An interactive platform for IA Games

Arturo de Salabert, Fernando Díez, Ruth Cobos, Manuel Alfonseca
arturo.desalabert@uam.es , fernando.diez@uam.es , ruth.cobos@uam.es , manuel.alfonseca@uam.es
Universidad Autónoma de Madrid, Madrid
Spain

***Abstract**: This paper introduces two platforms for the experimentation in games strate- gies within the study of Adversarial Search, part of the practical curriculum in Artificial Intelligence. The first one is a platform for performing tournaments. AI students are asked to send automated players for a given game, which are confronted against each other and ranked in order of performance. The platform was successfully used during the last two academic years by over 200 students per year, and performed over 100,000 confrontations/year. The second is a platform for executing P2P games in real time be- tween remote players. It is oriented to the performing of individual matches and in- cludes social network characteristics. This experience shows that the potential of such initiatives is very promising, since they not only stimulate the interest of the students for experimentation and learning, but also create a level of engagement that exceeded our and the student's expectations.*

Keywords: Adversarial search. Artificial Intelligence education. Action learning. P2P. Social networks. Self-learning. Learning virtualization. e-Learning.

Introduction

This paper reports a real educational experience in Artificial Intelligence that may be of interest for scholars in engineering education. It was inspired by ideas on active learning (Dilworth, 1998; Dunn, 2002) with the purpose of increasing the student's interest in our subject. There is no consensus about the influence of competition on motivation, but multiple experiences show that the influence can be positive (Marra & Wheeler, 2000; Cantador & Conde, 2010). This work takes place within the con- text of the European Higher Education Area (EHEA), which provides engineering educators the possi- bility to interact with and to provide continuous assessments to students (Pérez-Martinez et al., 2009; Crosier et al., 2006).

Artificial Intelligence (AI) has a rare characteristic that sets it apart from most disciplines in Science and Engineering: its objective is a moving target and its definition an ongoing one. Most of the prob- lems it deals with are open problems, where only incomplete or partial solutions exist. Certainly in AI there are well established concepts and techniques that could be taught using any other educational paradigm, but the study of Adversarial Search, where two or more contenders are competing to out- smart the other within the rules of a game, seems a good candidate for discovery learning. The subject to learn is the use of heuristics to estimate the goodness of a particular game movement considering the situation of the board and the options of the contrary. This is an open problem because no ultimate solution is known and because the goodness of any move is limited by the ability of the contrary. Therefore, the process of individual discovery seems to be as important as the techniques to be taught. In order to provide a realistic context where students can test their players' goodness and which could motivate them to work actively in this task, a tournament platform was implemented.

The next section introduces the objectives of this experience. The methods section describes the im- plemented platform and the data gathering process. The results section shows quantitative measures for the two years in which the platform has been used and presents our interpretation. The conclusion resumes the findings and their relevance. The last section describes our current activities including a P2P platform and our plans to use it in future work.

Research Objectives

The exploratory study described in this paper had three main objectives: 1. Stimulating the learning process; 2. Stimulating original thinking; 3. Understanding how time pressure affects both.

We can evaluate the first by measuring the amount of work produced by the students and comparing it with the previous situation. We can evaluate the second by confronting the algorithms produced for the tournament and the algorithms produced in previous years against some reference. It would be nice, but unfortunately we cannot confront directly the algorithms produced for the tournament in one year with those produced in previous years, because the games they are playing are not exactly the same. Indeed, in order to avoid plagiarism between years (or from external sources), we take two pre- cautions: firstly, the game we propose is a modified version of the game known in real life, so that comments, strategies or even the code available on the web may suggest ideas for successful players, but cannot be simply copy-pasted. Secondly, we modify the game from one year to the next in such way as to guarantee that the strategies used by successful players from previous years cannot be used so successfully as to represent a significant advantage in the current year. Finally, we can evaluate the effect of time pressure by observing the performance of the course under tighter time constraints. The time available for the competition was reduced to analyze this point.

Methodology

Our exploratory study took the form of a longitudinal case study conducted in the environment of an actual university class during two academic years. Students were requested to submit code, which was parsed and evaluated formally. If this evaluation was successful, the code was evaluated functionally by confronting it to its peers. Since the whole process is computerised, each action taken by the stu- dents or their players was recorded: timestamps of deliveries, identity of the opponents, result of each game, number of moves needed for game completion, primitives used within the code, errors & warn- ings, etc. As a result of each game, points were given or taken from the overall score of each player and the positions of the players in the ranking oscillate. In general, better players have an ascending trend and worse players a descending one. After completing the semester, students were asked to fill a questionnaire about their experiences with the tournament.

Participants

About two hundred university students participated each year in this study, grouped in teams of two, i.e. 100 teams/year. This exercise was part of their AI course in their 3rd or 4th year of studies. The total time available for the exercise was about 5 weeks. The work required to solve the proposed problem was estimated in 3-6 hours, but no limits were set nor indications were given in regard to the dedica- tion expected, although both, the data log and the questionnaire show that most students dedicated much more time to this task.

The Tournament Platform

A game is introduced and the students are asked to practice until they get familiar with its rules and basic strategies. Four automatic players (Lisp code) are provided with two goals: serve as counter players for the students and as examples of possible strategies. The students can play either by them- selves (manual input) or by writing code (automatic input). One of the four automatic players always plays randomly, so its games are unpredictable. The other three players are named bad, average and good, corresponding to their respective proficiency. Soon the students learn that playing by themselves they are only able to win systematically against the random player, and often against the bad player, but rarely against the average player. However, they can produce code that equals the average player. Defeating systematically the average player is a little more difficult and defeating the good player is much more difficult. The first objective of the exercise is quickly achieved: provided they work within a given framework that is part of the teaching, students can easily produce AI code able to perform a task that has proven very difficult for a human.

The second objective of the exercise is to explore options and produce the better possible playing strategy. Since their own code is already playing beyond their human capabilities, the only way to test new ideas is to confront algorithms with each other. To facilitate the games in an orderly form and to provide an impartial ranking of the players, a tournament platform was implemented. The tournament platform is permanently available for the submission of new players. Periodically the platform will confront the batch of new players with the older ones. The results of these games are used to update a ranking of the players based on their score, which is published on the web.

The first requirement of the system was to guarantee the correct timing, authenticity and non- repudiability of the submissions. Timing is not important per se (students have weeks to present their players), but is used to limit the number of submissions per day, so as to make it more difficult to per- form the activity by trial-and-error. Authenticity is needed to insure that each player really belongs to a certain student and not to a rival. Non-repudiability is required in case of student complaints, since grades are assigned according to their position in the final ranking. This has been achieved by coding the file names and using private passwords.

The tournament management system should allow any type of game. This is achieved by decoupling the game mechanics from the functions of submission management, opponent's selection, scoring, ranking updating and publishing. The implemented management system (majordomo) satisfies all these requirements and, in addition, can manage several different games simultaneously.

Since the student submissions (players) are computer code that must run on the same server where the tournament manager is running, system's security and stability is a major concern. A syntactical ana- lyser or parser was independently developed to inspect the code before its execution. The parser de- tects format errors as well as the use of functions that voluntarily or involuntarily could damage or jeopardise the system.

A major requirement to keep alive the interest on an on-line tournament is the ranking refresh rate, i.e. to publish results timely, shortly after the submissions. On the one hand, the only way to guarantee that each player occupies its true place in the ranking is to perform a round robin competition (Harary & Moser, 1966). However, there is a continuous flux of players entering and abandoning the competi- tion, so performing a round-robin competition before every update of the ranking is either not practical or impossible, due to the amount of computing involved. Computation grows geometrically as N2 with the number of players N, and therefore a round robin competition is not scalable. To avoid both prob- lems we implemented several scoring systems inspired on Elo's rating for chess, which were tested along thousands of simulated games. These simulations showed that given a pre-existent ranking pro- duced by a round robin competition amongst a small group of players, say 20, we could place any new player in its correct position (i.e. the position it would have reached in a round robin competition) with an average error of 5% after no more than 40 games against selected players. This situation keeps quite stable for a few hundreds of new entrants, but it shows a small but accumulative degradation in accuracy, especially for each old player that abandons the game or when there is a large difference between old and new players in number of games played. To overcome this degradation, the tourna- ment refreshes the ranking by performing a round robin tournament every several days, after midnight. In some cases this would produce significant changes in the position of a player, which was often fol- lowed by complaints (students followed very closely their positions in the ranking). The technique used exceeds the scope of this article, but we can affirm that it solves the problem in linear time and therefore is scalable to a large number of players, at least for a period of time. Every 60 minutes, the new entrants are confronted to the old ones and assigned a position in the ranking. This process takes only a few minutes, so it could be repeated much more often. It was chosen not to do so for the rea- sons expressed in the next paragraph, despite the sour complaints of some students.

Another issue related to the previous one is the number of submissions. In our experience, most stu- dents tend to delegate a great part of their intellectual tasks to the computer, disregarding careful plan- ning and reflexion. The result of this continual trial-and-error programming style is often code full of ad-hoc solutions, not well-thought and lacking structure. In an attempt to correct this tendency, the number of players that the students can submit to the platform is limited to 3 per day, and must wait up to one hour to see the results. This is not based on technical limitations but to reinforce the reflexion and careful planning of the strategies and to limit trial-and-error. This restriction also invites the students to start their experiments early: for each day that they don't play, they lose game options. Late- comers should have a difficult time catching up.

Data Analysis

The collected data have been aggregated using string processing functions (regular expressions) and processed using basic statistical spread sheet functions. To evaluate the work and results, datasets have been produced for deliveries per student and deliveries per day, which, combined with the points obtained, are an indication of the results achieved. In addition, we have a measure of the effectiveness of the players: the so called good player is not using any strategy related to the particular game; instead, it is just using the same plain strategy used by the average player, but with one more depth level. This serves two purposes: it is a rough sparring partner for the student players and a bar level that we can use across years. As mentioned earlier, in general we cannot confront players from different years.

Although the problem in its current form has been solved by the students since 2007, the tournament has only been available since 2009. We can compare however all four years by using the bar set by the good player. Unfortunately this is the only objective measure of quality, but the results are quite clear nevertheless.

Results and Discussion

The game strategy problem was presented to the students at the beginning of December and the clos- ing time for submissions was early January (15-1-2010 and 10-1-2011). One can appreciate the simi- larity between the curves and their adaptation to the major festivities, with clear drops in number of deliveries in Christmas day and January 1st. Likewise, almost nobody started working on the problem before December 20th (Fig.1) (begin of the holiday season). The real start is clearly the day after Christmas. The submissions per year remained constant at 900, averaging 9 per team (Fig.2).

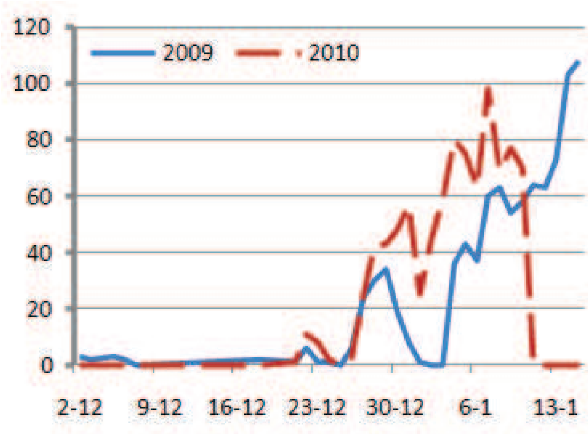

Figure 1: Deliveries per day

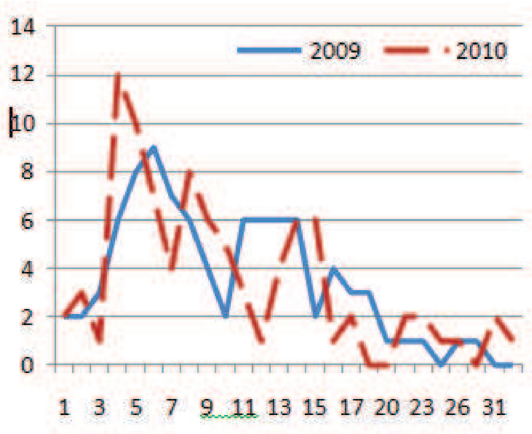

Figure 2: Deliveries per team

Despite the 5 days shorter deadline in 2010 the students did the same amount of work (900 players) in 20% less time, while delivering comparable quality: 5% of the players passing the high bar and 80% the low bar. Furthermore, the response pulse in 2010 shows better planning. The peak of deliveries happens a few days before the deadline, instead of the last minute as in 2009. Figures 3 and 4 confirm the same observation from a different perspective. Students in 2010 reach similar quality with fewer submissions in fewer days, i.e. their efficiency is quite higher. A questionnaire was sent to the most outstanding cases, which confirmed that the best performers did a lot of analysis before submitting every player. Finally, by plotting the starting date of the deliveries against the final points obtained we observe in both years a clear correlation between starting early and getting a higher score. In addition, the plot (not included) shows that in 2010 students started earlier although they had less time.

Comparing the two years with tournament with the situation before, where students had to produce just one player, which was graded not by comparing it to its peers but just as it performed against sev- eral bar levels, two things become obvious. Firstly, the quality of the players is much better, since no player produced earlier had ever passed the bar, while in the two years of the tournament 5% of the players did pass the bar each year. Secondly, and this is the cause of the previous result, the time in- vested in the exercise has increased from the estimated 6 hours before the tournament to an average of

20 hours or 10 days (11,5 in 2009 and 8 in 2010) since the tournament was introduced. Considering that the value of the player for the final grade has not changed, we must conclude that the increase in dedication was stimulated by the tournament.

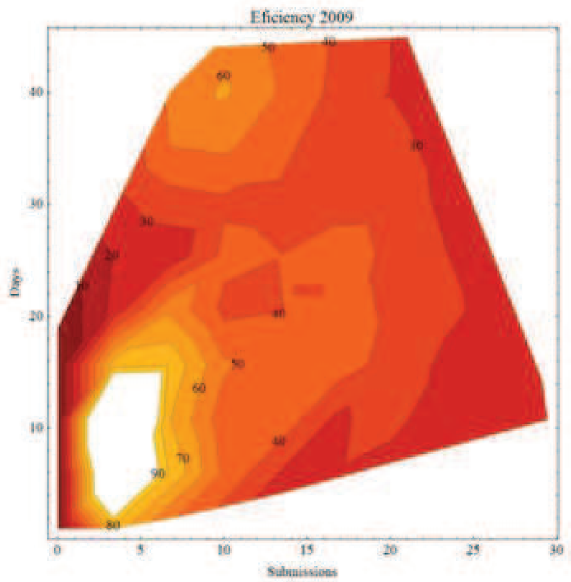

Figure 3: Efficiency 2009

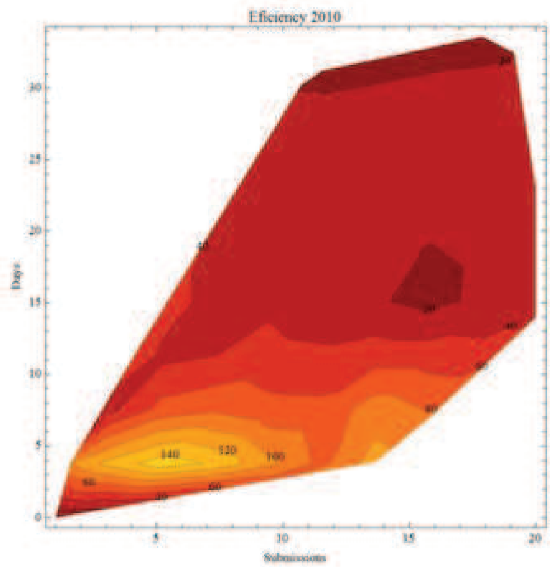

Figure 4: Efficiency 2010

Conclusions

This experience has created enormous interest among the students, some of which dedicated to these activities much more time and effort than both, faculty and students, expected. A result of the longer time dedicated to the activity is the quality of the results, clearly superior to previous years. The com- bination of a competition with a tournament platform creates a need and provides a medium to moti- vate work and to stimulate experimentation, resulting in better and more original results. The success of the first platform created requirements for a second platform (P2P platform, described below), which was developed partly based on student demand.

However, while the interest in competition is supported by 75% of the students, 20% of the students manifest strong reluctance to have their grades linked to such a competition. Both quantitative and qualitative data show that competition is a strong motivator for learning, but should not be too strongly coupled with grades.

Finally, we want to stress the fact that a moderate time pressure resulted to be a stimulus for careful planning and reflexion. This relevant circumstance is consistently confirmed by results from three different observations: submissions per day, efficiency and starting date.

Current and Future Plans

The tournament model has been well accepted by the students and has stimulated their interest in the subject. The utility of the tool from an educational perspective has been very positive, but it has a lim- ited capability as a research tool, since the information that can be obtained is reduced to several files of code per day. We can extract some approximated guess about the amount of work invested by the students and about the quality of their learning, as we have done in this article, but to fully understand the students learning process, their trials and perhaps the origins and flows of innovation, we need much richer information.

It was a suggestion of some students to have a system to challenge each other in a selective fashion. They asked to be able to confront their players against a particular opponent instead of just waiting for the ranking to give them an overall score. At the same time, we were considering to evolve the system to allow multiplayer and collaborative games. Therefore, during the current course we have imple- mented a new platform that fulfils both requirements and opens a new way of looking at the problem. While the tournament will remain the official ranking, the P2P platform will play the role of a chess club, where members can informally train themselves before going to the tournament.

The new platform is technically and functionally independent from the first, and it uses P2P tech- niques to allow individual games to take place in real time between two or more players from remote locations. Its most relevant characteristics are:

- It provides students with a tool to improve their players by selective confrontations, while protect- ing their code by exchanging only moves, not algorithms.
- It allows multiplayer and collaborative games, which enable many new types of strategies.
- It allows manual and automatic games and the possibility of focusing on specific moves or prob- lems.
- It has social network characteristics, where gaming plays the role of relationship. It is scalable to very large groups.
- It provides much richer information for education research purposes.

The P2P platform has been successfully tested during the academic year 2010/11 by a group of enthu- siastic students and will be used by all the students next year. This will allow us to gather much more data, which we expect to use in future research.

References

Cantador, I. & Conde, J.M. (2010). A simple e-learning system based on classroom competition. *Proceedings of the 5th European conference on Technology enhanced learning conference on Sustaining TEL: from innova- tion to learning and practice.* (vol 6383, pp 488-493). Springer Berlin, Heidelberg

Crosier, D.Purser, L. & Smidt H. (2006). *Trends V: Universities Shaping the European Higher Education Area.* EUA Report. Accessed at http://www.sowi-due.de/uploads/108.pdf on 2 Jun 2011.

Dilworth, R.L. (1998). Action Learning in a Nutshell*. Performance Improvement Quarterly 11, no. 1 (1998a):* 28-43.

Dunn, L. (2002). Theories of learning. *Learning and Teaching Briefing Papers Series. Oxford Centre for Staff and Learning Development OCSLD.*

Harary, F., Moser, L. Theory of Round Robin Tournaments*. The American Mathematical Monthly* 73(3), Mar., 1966, pp. 231-246.

Marra, R., & Wheeler, T. (2000). The impact of an authentic, student-centered engineering project on student- motivation. *Proceedings 30th Annual ASEE/IEEE Frontiers in Education Conference*, October 18-21, 2000, Kansas City, MO.

Pérez-Martínez, J.E., García-García. M.J., Hernández Perdomo, W., Villamide-Díaz M.J. (2009). Analysis of the results of the continuous assessment in the adaptation of the Universidad Politécnica de Madrid to the Euro- pean Higher Education Area. *Proceedings of the Research in Engineering Education Symposium*. Palm Cove, QLD, Australia.

Acknowledgements

The tournament platform was partially funded by a grant of the European Convergence Office in 2009 and used by all the students of 2009 and 2010. The P2P platform was partially funded by a grant of the European Convergence Office in 2010 and beta tested by a group of voluntary students. This re- search was also partially funded by the Spanish National Plan of R+D, project number TIN2008-02081/TIN and by the CAM (Autonomous Community of Madrid), project number S2009/TIC-1650.