is to help establish the characteristics of agent protocols that makes them either equitable or scrutable, and to investigate the feasibility of more ambitious scenarios.

## References

1. Sampson, D., Karagiannidis, C., and Kinshuk: 'Personalised Learning: Education, Technological and Standardisation Perspective.', *Interactive Education Multimedia*, 2002, 4, pp. 24-39
2. Bellifemine, F.L., Caire, G., and Greenwood, D.: 'Developing multi-agent systems with JADE' (John Wiley, 2007. 2007)
3. Kommers, P., and Aroyo, L.: 'Special Issue Preface - Intelligent Agents for Educational Computer-Aided Systems', *Journal of Interactive Learning Research*, 1999, 10, (3), pp. 235-242
4. Silveira, R., and Vicari, R.: 'Developing Distributed Intelligent Learning Environment with JADE — Java Agents for Distance Education Framework', In the *Proceedings of the 6th International Conference in Intelligent Tutoring System, ITS 2002*, 2002, pp. 105-118
5. Bariso, E.U.: 'Personalised eLearning in Further Education', in O'Donoghue, J. (Ed.): 'Technology-Supported Environments for Personalized Learning: Methods and Case Studies' (Information science reference, 2009), pp. 138-156
6. Conole, G.: 'Personalisation through Technology-Enhanced Learning', in O'Donoghue, J. (Ed.): 'Technology-Supported Environments for Personalized Learning: Methods and Case Studies' (Information science reference, 2009), pp. 30-44
7. Sampson, D.: 'Current and Future Research and Technology Developments in e-Learning'. *Proc. the 2nd International Conference on New Horizons in Industry and Education*, Milos Island, Greece 2001 pp. Pages
8. Coombs, P.H., and Ahmed, M.: 'Attacking rural poverty : how nonformal education can help' (Johns Hopkins University Press, 1974. 1974)
9. McGivney, V.: 'Informal learning in the community : a trigger for change and development' (National Institute of Adult Continuing Education, 1999. 1999)
10. Cross, K.P.: 'Adults as learners : increasing participation and facilitating learning' (Jossey-Bass Publishers, 1981. 1981)
11. Jennings, N.R., and Wooldridge, M.: 'Applications of intelligent agents': 'Agent technology: foundations, applications, and markets' (Springer-Verlag New York, Inc., 1998), pp. 3-28
12. Soh, L.-K., Jiang, H., and Ansorge, C.: 'Agent-based cooperative learning: a proof-of-concept experiment'. *Proc. Proceedings of the 35th SIGCSE technical symposium on Computer science education*, Norfolk, Virginia, USA 2004 pp. Pages
13. Wooldridge, M.J.: 'An introduction to multiagent systems' (Wiley ; Chichester, 2009, 2nd ed. edn. 2009)
14. Meo, P.D., Garro, A., Terracina, G., and Ursino, D.: 'Personalizing learning programs with X-Learn, an XML-based, "user-device" adaptive multi-agent system', *Information Sciences: an International Journal*, 2007, 177, (8), pp. 1729-1770
15. Shi, H., Shang, Y., and Chen, S.-S.: 'A multi-agent system for computer science education', In *Proceedings of the 5th Annual Conference on Innovation and Technology in Computer Science Education (ITiCSE 2000)*, 2000, 32, (3), pp. 1-4
16. Yang, C., Lin, H., and Lin, F.O.: 'Designing Multiagent-Based Education Systems for Navigation Training', *5th IEEE Int. Conf. on Cognitive Informatics (ICCI'06)*, 2006
17. Soh, L.-K., Jiang, H., and Ansorge, C.: 'Agent-based cooperative learning: a proof-of-concept experiment', *SIGCSE Bull.*, 2004, 36, (1), pp. 368-372
18. Beer, M., Inverno, M., Luck, M., Jennings, N.R., Preist, C., and Schroeder, M.: 'Negotiation in Multi-Agent Systems', *Knowledge Engineering Review*, 1998, 14, (3), pp. 285-289
19. Hensher, D.A., Rose, J.M., and Greene, W.H.: 'Applied choice analysis : a primer' (Cambridge University Press, 2005. 2005)